Financial Economics, Return Predictability, and Market Efficiency*

Stijn Van Nieuwerburgh
New York University Stern School of Business and NBER

Ralph S.J. Koijen
Tilburg University

May 7, 2007

Contents

1 Glossary **2**

2 Definition of the Subject and Its Importance **2**

3 Introduction **3**

4 Motivating Predictive Regressions **3**

5 Structural Model **8**

6 Future Directions **12**

---

*Koijen: Department of Finance, CentER and Netspar, Tilburg University, Tilburg, the Netherlands, 5000 LE; r.s.j.koijen@tilburguniversity.nl; Tel: 31-13-4663238; `http://center.uvt.nl/phd_stud/koijen/`. Van Nieuwerburgh: Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; svnieuwe@stern.nyu.edu; Tel: (212) 998-0673; `http://www.stern.nyu.edu/~svnieuwe`. We thank John Campbell for an enlightening presentation at the Copenhagen Business School Return Predictability Conference.

1 Glossary

Stock Return The stock return in this entry refers to the return on the portfolio of all stocks that are traded on the three largest equity markets in the US: the NYSE, NASDAQ, and AMEX. The return is measured as the price of the stock at the end of the year plus the dividends received during the year divided by the price at the beginning of the year. The return of each stock is weighted by its market capitalization when forming the portfolio. The source for the data is CRSP.

Dividend-Price Ratio and Dividend Yield The dividend-price ratio of a stock is the ratio of the dividends received during the year divided by the price of the stock at the *end* of the year. The dividend yield, instead, is the ratio of the dividends received during the year divided by the price of the stock at the *beginning* of the year. The stock return is the sum of the dividend yield and the capital gain yield, which measures the ratio of the end-of-year stock price to the beginning-of-year stock price

Predictability A stock return r_{t+1} is said to be predictable by some variable x_t if the expected return conditional on x_t, $E[r_{t+1} \mid x_t]$, is different from the unconditional expected return, $E[r_{t+1}]$. No predictability means that the best predictor of tomorrow's return is the constant, unconditional average return, i.e., $E[r_{t+1} \mid x_t] = E[r_{t+1}]$. When stock returns are unpredictable, stock prices are said to follow a random walk.

Market model The market model links the return on any asset i, r_{it} to the return on the market portfolio (r_t). Under joint normality of returns, it holds:

$$r_{it} = \alpha_i + \beta_i r_t + \varepsilon_{it}, \tag{1}$$

with $E[\varepsilon_{it}] = 0$ and $\mathrm{Var}[\varepsilon_{it}] = \sigma^2_{\varepsilon_i}$, see [12]. The typical assumption in the literature until the 1980s has been that $E[r]$ is constant.

2 Definition of the Subject and Its Importance

The efficient market hypothesis, due to [21], [22], and [23], states that financial markets are efficient with respect to a particular information set when prices aggregate all available information. Testing the efficient market hypothesis requires a "market model" which specifies how information is incorporated into asset prices. Efficiency of markets is then synonymous with the inability of investors to make economic, i.e., risk-adjusted, profits based on this information set ([36]). The question of market efficiency and return predictability is of tremendous importance for investors and academics alike. For investors, the presence of return predictability would lead to different optimal asset allocation rules. Failing to make portfolios conditional on this information may lead to substantial welfare losses. For academics, return predictability or the lack thereof has substantial implications for general equilibrium models that are able to accurately describe the risks and returns in financial markets.

3 Introduction

Until the 1980s, the standard market model assumed constant expected returns. The first empirical evidence, which showed evidence that returns were predictable to some extent, was therefore interpreted as a sign of market inefficiency ([54] and [56]). [25] proposed the alternative explanation of time-varying expected returns. This prompted the question of why aggregate stock market returns would be time varying in equilibrium. [23] provides a summary of this debate.

Recently developed general equilibrium models show that expected returns can indeed be time varying, even if markets are efficient. Time-variation in expected returns can result from time-varying risk aversion ([11]), long-run consumption risk ([5]), or time-variation in risk-sharing opportunities, captured by variation in housing collateral ([44]). Predictability of stock returns is now, by-and-large, interpreted as evidence of time-varying expected returns rather than market inefficiency.

4 Motivating Predictive Regressions

Define the gross return on an equity investment between period t and period $t+1$ as

$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t},$$

where P denotes the stock price and D denotes the dividend. [9] log-linearizes the definition of a return to obtain:

$$r_{t+1} = k + \Delta d_{t+1} + \rho dp_{t+1} - dp_t. \tag{2}$$

All lower-case letters denote variables in logs; d_t stands for dividends, p_t stands for the price, $dp_t \equiv d_t - p_t$ is the log dividend-price ratio, and r_t stands for the return. The constants k and $\rho = (1 + \exp(\overline{dp}))^{-1}$ are related to the long-run average log dividend-price ratio $\overline{dp}$. By iterating forward on equation (2) and by imposing a transversality condition (i.e., we rule out rational bubbles), one obtains

$$dp_t = \overline{dp} + E_t \sum_{j=1}^{\infty} \rho^{j-1} \big[(r_{t+j} - \overline{r}) - (\Delta d_{t+j} - \overline{d})\big]. \tag{3}$$

Since this equation holds both ex-post and ex-ante, an expectation operator can be added on the right-hand side. This equation is one of the central tenets of the return predictability literature, the so-called Campbell and Shiller ([13] and [14]) equation. It says that, as long as the expected returns and expected dividend growth are stationary, deviations of the dividend-price ratio (dp_t) from its long-term mean ($\bar{dp}$) ought to forecast either future returns, or future dividend growth rates, or both.

This accounting identity has motivated some of the earliest empirical work in return predictability, which regressed returns on the lagged dividend-price ratio, as in equation (4):

$$(r_{t+1} - \bar{r}) = \kappa_r(dp_t - \overline{dp}) + \tau^r_{t+1}, \tag{4}$$

$$(\Delta d_{t+1} - \overline{d}) = \kappa_d(dp_t - \overline{dp}) + \tau^d_{t+1}, \tag{5}$$

$$(dp_{t+1} - \overline{dp}) = \phi(dp_t - \overline{dp}) + \tau^{dp}_{t+1}, \tag{6}$$

where $\bar{r}$ is the long-run mean return and τ^r is a mean-zero innovation. The logic of (3) suggests that the dividend-price ratio could predict future dividend growth rates instead of, or in addition to, future returns. Testing for dividend growth predictability would lead one to estimate equation (5), where $\overline{d}$ denotes the long-run mean log dividend growth.

The empirical return predictability literature started out by estimating equation (4) with the dividend-price ratio on the right-hand side; see [53], [24], [13], [17], [29], [34], and [42], among others. It found evidence for return predictability, i.e., $\kappa_r > 0$. This finding was initially interpreted as evidence against the efficient market hypothesis.

Around the same time, [25] and [52] document a negative autocorrelation in long-horizon returns. Good past returns forecast bad future returns. [12] and [18] summarize the evidence based on long-horizon autocorrelations and variance ratios, and conclude that the statistical evidence in favor of mean reversion in long-horizon returns is weak, possibly due to small sample problems. This motivates [4] to use a large cross-section of countries and use a panel approach instead. They in turn document strong evidence in favor of mean-reversion of long-horizon returns with an estimated half-life of 3-3.5 years.

Second, other financial ratios, such as the earnings-price ratio or the book-to-market ratio, or macro-economic variables such as the consumption-wealth ratio, the labor income-to-consumption ratio, or the housing collateral ratio, as well as corporate decisions, and the cross-sectional price of risk have subsequently been shown to predict returns as well; see [38], [3], [39], [45], [43], [50], and [51], among others.

Third, long-horizon returns are typically found to be more predictable than one-period ahead returns. The coefficient $\kappa_r(H)$ in the H-period regression

$$\sum_{j=1}^{H} r_{t+j} = \kappa_r(H)\ dp_t + \tau^r_{t,t+H} \tag{7}$$

exceeds the coefficient κ_r in the one-period regression. This finding is interpreted as evidence for the fact that the time-varying component in expected returns is quite persistent.

Fourth, these studies conclude that growth rates of fundamentals, such as dividends or earnings, are much less forecastable than returns using financial ratios. This suggests that most of the variation of financial ratios is due to variation in expected returns.

Fifth, predictability of stock returns does not only arise for the US. Studies by [26], [10], [33], and [2] analyze a large cross-section of countries and find evidence in favor of predictability by financial ratios in some countries, even though the evidence is mixed. More robust results are documented for the predictive ability of term structure variables.

These conclusions regarding predictability of stock returns are controversial because the forecasting relationship of financial ratios and future stock returns exhibits three disconcerting statistical features. First, correct inference is problematic because financial ratios are extremely persistent. The empirical literature typically augments equation (4) with an auto-regressive specification for the predictor variable, as in equation (6), where $\bar{dp}$ is the long-run mean of the dividend-price ratio. The estimated autoregressive parameter ϕ is near unity and standard tests leave the possibility of a unit root open (i.e., $\phi = 1$). [46], [55], [2], [27], and [58] conclude that the statistical evidence of forecastability is weaker once tests are adjusted for high persistence. [1], [2], [16], [42], [57], and [20] derive asymptotic distributions for predictability coefficients under the assumption that the forecasting variable follows a local-to-unit root, yet stationary, process.

Second, financial ratios have poor out-of-sample forecasting power, as shown in [7], [31], and [32], but see [35] and [15] for different interpretations of the out-of-sample tests and evidence.

Third, the forecasting relationship of returns and financial ratios exhibits significant instability over time. Figure 1 shows that in rolling 30-year regressions of annual log CRSP value-weighted returns on lagged log dividend-price ratios, the ordinary least squares (OLS) regression coefficient varies between zero and 0.5 and the associated R^2 ranges from close to zero to 30% depending on the subsample.

[60] and [49] report evidence in favor of breaks in the OLS coefficient in the forecasting regression of returns on the lagged dividend-price ratio, while [41] report evidence for structural shifts in $\bar{dp}$. [47] use Bayesian methods to estimate structural breaks in the equity premium.

Figure 1: Parameter Instability in Return Predictability Coefficient

The figure plots estimation results for the equation $r_{t+1} - \bar{r} = \kappa_r(dp_t - \overline{dp}) + \tau^r_{t+1}$. It shows the estimates for κ_r using 30-year rolling windows. The dashed line in the left panels denote the point estimate plus or minus one standard deviation. The standard errors are asymptotic. The parameters $\bar{r}$ and $\overline{dp}$) are the sample means of log returns r and the log dividend-price ratio dp. The data are annual for 1927-2004.

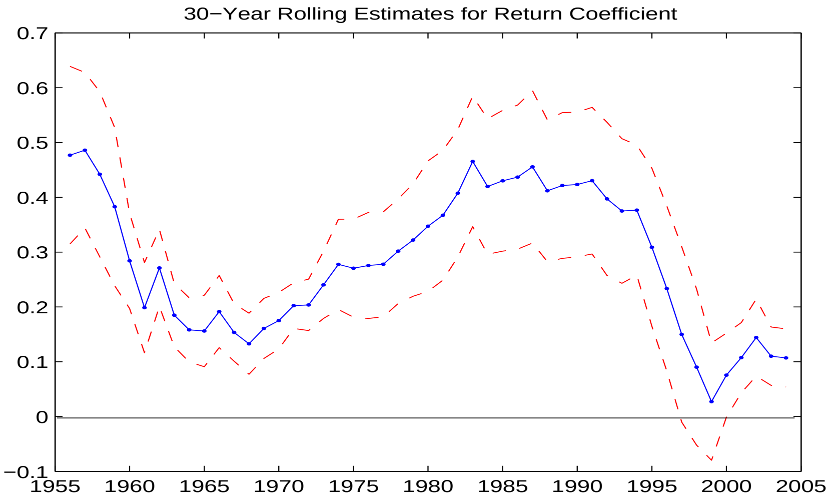

Empirical Evidence Revisited Table 1 reviews the empirical evidence using annual value-weighted CRSP log return, dividend growth, and dividend-price ratio data for 1927-2004. In Panel A, the system of equations (4) and (5) is estimated by GMM. The first row indicates that a higher dividend-price ratio leads to a higher return ($\kappa_r = .094$ in Column 2) and a higher dividend growth rate ($\kappa_d = .005$ in Column 1). The latter coefficient has the wrong sign, but the coefficient is statistically indistinguishable from zero. The asymptotic standard error on the estimate for κ_r is .046. The corresponding asymptotic p-value is 3.6% so that κ_r is statistically different from zero at conventional levels. In other words, the dividend-price ratio seems to predict stock returns, but not dividend growth. A similar result holds if returns in excess of a risk-free rate are used, or real returns instead of nominal returns.

[41] conduct an extensive Monte Carlo analysis to investigate the small-sample properties of estimates for κ_r and κ_d. Consistent with [55], the estimate for κ_r displays an upward small-sample bias. In addition, the standard error on κ_r is understated by the asymptotic standard error. As a result, one can no longer reject the null hypothesis that κ_r is zero. Based on this evidence, one is tempted to conclude that neither returns nor dividend growth are forecastable.

The second and third rows implement the suggestion of [41] to correct the long-run mean dividend-price ratio, $\bar{dp}$, for structural breaks. The data strongly suggest either one break in 1991, or two breaks in 1954 and 1994 in favor of either no breaks or three breaks. This break-adjusted dividend-price ratio is less persistent and less volatile. Its lower persistence alleviates the econometric issues mentioned above.

The second row of Table 1 uses the one-break adjusted dividend-price ratio as a regressor in the return and dividend growth predictability equations. The evidence in favor of return predictability is substantially strengthened. The point estimate for κ_r more than doubles to .235, and is highly significant. In the two-break case in the third row, the point estimate further doubles to 0.455. The small-sample bias in κ_r is negligible relative to the size of the coefficient. The R^2 of the return equation is 10% in the one-break case and even 23% in the two-break case. This compares to 3.8% in the no-break case. Furthermore, rolling regression estimates of κ_r indicate that it is much more stable over time when the break-adjusted *dp* series is used as a regressor. The dividend growth coefficient κ_d remains statistically indistinguishable from zero. This evidence strengthens the view that returns are predictable and dividend growth is not, and that these findings are not an artefact of statistical issues.

Table 1: Return and Dividend Growth Predictability in the Data

This table reports GMM estimates for the parameters $(\kappa_d, \kappa_r, \phi)$ and their asymptotic standard errors (in parentheses). The results in panel A are for the system with one-year ahead equations for dividend growth and returns ($H = 1$, $N = 0$). The results in panel B are for the system with one-year, three-year and five-year ahead equations for dividend growth and returns ($H = \{1, 3, 5\}$, $N = 2$). The first-stage GMM weighting matrix is the identity matrix. The asymptotic standard errors and p-values are computed using the Newey-West HAC procedure (second stage weighting matrix) with four lags in panel A and $H = 5$ lags in panel B. The last column denotes the present-value constraint violation of the univariate OLS slope estimators: $(1 - \rho\phi^{ols})^{-1}(\kappa_r^{ols} - \kappa_d^{ols})$. It is expressed in the same units as κ_d and κ_r. In panel B this number is the average violation of the three constraints, one constraint at each horizon. The dividend-price ratio in rows 1 and 4 is the unadjusted one. In rows 2 and 5, the dividend-price ratio is adjusted for one break in 1991, and in rows 3 and 6, it is the series adjusted for two breaks in 1954 and 1994. All estimation results are for the annual sample 1927-2004.

	κ_d	κ_r	ϕ	PV violation
Panel A: No Long-Horizon Moments $H = \{1\}$				
No Break	.005	.094	.945	-.046
	(.037)	(.046)	(.052)	
1 Break ('91)	.019	.235	.813	.004
	(.047)	(.055)	(.052)	
2 Breaks ('54, '94)	.124	.455	.694	-.001
	(.073)	(.079)	(.070)	
Panel B: Long-Horizon Moments $H = \{1, 3, 5\}$				
No Break	.021	.068	.990	.189
	(.018)	(.038)	(.032)	
1 Break ('91)	.012	.210	.834	.076
	(.019)	(.043)	(.042)	
2 Breaks ('54, '94)	.080	.409	.697	.100
	(.065)	(.078)	(.060)	

5 Structural Model

What are researchers estimating when they run the return predictability regression (4)? How are the return and dividend growth predictability regressions in (4) and (5) related? To answer these important questions, we set up a simple structural model with time-varying expected returns and expected dividend growth rates. This structural model has the system of equations (4)-(6) as its reduced-form. The main purpose of this model is to show that (i) the dividend-price ratio is a contaminated predictor of returns and dividend growth rates, (ii) that the parameters in (4)-(6) have to satisfy a cross-equation restriction, which we call the *present-value constraint*, and (iii) this restriction enables decomposing the dividend-price ratio into expected returns and expected dividend growth. Similar models can be derived for financial ratios other than the dividend-price ratio (e.g., [61]). [6] show how stock returns and book-to-market ratios are related in a general equilibrium model.

A Present-Value Model We assume that expected dividend growth, z, and expected returns, x, follow an AR(1) process with autoregressive coefficient ϕ:

$$\Delta d_{t+1} - \overline{d} = z_t + \epsilon_{t+1}, \qquad z_{t+1} = \phi z_t + \zeta_{t+1}, \tag{8}$$

$$r_{t+1} - \bar{r} = x_t + \eta_{t+1}, \qquad x_{t+1} = \phi x_t + \xi_{t+1}. \tag{9}$$

The model has three fundamental shocks: an innovation in unexpected dividends ϵ_{t+1}, an innovation in expected dividends ζ_{t+1}, and an innovation in expected returns ξ_{t+1}. We assume that all three errors are serially uncorrelated and have zero cross-covariance at all leads and lags: $\text{Cov}(\epsilon_{t+1}, \zeta_{t+j}) = 0$, $\forall j \neq 1$, $\text{Cov}(\xi_{t+1}, \zeta_{t+j}) = 0$, $\forall j \neq 1$, and $\text{Cov}(\epsilon_{t+1}, \xi_{t+j}) = 0$, $\forall j$, except for a contemporaneous correlation between expected return and expected dividend growth innovations $\text{Cov}(\zeta_t, \xi_t) = \chi$, and a correlation between expected and unexpected dividend growth innovations $\text{Cov}(\zeta_t, \epsilon_t) = \lambda$. We discuss innovations to unexpected returns η below.

In steady-state, the log dividend-price ratio is a function of the long-run mean return and dividend growth rate $\overline{dp} = \log\left(\frac{\bar{r}-\bar{d}}{1+\bar{d}}\right)$. The log dividend-price ratio in (3) can then be written as:

$$dp_t - \overline{dp} = \frac{x_t - z_t}{1 - \rho\phi}. \tag{10}$$

The dividend-price ratio is the difference of two AR(1) processes with the same root ϕ, which is again an AR(1) process. I.e., we recover equation (6).

The return decomposition in [9] implies that the innovation to unexpected returns follows from the three fundamental shocks (i.e., combine (2) with (8)-(10)):

$$\eta_{t+1} = \frac{-\rho}{1-\rho\phi}\xi_{t+1} + \frac{\rho}{1-\rho\phi}\zeta_{t+1} + \epsilon_{t+1}. \tag{11}$$

Since both ρ and ϕ are positive and $\rho\phi < 1$, a positive shock to expected returns leads, ceteris paribus, to a negative contemporaneous return. Likewise, a shock to expected or unexpected dividend growth induces a positive contemporaneous return.

Contaminated Predictor The first main insight from the structural model is that the demeaned dividend-price ratio in (10) is an imperfect forecaster of both returns and dividend growth. Returns are predicted by x_t (see equation (9)), but variation in the dividend-price ratio is not only due to variation in x, but also in expected dividend growth z_t. The same argument applies to dividend growth which is predicted by z_t (see equation 8). This implies that the regressions in the reduced-form model in (4) and (5) suffer from an errors-in-variables problem ([24], [37], and [30]).

To illustrate the bias, we can link the regression coefficients κ_r and κ_d explicitly to the underlying structural parameters:

$$\kappa_r = \frac{\text{Cov}(r_{t+1}, dp_t)}{\text{Var}(dp_t)} = \frac{(1-\rho\phi)(\sigma_\xi^2 - \chi)}{\sigma_\xi^2 + \sigma_\zeta^2 - 2\chi}, \tag{12}$$

$$\kappa_d = \frac{\text{Cov}(\Delta d_{t+1}, dp_t)}{\text{Var}(dp_t)} = \frac{-(1-\rho\phi)(\sigma_\zeta^2 - \chi)}{\sigma_\xi^2 + \sigma_\zeta^2 - 2\chi}. \tag{13}$$

If growth rates are constant, i.e., $\chi = 0$ and $\sigma_\zeta = 0$, then the dividend-price ratio is a perfect predictor of returns and $\kappa_r^\star = 1 - \rho\phi$. In all other cases, there is a bias in the return predictability coefficient:

$$\kappa_r^\star - \kappa_r = \frac{(1-\rho\phi)(\sigma_\zeta^2 - \chi)}{\sigma_\xi^2 + \sigma_\zeta^2 - 2\chi}. \tag{14}$$

[24] argue that κ_r is downward biased ($\kappa_r^\star - \kappa_r > 0$). In fact, the structural parameters that are implied by the reduced-form model parameters indicate an upward bias. This occurs because the correlation between expected dividend growth and expected returns is sufficiently high.

A similar argument applies to κ_d. [40] construct a variable based on the co-integrating relationship between consumption, dividends from asset wealth, and dividends from human wealth. They show that this variable has strong predictive power for dividend growth, and they show that expected returns and expected growth rates are highly positively correlated. This implies that expected growth rates and expected returns have an offsetting effect on financial ratios, which makes it hard to reliably detect time-varying growth rates using such financial ratios.

Present-Value Constraint The second main insight from the structural model is that there is a cross-equation restriction on the three innovations $\tau = (\tau^d, \tau^r, \tau^{dp})$ of the reduced-form model (4)-(6). Expressed in terms of the structural parameters, these innovations are:

$$\tau^d_{t+1} = \epsilon_{t+1} + x_t \left(\frac{-\kappa_d}{1-\rho\phi} \right) + z_t \left(\frac{\kappa_r}{1-\rho\phi} \right) \tag{15}$$

$$\tau^r_{t+1} = \epsilon_{t+1} + x_t \left(\frac{-\kappa_d}{1-\rho\phi} \right) + z_t \left(\frac{\kappa_r}{1-\rho\phi} \right) - \rho \left(\frac{\xi_{t+1} - \zeta_{t+1}}{1-\rho\phi} \right) \tag{16}$$

$$\tau^{dp}_{t+1} = \frac{\xi_{t+1} - \zeta_{t+1}}{1-\rho\phi}. \tag{17}$$

They imply the present value restriction:

$$\rho\tau^{dp}_{t+1} = \tau^d_{t+1} - \tau^r_{t+1} \Leftrightarrow \kappa_r - \kappa_d = 1 - \rho\phi. \tag{18}$$

Another way to write this restriction is as a restriction on a weighted sum of κ_r and κ_d: Any two equations from the system (4)-(6) implies the third. Evidence that dividend growth is not forecastable is evidence that returns are forecastable: if $\kappa_d = 0$ in equation (18), then $\kappa_r > 0$ because $\rho\phi < 1$. If estimating (5) uncovers that a high dividend-price ratio forecasts a higher future dividend growth ($\kappa_d > 0$), as we showed it does, then this strengthens the evidence for return predictability. [19] makes an important and closely related point: That it is important to impose the present-value relationship when testing the null hypothesis of no return predictability. That null ($\kappa_r = 0$) is truly a joint hypothesis, because it implies a negative coefficient in the dividend growth equation ($\kappa_d < 0$). [19], too, finds strong evidence for return predictability.

Returning to Panel A of Table 1, Column 3 backs out the AR(1) coefficient ϕ from the estimated κ_d and κ_r, and from the present-value constraint (18).[1] In the first row, $\phi = .945$, and is statistically undistinguishable from a unit root. This high persistence is a familiar result in the literature. The last column reports the left-hand side and the right-hand side of equation (18) for *univariate* OLS regressions of (4)-(6). It shows the violation of the present-value constraint. In the first row, the violation is half as large as the actual point estimate κ_r. The standard OLS point estimates do not satisfy the present-value constraint, which can lead to faulty inference.

However, when we use the break-adjusted dividend-price ratio series in rows 2 and 3, we find that (1) the persistence of the break-adjusted *dp* ratio is much lower than the unadjusted series (.81 and .69 versus .95), and (2) the present-value constraint is satisfied by the OLS coefficients.

A similar present-value constraint can be derived for long-horizon return and dividend growth

---

[1] The linearization parameter ρ is tied to the average dividend-price ratio, and is held fixed at 0.9635.

regressions:

$$\kappa_r(H) = \kappa_r \left(\frac{1 - \phi^H}{1 - \phi} \right)$$
$$\kappa_d(H) = \kappa_d \left(\frac{1 - \phi^H}{1 - \phi} \right)$$

Not only are the coefficients on the long-horizon return predictability regressions for all horizons linked to each other (see [8]), all long-horizon regression coefficients in the return equations are also linked to those from the dividend growth equations. I.e., there is one present-value constraint for each horizon H. Imposing these restrictions in a joint estimation procedure improves efficiency.

Panel B of Table 1 shows the results from a joint estimation of 1-year, 3-year, and 5-year cumulative returns and dividend growth rates on the lagged dividend-price ratio. Because of the restrictions, there are only two parameters to be estimated from these six equations. The results are close to those from the one-year system in Panel A, confirming the main message of [8]. The main conclusion remains that returns are strongly predictable, and dividend growth rates are not.

Exploiting Correlation in Innovations The present-value model implies a restriction on the innovations in returns and the dividend-price ratio (see equation 18). A third main insight from the structural model is that this correlation contains useful information for estimating the structural parameters, and hence for how much return predictability and dividend growth predictability there truly is. [48] show that exploiting the correlation between expected and unexpected stock returns can lead to substantially more accurate estimates. The information in correlations is incorporated by specifying a prior belief about the correlation between expected and unexpected returns, and updating that prior in a Bayesian fashion using observed data. Their method ignores the present-value constraint. The structural parameters in Panel B of Table 1, which impose the present-value constraint, imply that two-thirds of the variability in the price-dividend ratio is due to expected future returns and one-third is due to expected future dividend growth rates.

Likewise, [59] write down a model like (8)-(9) where expected returns and growth rates of dividends are auto-regressive, exploiting the present-value constraint. Because the price-dividend ratio is linear in expected returns x and expected dividend growth z (see equation (10)), its innovations in (17) can be attributed to either innovations in expected returns or expected growth rates. The present-value constraint enables one to disentangle the information in price-dividend ratios about both expected returns and growth rates, and therefore to undo the contamination coming from correlated innovations. With this decomposition in hand, it is then possible to recover the full time-series of expected returns, x, and expected growth rates, z. [59] show that the resulting processes are strong predictors of realized returns and realized dividend growth rates, respectively. This underscores the importance of specifying a present-value model to address return predictability.

Geometric or Arithmetic Returns As a final comment, most predictive regressions are estimated using geometric, i.e. log returns, instead of arithmetic, i.e. simple returns. This choice is predominantly motivated by the [13] log-linearization discussed before. Since investors are ultimately interested in arithmetic instead of log returns, [59] specify a process for expected simple returns instead. This is made possible by applying the techniques of linearity-inducing models, recently introduced by [28].

6 Future Directions

The efficient market hypothesis, which states that markets efficiently aggregate all information, was first interpreted to mean that returns are not predictable. Early evidence of predictability of stock returns by the lagged dividend-price ratio seemed to be evidence against the efficient market hypothesis. However, return predictability and efficient markets are not incompatible because return predictability arises naturally in a world with time-varying expected returns. In the last 15 years, the empirical literature has raised a set of statistical objections to return predictability findings. Meanwhile, the theoretical literature has progressed, seemingly independently, in its pursuit of new ways to build models with time-varying expected returns. Only very recently has it become clear that theory is necessary to understand the empirical facts.

In this entry, we have set up a simple present-value model with time-varying expected returns that generates the regression that is the focus of the empirical literature. The model also features time-varying expected dividend growth. It shows that the dividend-price ratio contains information about both expected returns and expected dividend growth. A regression of returns on the dividend-price ratio may therefore be a poor indicator of the true extent of return predictability. At the same time, the present-value model provides a solution to this problem: It disentangles the two pieces of information in the price-dividend ratio. This allows us to interpret the standard predictability regressions in a meaningful way. Combining data with the present-value model, we conclude that there is strong evidence for return predictability. We interpret this as evidence for the presence of time-varying expected returns, not evidence against the efficient market hypothesis. The main challenge for the future is to better understand the underlying reasons for this time-variation.

References

[1] Y. Amihud and C. M. Hurvich. Predictive regressions: A reduced-bias estimation method. *Financial and Quantitative Analysis*, 39:813–841, 2004.

[2] A. Ang and G. Bekaert. Stock return predictability: Is it there? *Forthcoming Review of Financial Studies*, 2006.

[3] M. Baker and J. Wurgler. The equity share in new issues and aggregate stock returns. *Journal of Finance*, 55:2219–2258, 2000.

[4] R. Balvers, Y. Wu, and E. Gilliland. Mean reversion across national stock markets and parametric contrarian investment strategies. *Journal of Finance*, 55:745–772, 2000.

[5] R. Bansal and A. Yaron. Risks for the long-run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, August 2004.

[6] J. B. Berk, R. C. Green, and V. Naik. Optimal investment, growth options and security returns. *Journal of Finance*, 54:1153–1607, 1999.

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

[8] J. Boudoukh, M. Richardson, and R. F. Whitelaw. The myth of long-horizon predictability. Forthcoming Review of Financial Studies, 2007.

[9] J. Y. Campbell. A variance decomposition for stock returns. *Economic Journal*, 101:157–179, 1991.

[10] J. Y. Campbell. Consumption-based asset pricing. In G. Constantinides, M. Harris, and R. Stulz, editors, *Handbook of the Economics of Finance* forthcoming. North-Holland, Amsterdam, 2003.

[11] J. Y. Campbell and J. H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107:205–251, 1999.

[12] J. Y. Campbell, A. W. Lo, and C. MacKinlay. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ, 1997.

[13] J. Y. Campbell and R. J. Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1:195–227, 1988.

[14] J. Y. Campbell and R. J. Shiller. Yield spreads and interest rates: A bird's eye view. *Review of Economic Studies*, 58:495–514, 1991.

[15] J. Y. Campbell and S. Thompson. Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 2007. Forthcoming.

[16] J. Y. Campbell and M. Yogo. Efficient tests of stock return predictability. Unpublished paper, Harvard University, 2002.

[17] J. H. Cochrane. Explaining the variance of price-dividend ratios. *Review of Financial Studies*, 5(2):243–280, june 1991.

[18] J. H. Cochrane. *Asset Pricing*. Princeton University Press, Princeton, NJ, 2001.

[19] J. H. Cochrane. The dog that did not bark: A defense of return predictability. Unpublished paper, University of Chicago Graduate School of Business, 2006.

[20] P. Eliasz. Optimal median unbiased estimation of coefficients on highly persistent regressors. Unpublished paper, Department of Economcis, Princeton University, 2005.

[21] E. F. Fama. The behavior of stock market prices. *Journal of Business*, 38:34–101, 1965.

[22] E. F. Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25:383–417, 1970.

[23] E. F. Fama. Efficient markets: II. *Journal of Finance*, 46(5):1575–1618, 1991.

[24] E. F. Fama and K. R. French. Dividend yields and expected stock returns. *Journal of Financial Economics*, 22:3–27, 1988.

[25] E. F. Fama and K. R. French. Permanent and temporary components of stock prices. *Journal of Political Economy*, 96(2):246–273, April 1988.

[26] W. E. Ferson and C. R. Harvey. The risk and predictability of international equity returns. *Review of Financial Studies*, 6:527–566, 1993.

[27] W. E. Ferson, S. Sarkissian, and T. T. Simin. Spurious regressions in financial economics? *Journal of Finance*, 58(4):1393–1413, August 2003.

[28] X. Gabaix. Linearity-generating processes: A modelling tool yielding closed forms for asset prices. working paper MIT, 2007.

[29] W. N. Goetzman and P. Jorion. Testing the predictive power of dividend yields. *Journal of Finance*, 48:663–679, 1993.

[30] W. N. Goetzman and P. Jorion. A longer look at dividend yields. *Journal of Business*, 68:483–508, 1995.

[31] A. Goyal and I. Welch. Predicting the equity premium with dividend ratios. *Management Science*, 49(5):639–654, May 2003.

[32] A. Goyal and I. Welch. A comprehensive look at the empirical performance of the equity premium prediction. *Review of Financial Studies*, 2006. Forthcoming.

[33] E. Hjalmarsson. On the predictability of global stock returns. Unpublished paper, Yale University, 2004.

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

[35] A. Inoue and L. Kilian. In-sample or out-of-sample tests of predictability: Which one should we use? *Econometric Reviews*, 23:371402, 2004.

[36] M. C. Jensen. Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, 6:95–101, 1978.

[37] S. Kothari and J. Shanken. Stock return variation and expected dividends: A time-series and cross-sectional analysis. *Journal of Financial Economics*, 31:177–210, 1992.

[38] O. Lamont. Earnings and expected returns. *Journal of Finance*, 53:1563–87, 1998.

[39] M. Lettau and S. C. Ludvigson. Consumption, aggregate wealth and expected stock returns. *Journal of Finance*, 56(3):815–849, June 2001.

[40] M. Lettau and S. C. Ludvigson. Expected returns and expected dividend growth. *Journal of Financial Economics*, 76:583–626, 2005.

[41] M. Lettau and S. Van Nieuwerburgh. Reconciling the return predictability evidence. Forthcoming Review of Financial Studies, 2006.

[42] J. W. Lewellen. Predicting returns with financial ratios. *Journal of Financial Economics*, 74(2):209–235, 2004.

[43] H. Lustig and S. Van Nieuwerburgh. Housing collateral, consumption insurance and risk premia: An empirical perspective. *Journal of Finance*, 60(3):1167–1219, 2005.

[44] H. Lustig and S. Van Nieuwerburgh. Can housing collateral explain long-run swings in asset returns? Unpublished manuscript, University of California at Los Angeles and New York University, 2006.

[45] L. Menzly, T. Santos, and P. Veronesi. Understanding predictability. *Journal of Political Economy*, 112(1):1–47, February 2004.

[46] C. C. Nelson and M. J. Kim. Predictable stock returns: The role of small sample bias. *Journal of Finance*, 43:641–661, 1993.

[47] L. Pastor and R. F. Stambaugh. The equity premium and structural breaks. *Journal of Finance*, 56(4):1207–1239, 2001.

[48] L. Pastor and R. F. Stambaugh. Predictive systems: Living with imperfect predictors. Graduate School of Business, University of Chicago, Working Paper, 2006.

[49] B. S. Paye and A. Timmermann. Instability of return prediction models. Forthcoming *Journal of Empirical Finance*, 2005.

[50] M. Piazzesi, M. Schneider, and S. Tuzel. Housing, consumption, and asset pricing. Forthcoming *Journal of Financial Economics*, 2006.

[51] C. Polk, S. Thompson, and T. Vuolteenaho. Cross-sectional forecasts of the equity risk premium. *The Journal of Financial Economics*, 81:101–141, 2006.

[52] J. M. Poterba and L. H. Summers. Mean reversion in stock returns: Evidence and implications. *Journal of Financial Economics*, 22:27–60, 1988.

[53] M. S. Rozeff. Dividend yields are equity risk premia. *Journal of Portfolio Managment*, 49:141–160, 1984.

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

[55] R. F. Stambaugh. Predictive regressions. *Journal of Financial Economics*, 54:375–421, 1999.

[56] L. H. Summers. Does the stock market rationally reflect fundamental values? *Journal of Finance*, 41:591–601, 1986.

[57] W. Torous, R. Volkanov, and S. Yan. On predicting returns with nearly integrated explanatory variables. *Journal of Business*, 77:937–966, 2004.

[58] R. Valkanov. Long-horizon regressions: Theoretical results and applications. *Journal of Financial Economics*, 68:201–232, 2003.

[59] J. van Binsbergen and R. S. Koijen. Predictive regressions: A present-value approach. Duke University, working paper, 2007.

[60] L. Viceira. Testing for structural change in the predictability of asset returns. Unpublished manuscript, Harvard University, 1996.

[61] T. Vuolteenaho. Understanding the aggregate book-market ratio and its implications to current equity-premium expectations. Unpublished paper, Harvard University, 2000.

Books and Reviews

J.Y. Campbell, A.W. Lo, and C. MacKinlay. *The Econometrics of Financial Markets.* Princeton University Press, Princeton, NJ, 1997.

J.H. Cochrane. *Asset Pricing.* Princeton University Press. Princeton. NJ, 2005.

B.G. Malkiel. *A Random Walk Down Wall Street.* W.W. Norton & Company, Inc. NY, 2004.