

Predictability of Returns and Cash Flows*

Ralph S.J. Koijen Stijn Van Nieuwerburgh
University of Chicago New York University
Booth School of Business, Stern School of Business,
and NBER NBER, and CEPR

January 6, 2011

Abstract

We review the literature on return and cash flow growth predictability from the perspective of the present-value identity. We focus predominantly on recent work. Our emphasis is on U.S. aggregate stock return predictability, but we also discuss evidence from other asset classes and countries.

JEL classification: G10, G12, G14, G35.

Keywords: return predictability, present-value model, dividend growth predictability.

*Koijen: University of Chicago, Booth School of Business, Chicago, 60637 IL; ralph.koijen@chicagobooth.edu; Tel: (773) 834-4199; <http://faculty.chicagobooth.edu/ralph.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 would like to thank Jules van Binsbergen, Martin Lettau, Sydney Ludvigson, and Motohiro Yogo for helpful comments.

1 Introduction

Stock prices move substantially relative to measures of fundamentals, such as dividends. The question of what moves stock prices has received an enormous amount of attention in the financial economics literature. Time variation in the price-dividend ratio is intimately linked to time variation in expected returns and in expected dividend growth rates. Therefore, the question can alternatively be posed as whether stock returns and dividend growth rates are predictable. The answer is of paramount importance to our understanding of the risk-return relationship, which lies at the core of finance theory.

Until the early 1980s, the standard model assumed constant expected returns for stocks. Then, empirical evidence was uncovered showing that returns were predictable by financial ratios, such as the price-dividend or price-earnings ratio. Later other variables, such as the spread between long-term and short-term bond yields, the consumption-wealth ratio, macroeconomic variables, and corporate decision variables were also shown to have predictive ability. The literature has expanded its interest to returns on other asset classes, such as government bonds, currencies, real estate, and commodities, and to many countries.

Initially, the finding of predictability was interpreted as evidence against the efficient market hypothesis.¹ Fama (1991) proposed the alternative explanation of time-varying expected returns. In the past twenty years, research in asset pricing has proposed several equilibrium models with efficient markets that generate time variation in expected returns: models with time-varying risk aversion (Campbell and Cochrane, 1999), time-varying aggregate consumption risk (Bansal and Yaron, 2004; Bansal, Kiku, and Yaron, 2009), time-varying consumption disasters (Gabaix, 2009), time-variation in risk-sharing opportunities among heterogeneous agents (Lustig and Van Nieuwerburgh, 2005), or time-variation in beliefs (Timmermann, 1993; Detemple and Murthy, 1994). The dominant view today is that predictability of asset returns is no longer *prima facie* evidence of market inefficiency.

A large partial equilibrium literature takes time variation in expected returns as given and asks how it affects optimal asset allocation decisions (see Wachter, 2010, for a recent review). Return predictability is of considerable interest to practitioners who can develop market-timing portfolio strategies that exploit predictability to enhance profits, if indeed such predictability is present.

Despite the theoretical developments, return predictability is a subtle feature of the data. A parallel literature developed in the 1990s questioning the strength of the statistical evidence. This literature points out problems such as biased regression coefficients, in-sample instability of

¹The efficient market hypothesis states that financial markets are efficient with respect to a particular information set when prices aggregate all available information (see Fama, 1965, 1970; Jensen, 1978; Shiller, 1984; Summers, 1986; Fama, 1991). Testing the efficient market hypothesis requires a “market model” that specifies how information gets incorporated into asset prices.

estimates indicating periods with and without predictability, and poor out-of-sample performance.

Recently, the literature has turned to the question of whether measures of cash-flow growth, such as dividend growth, are predictable as well. While interesting in its own right and important for model design, dividend growth directly speaks to return predictability. Through the present-value relationship, which links asset prices today to future returns and future dividend growth, dividend growth and return predictability are two sides of the same coin (Lettau and Van Nieuwerburgh, 2008; Cochrane, 2008; Binsbergen and Koijen, 2010).

In this article, we survey the return and dividend growth predictability debate, with an emphasis on recent developments.² The organizing principle is the present-value relationship between asset prices and their cash flows, broadening the discussion from return predictability to include cash flow growth predictability. While our emphasis is on U.S. stock return and dividend growth predictability by the price-dividend ratio, we discuss other return forecasters, return predictability in other asset classes, as well as in other countries. We conclude by offering some thoughts for future work.

2 U.S. Stock Return and Dividend Growth Predictability

We start by revisiting the basic findings on aggregate U.S. stock return predictability.

2.1 Motivating Predictive Regressions

Define the gross return on an 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 cash flow, hereafter dividend. Campbell and Shiller (1988b) log-linearize the return to obtain:

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

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 $\rho = \exp(-\overline{dp}) / (1 + \exp(-\overline{dp}))$ and $k = \log(1 + \exp(-\overline{dp})) + \rho \overline{dp}$ are a function of the long-run average log dividend-price ratio \overline{dp} .

²Earlier reviews include the special issue in the Review of Financial Studies (Spiegel, 2008), Koijen and Van Nieuwerburgh (2009), Lettau and Ludvigson (2010), and discussions in Campbell, Lo, and MacKinlay (1997); Cochrane (2001, 2006).

Iterating forward on equation (1) and by imposing a transversality condition, one obtains:

$$dp_t = \bar{dp} + E_t \sum_{j=1}^{\infty} \rho^{j-1} [(r_{t+j} - \bar{r}) - (\Delta d_{t+j} - \bar{d})]. \quad (2)$$

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 present-value equation. As long as the (expected) returns and dividend growth are stationary, the dividend-price ratio is 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. One alternative possibility we return to later is low-frequency changes in the long-run mean of the dividend-price ratio \bar{dp} . We also discuss the literature that argues that the dp ratio displays (near-) unit root behavior.

Equation (2) motivates some of the earliest empirical work in the stock return predictability literature, which regresses returns on the lagged dividend-price ratio, as in equation (3):

$$(r_{t+1} - \bar{r}) = \kappa_r (dp_t - \bar{dp}) + \tau_{t+1}^r, \quad (3)$$

$$(\Delta d_{t+1} - \bar{d}) = \kappa_d (dp_t - \bar{dp}) + \tau_{t+1}^d, \quad (4)$$

where \bar{r} is the long-run mean return and τ^r is a mean-zero innovation. However, the logic of (2) 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 (4), where \bar{d} denotes the long-run mean log dividend growth.

2.2 Data

The U.S. stock market return is the CRSP value-weighted market return containing all NYSE, AMEX, and NASDAQ stocks. Our sample is January 1926 to December 2009. We start from monthly cum-dividend and ex-dividend returns. Their difference, multiplied by the lagged ex-dividend price is the monthly dividend:

$$D_t = (R_t^{cum} - R_t^{ex}) P_{t-1}.$$

This dividend has a strong seasonal pattern. To de-seasonalize the data, the typical solution is to aggregate the dividends paid out over the year.

One important question, recently highlighted by Chen (2009) and Binsbergen and Koijen (2010), is what to assume about the reinvestment rate on these monthly dividends received within the year. We study three alternatives we have been used interchangeably in the literature. The first option is

to reinvest them at a zero rate. This amounts to adding up the dividends in the current month and the past eleven months. The second option is to reinvest the dividends at the risk-free rate. For example, the dividend received in month $t-2$ is reinvested at the one month interest rate prevailing at $t-2$ and again reinvested one more month at the one month rate prevailing at $t-1$. We use the CRSP 30-day T-bill rate as our risk-free rate. The third option is to reinvest the dividends at the cum-dividend stock market return. CRSP computes quarterly or annual return series under the stock market reinvestment assumption.

We believe that this third approach is problematic because it imparts some of the properties of returns to cash flows. To underscore this point, we compute annual returns $(D_t + P_t)/P_{t-12} - 1$ and dividend growth rates $D_t/D_{t-12} - 1$, where D is the 12-month dividend annualized under each of the three reinvestment assumptions. This generates a time series of 84 observations (December 1926 until December 2009). While all three approaches generate nearly identical annualized return series, they generate quite different annual dividend growth series, and dividend-price series. While the first two approaches generate dividend growth series with the same mean (5.12%), volatility (11.9% vs. 12.0%), and a high correlation of 99.88%, the dividend growth series with market return-reinvestment has a mean of 5.62%, a much higher volatility of 15.1%, and a correlation of only 40% with the other two series. To gauge the discrepancy, it suffices to look at the year 2009: The cash-reinvested dividend growth rate was -17.0% while the market return-reinvested growth rate was +30.5%. In what follows, we use the annual return series constructed from monthly data under the assumption of reinvestment at the risk-free rate as our default, and compare results with those of the market return reinvestment strategy.

For the return predictability exercises below, we *define* log returns by equation (1) above, unless explicitly stated otherwise. Hence, we assume that the approximation is accurate. In addition to the full sample of firms, we also study the universe of firms without NASDAQ (CRSP data).

2.3 Univariate DP-Predictability Evidence

The return predictability equation (3) was tested among others by (Rozeff, 1984; Campbell and Shiller, 1988b; Fama and French, 1988; Cochrane, 1991b; Hodrick, 1992; Goetzmann and Jorion, 1993; Lewellen, 2004). It generally found statistically significant evidence for return predictability, i.e., $\kappa_r > 0$.

We review and update the empirical evidence on U.S. stock return predictability by the dividend-price ratio using the latest available data. The estimation is by Ordinary Least Squares with Newey-West t-statistics with one lag. The return predictability coefficient $\kappa_r = 0.077$ with a t-statistic of 1.31; the R^2 is 2.90%. A high dividend-price ratio at the end of a year predicts a higher return over the following year, as suggested by the present-value equation, but the point

estimate suggests that there is no statistically significant evidence for return predictability over the full 1926-2009 sample. See the three most left columns of Table 1, first row. The results are very similar if we use actual log returns instead of log-linearized returns (second row of Table 1), if we predict log returns in excess of the log risk-free rate (row 2),³ real returns instead of nominal returns (row 4), if we exclude NASDAQ stocks (row 5).⁴

The results are quite different if we focus on the post-1945 sample (first observation on dp is 1945, first return is 1946). In the post-war sample, the evidence for stock return predictability is much stronger. Row 6 of Table 1 shows a point estimate that is almost twice as high as in the full sample, at $\kappa_r = 0.130$, with a t-statistic of 2.56, and an R^2 of 10.84%. One important take-away is that the sample period matters for the evidence on return predictability, using simple predictability regressions with dp as the predictor. We return to the sample specificity and instability of the results in the next section and to a different method in section 4.

The full sample results stand in sharp contrast with the ones obtained using the market-return reinvestment strategy instead of the risk-free rate-reinvestment strategy. The right columns of Table 1, panel A, show that the evidence for predictability is considerably stronger. The point estimate of $\kappa_r = 0.104$ is one-third higher with a t-statistic of 2.08 and a respectable R^2 of 4.82%. The difference between the left and right columns arises because the dividend-price ratio is different across reinvestment methodologies, while the return series are almost identical. Thus, the evidence on return predictability in the full sample depends strongly on the reinvestment assumption. Our interpretation is that the stronger evidence of return predictability arises from assigning some of the intra-year movement in prices to dividends (and therefore to the return predictor dp) rather than to cash-flows. Because the results are similar across reinvestment strategies in the post-1945 sample, the contamination is concentrated in the pre-1945 era.

To interpret the economic magnitudes of the coefficients, consider the 0.13 point estimate for κ_r in the post-1945 sample. A swing in dp from one standard deviation below its mean to one standard deviation above its mean (a 0.89 change) results in an increase in annual stock market return of 11.5% (or 0.58 return standard deviations). Given an average Sharpe ratio of 36% on the U.S. stock market, Campbell and Thompson (2008) calculate that an annual R^2 of 10.80% corresponds to a doubling in the expected return of a dynamic equity strategy compared a static one that does not exploit predictability. Or, it corresponds to an increase of 6.7% in expected return for an investor with a relative risk aversion coefficient of two.⁵ Return predictability significantly affects

³The annual risk-free rate is calculated as the return on rolling over the one-month Treasury bill rate for 12 months.

⁴NASDAQ stocks enter the sample in 1972 and comprise about 15% of the market's capitalization afterwards, on average.

⁵The proportional/percentage increase in the expected return is $\frac{R^2}{1-R^2} \frac{1+S^2}{S^2}$, where S is the unconditional Sharpe ratio of the asset and R^2 is the R-squared of the return predictability regression. In our case, $S = 0.36$, $R^2 = 0.108$, and we obtain an increase of 105.53%. The additive/percentage point increase in the expected return is

Table 1: Return and Dividend Growth Predictability by the Dividend-Price Ratio

The table reports Ordinary Least Squares point estimates, Newey-West t-statistics (1 lag), and R-squared statistics for regressions of annual log stock returns (Panel A) or annual log dividend growth rates on the lagged log dividend-price ratio. In the left columns, returns, dividend growth rates, and dividend-price ratios as computed assuming that dividends are reinvested at the risk-free rate within the year, whereas the right columns assume reinvestment at the stock market return. The sample runs from December 1926 until December 2009 (84 annual observations).

		Panel A: Return Predictability					
		Div. Reinv. at R^f			Div. Reinv. at R^m		
		κ_r	$t - stat$	R^2	κ_r	$t - stat$	R^2
1	benchmark 1926-2009	0.077	1.31	2.90	0.104	2.08	4.82
2	no approximation	0.075	1.29	2.79	0.102	2.03	4.58
3	excess returns	0.087	1.51	3.64	0.115	2.30	5.73
4	real returns	0.085	1.57	3.46	0.105	2.10	4.86
5	no NASDAQ	0.084	1.35	2.92	0.120	2.33	5.24
6	benchmark 1945-2009	0.130	2.56	10.84	0.126	2.58	10.02
		Panel B: Dividend Growth Predictability					
		Div. Reinv. at R^f			Div. Reinv. at R^m		
		κ_d	$t - stat$	R^2	κ_d	$t - stat$	R^2
7	benchmark 1926-2009	-0.078	-1.48	7.64	0.008	0.20	0.05
8	real div gr	-0.070	-1.63	7.61	0.009	0.23	0.07
9	no NASDAQ	-0.100	-1.74	10.55	0.008	0.22	0.05
10	benchmark 1945-2009	0.017	0.68	1.13	0.044	1.10	2.03

an individual's expected portfolio returns, at least in the post-war period.

Dividend growth predictability has received relatively little interest; exceptions are Fama and French (1988) and Lettau and Ludvigson (2005). To the limited extent that it has, the consensus view seems to be that growth rates of fundamentals, such as dividends or earnings, are much less forecastable than returns when using financial ratios as predictors. This view implies that most of the variation in the price-dividend or price-earnings ratio comes from variation in expected returns, not expected future cash-flow growth. We retrieve this result in the right columns of Panel B of Table 1 for the dp ratio. Indeed, under the market-reinvestment assumption, the point estimate on κ_d (equation 4) is not statistically different from zero (right columns). The R^2 is essentially zero, whether we look at the pre-war, post-war, or full sample. We have argued that this case suffers from mismeasurement of dividend growth and the dividend-price ratio. Also, Binsbergen and Koijen (2010) show that lagged market returns mechanically predict dividend growth in this case.

In contrast, the evidence for dividend growth predictability by the dividend-price ratio is substantially stronger under the cash- (or no) reinvestment assumption; left columns of Panel B. Indeed, the full sample point estimate is large (as large as the return coefficient) and negative, as suggested by equation (2). The regression's R^2 is as high as 7.64%, 2.6 times the value for the return equation's R^2 . While the R^2 is high in economic terms, the standard error remains too large for statistical significance at conventional levels, at least when using simple predictive regressions. The results are similar with real dividend growth (row 8) or in the no-NASDAQ sample (row 9). In the post-1945 sample, when we find stronger evidence for return predictability, there is no evidence for dividend growth predictability (row 10). While based on a short sample, we find strong dividend growth predictability evidence in the 1927-1945 sample. The point estimate for κ_d is -0.642 with a t-stat of -6.24 and an R^2 of 74.4% (not reported in the table). These results are consistent with Chen (2009), who shows that dividend growth is predictable in the 1871-1945 period. In contrast, he shows that earnings growth is predictable in both eras (implying that the dividend-earnings ratio is predictable post-war). Chen, Da, and Priestley (2010) argue that dividend growth predictability disappears once firms started smoothing dividends in the post-war period. We return to alternative cash flow measures below.

In sum, when our favorite dividend reinvestment procedure is used, returns do not appear to be predictable while dividend growth rates do (marginally) for the full 1927-2009 sample. These conclusions reverse in the post-1945 sample. In section 4, we argue that the return and dividend growth predictability results are closely related to one another and explore their link in depth.

$\frac{1}{\gamma} \frac{R^2}{1-R^2} (1+S^2)$, where γ is the coefficient of relative risk aversion of the investor. We note that these calculations assume that the return predictability parameters are known with certainty by the investor.

3 Statistical Issues

Findings regarding the predictability of stock returns, like the ones presented in the previous section, 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 (3) with an auto-regressive specification for the predictor variable,

$$(dp_{t+1} - \bar{dp}) = \delta(dp_t - \bar{dp}) + \tau_{t+1}^{dp}, \quad (5)$$


where \bar{dp} is the long-run mean of the dividend-price ratio. The estimated autoregressive parameter δ is near unity and standard tests leave the possibility of a unit root open (i.e., $\delta = 1$). Nelson and Kim (1993), Stambaugh (1999), Ang and Bekaert (2007), Ferson et al. (2003), and Valkanov (2003) conclude that the statistical evidence of forecastability is weaker once tests are adjusted for high persistence. Amihud and Hurwicz (2004), Ang and Bekaert (2007), Campbell and Yogo (2006), Lewellen (2004), Torous et al. (2004), and Elias (2005) derive asymptotic distributions for predictability coefficients under the assumption that the forecasting variable follows a local-to-unit root, yet stationary, process. We note that in our annual 1926-2009 sample, the persistence of the cash-reinvested dividend-price ratio is 0.86, smaller than the 0.93 persistence of the market return-reinvested ratio. Hence, the predictor we advocate is less persistent than the measure that is traditionally considered in the literature.

Second, the forecasting relationship of returns and financial ratios exhibits significant instability over time. The left panel of 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 κ_r varies between zero and 0.5 (solid line). The right panel shows that the corresponding t-statistic fluctuates between 0 and 5 (solid line). In other words, for thirty-year samples ending between 1965 and 1995, there was evidence for stock return predictability, but this evidence disappeared after 1995. It was absent for the pre-war period as well. The figure also investigates how dividend growth predictability evidence depends on the sample (dashed line). In contrast with the return predictability picture, there was considerable evidence for dividend growth predictability for samples ending before 1965. Afterwards, the κ_d coefficient is indistinguishable from zero.

Recognizing this instability, Viceira (1996) and Paye and Timmermann (2006) report evidence in favor of breaks in the OLS coefficient in the forecasting regression of returns on the lagged dividend-price ratio, while Lettau and Van Nieuwerburgh (2008) report evidence for structural shifts in the mean of the dividend-price ratio. We return to the latter idea below. Constantinides and Ghosh (2010) study return predictability in the presence of structural breaks in consumption

Figure 1: Parameter Instability in Return Predictability Coefficient

The figure plots estimation results for the equations $r_{t+1} - \bar{r} = \kappa_r(dp_t - \bar{dp}) + \tau_{t+1}^r$ and $\Delta d_{t+1} - \bar{g} = \kappa_d(dp_t - \bar{dp}) + \tau_{t+1}^d$. The left panel shows the estimates for κ_r (solid line) and κ_d (dashed line) using 30-year rolling windows. The right panel shows (asymptotic) t-statistics computed with Newey-West with one lag. The data are annual for 1926-2009. The first rolling window ends in 1957.



and dividend growth. Working in a Bayesian setting, Pastor and Stambaugh (2001) estimate structural breaks in the equity premium, while Dangl and Halling (2009) and Johannes, Korteweg, and Polson (2008) consider drifting regression coefficients to address parameter instability.

Third, financial ratios have poor out-of-sample forecasting power, as shown in Bossaerts and Hillion (1999) and Goyal and Welch (2003, 2008). Inoue and Kilian (2004), Campbell and Thompson (2008), and Kelly and Pruitt (2010) have a different take on the out-of-sample tests and evidence. Lettau and Van Nieuwerburgh (2008) attribute most of the poorer out-of-sample performance with the difficulty of estimating the new mean of the dp ratio after the economy undergoes a structural break, rather than with detecting the break itself.

3.1 Structural Break Adjustment

Lettau and Van Nieuwerburgh (2008) report evidence for structural shifts in the mean of the dividend-price ratio \bar{dp} in 1954 and 1994. They advocate using a demeaned dp series as return predictor, where regime-specific means are used in the demeaning procedure. The correspondingly break-adjusted dp ratio displays much less persistence. The first order autocorrelation of the break-adjusted dp ratio is 0.502 under the cash reinvestment assumption and 0.641 under the market return reinvestment assumption. The reduced persistence alleviates the upward small-sample bias

in the return predictability coefficient (see Stambaugh, 1999). The traditional dp ratio's persistence is 0.858 and 0.927, respectively. The break-adjusted dp series are also only half as volatile.

Table 2 shows that the adjusted dp ratio leads to much stronger return predictability and reduced in-sample coefficient instability. For example, in the 1926-2009 sample, the return predictability coefficient is $\kappa_r = 0.212$, three times higher than the benchmark 0.077 slope in row 1 of Table 1. It has a t-stat of 2.32 compared to 1.31. The R^2 is 6.2% compared to 2.9%. Panel B shows that dividend growth predictability is now also significantly negative with a point estimate of $\kappa_d = -0.240$, a t-statistic of -2.53, and an R^2 of 20.5% per annum. These numbers are almost three times higher as in the benchmark case without \bar{dp} adjustment.⁶ In sum, the low-frequency movements in the unadjusted dp ratio mask evidence of return predictability as well as dividend growth predictability.

Low-frequency changes in expected returns or expected dividend growth can reflect a variety of structural changes to the economy. Lettau, Ludvigson, and Wachter (2008) argue that a persistent decline in the volatility of aggregate consumption growth leads to a decline in the equity premium. The latter decline also arises in models with persistent improvements in the degree of risk-sharing among households or regions, for example due to developments in the market for housing-collateralized debt (Lustig and Van Nieuwerburgh, 2007, 2009; Favilukis, Ludvigson, and Van Nieuwerburgh, 2010), if the tax code changes persistently (McGrattan and Prescott, 2005), or if there is a gradual entry of new participants in stock markets (Calvet, Gonzalez-Eiras, and Sodini, 2004). Other models argue that there was a persistent increase (or at least a perception thereof) in the long-run growth rate of the economy in the 1990s (Quadrini and Jermann, 2007). Either effect lowers the steady-state level of the dividend-price ratio \bar{dp} . Other candidates for the non-stationarity in the traditional dp measure are persistent changes in firms' payout policies, including dividend smoothing, discussed above, or share repurchases, discussed below. While such slow-moving changes are obviously important for long-run asset prices, they mask the time variation in expected returns and expected dividend growth at higher frequencies.

4 The Present-Value Relationship

As indicated by the present-value equation (2), return and dividend growth predictability by the dp ratio are tightly connected to one another. In fact, weaker evidence for dividend growth predictability points to stronger evidence for return predictability and vice versa (Lettau and Van Nieuwerburgh, 2008; Cochrane, 2008; Binsbergen and Koijen, 2010). The two slope coefficients in equations (3) and (4) are closely connected to each other.

⁶Using the market return reinvested dividend price ratio (in the right columns), the κ_r point estimate is 0.393 compared to 0.103 and the t-statistic is 4.29 compared to 2.08. Dividend growth remains unpredictable.

Table 2: Structural-Break Adjusted Results

The table reports Ordinary Least Squares point estimates, Newey-West t-statistics (1 lag), and R-squared statistics for regressions of annual log stock returns (Panel A) or annual log dividend growth rates on the lagged log adjusted dividend-price ratio. The adjusted dividend-price ratio is computed from the raw dividend-price ratio by demeaning, and adding back in a regime-specific mean for 1926-1953, 1954-1994, and 1994-2009. The 1954 and 1994 breakpoints are estimated in Lettau and Van Nieuwerburgh (2008). In the left columns, returns, dividend growth rates, and dividend-price ratios as computed assuming that dividends are reinvested at the risk-free rate within the year, whereas the right columns assume reinvestment at the stock market return. The sample runs from December 1926 until December 2009 (84 annual observations).

		Panel A: Return Predictability					
		Div. Reinv. at R^f			Div. Reinv. at R^m		
		κ_r	$t - stat$	R^2	κ_r	$t - stat$	R^2
1	benchmark 1926-2009	0.212	2.32	6.20	0.393	4.29	14.91
2	no approximation	0.213	2.36	6.36	0.393	4.31	15.00
3	excess returns	0.195	2.15	5.16	0.365	3.65	12.64
4	real returns	0.215	2.61	6.33	0.362	3.74	12.57
5	no NASDAQ	0.206	2.28	5.96	0.398	4.42	15.13
6	benchmark 1945-2009	0.322	4.47	17.25	0.357	4.17	17.72
		Panel B: Dividend Growth Predictability					
		Div. Reinv. at R^f			Div. Reinv. at R^m		
		κ_d	$t - stat$	R^2	κ_d	$t - stat$	R^2
7	benchmark 1926-2009	-0.240	-2.53	20.52	0.107	1.37	2.15
8	real div gr	-0.237	-3.11	24.42	0.076	1.00	1.12
9	no NASDAQ	-0.255	-2.72	23.23	0.106	1.39	2.26
10	benchmark 1945-2009	-0.021	-0.33	0.42	0.133	1.86	4.08

To better understand this connection, we introduce a simple model that assumes that expected returns, μ_t , and expected dividend growth, g_t each follow AR(1) models.

$$\Delta d_{t+1} - \bar{d} = g_t + \varepsilon_{t+1}^d, \quad g_{t+1} = \gamma g_t + \varepsilon_{t+1}^g, \quad (6)$$

$$r_{t+1} - \bar{r} = \mu_t + \varepsilon_{t+1}^r, \quad \mu_{t+1} = \delta \mu_t + \varepsilon_{t+1}^\mu. \quad (7)$$

The model has three shocks: an innovation in unexpected dividends ε_{t+1}^d , an innovation in expected dividend growth ε_{t+1}^g , and an innovation in expected returns ε_{t+1}^μ . We assume that all three shocks are serially uncorrelated and have zero cross-covariance at all leads and lags, except for a contemporaneous correlation between expected return and expected dividend growth innovations $\text{Corr}(\varepsilon_t^g, \varepsilon_t^\mu) = \rho_{\mu g}$, and a correlation between expected and unexpected dividend growth innovations $\text{Corr}(\varepsilon_t^d, \varepsilon_t^\mu) = \rho_{\mu d}$.

In steady-state, the log dividend-price ratio is a function of the long-run mean return and dividend growth rate $\bar{dp} = \log\left(\frac{\bar{r}-\bar{d}}{1+\bar{d}}\right)$. The log dividend-price ratio in (2) can then be written as the difference between an expected return and an expected dividend growth term:

$$dp_t - \bar{dp} = \frac{\mu_t}{1 - \rho\delta} - \frac{g_t}{1 - \rho\gamma}. \quad (8)$$

Campbell (1991)'s return decomposition implies that the innovation to unexpected returns follows from the three fundamental shocks (i.e., combine (1) with (6)-(8)):

$$\varepsilon_{t+1}^r = \frac{-\rho}{1 - \rho\delta} \varepsilon_{t+1}^\mu + \frac{\rho}{1 - \rho\gamma} \varepsilon_{t+1}^g + \varepsilon_{t+1}^d. \quad (9)$$

Since ρ , δ , and γ are positive, and $\rho\delta < 1$ and $\rho\gamma < 1$, 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.

The simplest case of this setup is considered in Lettau and Van Nieuwerburgh (2008), where the persistence of expected dividend growth is assumed equal to that of expected returns: $\delta = \gamma$. In that case, the log dividend-price ratio in equation (8) follows an AR(1) process, as in equation (5), with persistence δ . In the more general case, considered in Binsbergen and Koijen (2010), $\delta \neq \gamma$ and dp follows an ARMA(1,1). Similar models can be derived for financial ratios other than the dividend-price ratio (Vuolteenaho, 2000).

4.1 Present-Value Relationship

The second insight from the model is that there is a linear relationship between the three innovations $\tau = (\tau^d, \tau^r, \tau^{dp})$. Starting from the definition of a return in (1), and subtracting the linear

projection on dp_t , we can write return innovations in terms of dividend growth and dp innovations:

$$r_{t+1} - E(r_{t+1} | dp_t) = \Delta d_{t+1} - E(\Delta d_{t+1} | dp_t) - \rho(dp_{t+1} - E(dp_{t+1} | dp_t)),$$

or, equivalently:

$$\tau_{t+1}^r = \tau_{t+1}^d - \rho\tau_{t+1}^{dp}.$$

By the same logic, there is a tight relationship between the predictive coefficients of returns and dividend growth rates, and the persistence of the dividend yield. In the simplest case where the dp ratio is assumed to be an AR(1), this relationship is:

$$\kappa_r = \frac{Cov(r_{t+1}, dp_t)}{Var(dp_t)} = \frac{Cov(\Delta d_{t+1} - \rho dp_{t+1} + dp_t, dp_t)}{Var(dp_t)} = \kappa_d - \rho\delta + 1. \quad (10)$$

Evidence that dividend growth is not forecastable is evidence that returns are forecastable: if $\kappa_d = 0$ in equation (10), then $\kappa_r > 0$ because $\rho\delta < 1$. If estimating (4) uncovers that a high dividend-price ratio forecasts lower future dividend growth ($\kappa_d < 0$), then this is indirect evidence of return predictability. Cochrane (2008) dubs the lack of dividend growth predictability “the dog that did not bark,” highlighting that the null hypothesis of no return predictability ($\kappa_r = 0$) is a joint hypothesis, because it implies a negative coefficient in the dividend growth equation ($\kappa_d < 0$). The lack of evidence for $\kappa_d < 0$ is evidence in favor of $\kappa_r > 0$.

4.2 Price Variation

We can use the present-value model to understand why asset prices fluctuate over time. Variation in the price-dividend ratio reflects variation in expected returns (discount rate news), variation in expected dividend growth (cash-flow news), or their covariance:

$$\begin{aligned} \text{Var}(pd_t) &= \frac{\text{Var}(\mu_t)}{(1 - \rho\delta)^2} + \frac{\text{Var}(g_t)}{(1 - \rho\gamma)^2} - \frac{2\text{Cov}(\mu_t, g_t)}{(1 - \rho\delta)(1 - \rho\gamma)}, \\ &= \frac{\sigma_\mu^2}{(1 - \delta^2)(1 - \rho\delta)^2} + \frac{\sigma_g^2}{(1 - \gamma^2)(1 - \rho\gamma)^2} - \frac{2\rho\mu_g\sigma_\mu\sigma_g}{(1 - \delta\gamma)(1 - \rho\delta)(1 - \rho\gamma)}, \end{aligned} \quad (11)$$

$$= -\frac{\text{Cov}(pd_t, \mu_t)}{1 - \rho\delta} + \frac{\text{Cov}(pd_t, g_t)}{1 - \rho\gamma}. \quad (12)$$

Alternatively, we can decompose the variance of the pd ratio into the covariance with expected returns and the covariance with expected growth rates (equation 12). The first covariance equals the first term in equation (11) plus half the third term, while the second covariance equals the second term plus half the third term.

4.3 Long-Horizon Regressions

The evidence favoring return predictability, as measured by the R^2 statistic, tends to be stronger at longer horizons.⁷ We investigate long-horizon return R^2 in the context of the present-value model. We start from the h -period stock return:

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

and log-linearize as before:

$$r_{t+h}(h) = k + \sum_{i=1}^h \Delta d_{t+i} - \rho dp_{t+h} + dp_t. \quad (13)$$

We compute the R^2 value of the h -period return predictability regression

$$r_{t+h}(h) = \bar{r}_h + \kappa_{r,h} dp_t + \tau_{t+h}^r,$$

as $R^2(h) = \kappa_{r,h}^2 \text{Var}(dp_t) / \text{Var}(r_{t+h}(h))$.

To simplify the exposition, we assume that dividend growth is unpredictable ($\sigma_g = 0$) and that dividend growth shocks are independent of discount rate shocks ($\rho_{\mu d} = 0$). Then the per-period variance of returns is

$$\frac{\text{Var}(r_{t+h}(h))}{h} = \underbrace{\sigma_d^2}_{\text{Cash flows}} + \underbrace{\frac{1}{h} \frac{(\rho^2 + 1 - 2\rho\delta^h) \sigma_\mu^2}{(1 - \delta^2)(1 - \rho\delta)^2}}_{\text{Discount rates}}.$$

This equation clearly shows that, as $h \rightarrow \infty$, all variation in returns reflects cash-flow risk (see Hansen, Roberds, and Sargent, 1991; Lustig and Van Nieuwerburgh, 2008).

Returning to the return predictability R^2 , we get

$$R^2(h) = \frac{1 + \rho^2 (\delta^h)^2 - 2\rho\delta^h}{h\sigma_d^2/\text{Var}(pd_t) + 1 + \rho^2 - 2\rho\delta^h} \quad (14)$$


It is easy to see that $\lim_{h \rightarrow \infty} R^2(h) = 0$. However, in an intermediate range, the R^2 first increases in h before decreasing. Figure 2 reports model-implied R^2 for long-horizon return predictability equations. It uses the observed volatility of dividend growth for σ_d , as well as the estimated

⁷See Fama and French (1988); Campbell and Shiller (1988b). However, Boudoukh et al. (2008) argue that long-horizon and short-horizon predictability estimators are highly correlated, presenting little independent information on the presence of return predictability. Cochrane (2008) reconciles both views arguing that the power advantage arises from horizons beyond five years.

persistence of expected returns reported below. The return R^2 displays a hump shape and peaks around 35 years.

Figure 2: Long-horizon R-squared in Present-value Model

The figure plots the long-horizon return predictability R^2 from the equations $r_{t+h}(h) = \bar{r}_h + \kappa_{r,h}(dp_t - \bar{dp}) + \tau_{t+h}^r$, as implied by the present-value model. It uses equation (14). The parameter values are based on the post-war sample. We match $\sigma_d = 0.0688$ to the volatility of dividend growth, $\delta = 0.9112$ to the AR(1)-coefficient of the dividend yield, the variance of the dividend yield equals 18.19%, and the linearization constant equals $\rho = 0.9691$.



4.4 Estimation Results

We estimate the present-value model in (6) - (7), following Binsbergen and Koijen (2010). We use the Kalman filter to construct the model's likelihood, maximize that likelihood, and obtain the filtered expected return and dividend growth series $\hat{\mu}_t$ and \hat{g}_t as a byproduct.

Table 3 contains the point estimates. The first column uses the raw data (without adjustments for structural breaks) for the full sample. All results in this section assume that dividends are reinvested at the risk-free rate. We find expected returns to be more persistent than expected growth rates ($\delta = 0.93 > \gamma = 0.26$). The volatility of annual expected returns is 4.2%, while the volatility of expected dividend growth rates is 12.2%. Expected returns and expected growth rates have a modest positive correlation of 18.5%, consistent with Lettau and Ludvigson (2005) and Menzly et al. (2004).

We estimate the total amount of return predictability, measured as $1 - (Var[r_{t+1} - \hat{\mu}_t] / Var[r_{t+1}])$, to be 3.0% in the full sample. The total amount of dividend growth predictability, measured as

$1 - (Var[\Delta d_{t+1} - \hat{g}_t]/Var[\Delta d_{t+1}])$, is estimated to be a very large 46.8% in the full sample. In comparison with the simple predictability results reported in Table 1, we see that using the dp ratio as a predictor captures most of the total return predictability (2.9% versus 3.0%), but substantially understates the total dividend growth predictability (7.6% versus 46.8%). This reinforces our conclusion that returns are only modestly predictable over the full sample, while dividend growth rates are strongly predictable.

Despite the strong predictability of dividend growth rates, most of the variation in the price-dividend ratio arises from discount rate news. This happens because of the much higher persistence of expected returns (see also Lettau and Ludvigson, 2005). Our benchmark full-sample estimates imply a variance decomposition of the price-dividend ratio that assigns 93% of the variability to expected returns, 13% to expected growth rates, and -6% to their covariance (see equation 11). The last three rows of Table 3 report these shares. Alternatively, we can think of the variance of dp as coming 90% from covariance with expected returns and 10% from covariance with expected growth rates (see equation 12). The discrepancy between the amount of dividend growth predictability in simple predictability regressions and in the present-value model is caused by a severe errors-in-variables problem in the standard predictability regression (see also Fama and French, 1988; Kothari and Shanken, 1992; Goetzmann and Jorion, 1995; Binsbergen and Koijen, 2010). Indeed, the dividend-price ratio is not a good proxy for the expected dividend growth rate because it mostly captures fluctuations in expected returns.

Column 2 of Table 3 re-estimates the parameters on the post-war sample. In this sample, the volatility of expected growth rates is only half as big as in the full sample (6.9% versus 12.2%), while the volatility of expected returns is similar (4.6% versus 4.2%). Returns are substantially more predictable in the post-war period, with an R^2 of 9.1%. Dividend growth remains strongly predictable as well, with an R^2 of 18.9%. The latter result stands in strong contrast with that in Table 1, where we found no return predictability *by the dp ratio* in the post-war sample (1.1% R^2). The discrepancy is again caused by the errors-in-variables problem. The increased volatility and undiminished persistence of expected return raise its contribution to the variation in dp to 98% (103%-5%), leaving only 2% for the covariation of the dp ratio with expected dividend growth.

In Columns 3 and 4, we adjust the dividend-price ratio for structural breaks, as in Section 3. Since the dividend-price ratio is now less persistent, either expected returns or expected growth rates must be less persistent. We find that it is mostly the persistence of expected returns that is lower (0.66 versus 0.93). At the same time, the volatility of shocks to expected returns (σ_μ) more than triples. The resulting standard deviation of expected returns is 7.8% (8.5%) in the full (post-war) sample while the volatility of expected dividend growth rates is 12.3% (6.8%). Expected returns are much more volatile once the low-frequency component in the dp ratio is removed, while expected dividend growth rates are unaffected. Consistent with the results in Table 2, the

Table 3: Present Value Parameter Estimates

This table reports maximum likelihood estimates of the present-value model. The parameters are reported in the first column. Columns 2 and 4 report results for the 1926 to 2009 sample, while columns 3 and 5 report results for the 1945 to 2009 sample. Columns 2 and 3 use raw returns and dividend growth rates, while columns 4 and 5 use data that are adjusted for structural breaks in the mean dp ratio in 1954 and 1994.

	benchmark	benchmark	break-adjusted	break-adjusted
	1926-2009	1945-2009	1926-2009	1945-2009
\bar{r}	0.0725	0.0824	0.0767	0.0903
δ	0.9338	0.9243	0.6604	0.6424
\bar{d}	0.0431	0.0551	0.042	0.0549
γ	0.2564	0.3822	0.2935	0.3453
σ_μ	0.0150	0.0174	0.0585	0.0649
σ_g	0.1182	0.0639	0.1173	0.0638
σ_D	0.0026	0.0012	0.0086	0.0094
$\rho_{\mu g}$	0.1845	0.3395	0.2576	0.3269
$\rho_{\mu D}$	-0.8906	-0.2911	0.8662	0.9394
$Std[\mu_t]$	4.2%	4.6%	7.8%	8.5%
$Std[g_t]$	12.2%	6.9%	12.3%	6.8%
R^2 returns	3.0%	9.1%	6.7%	14.1%
R^2 div. gr.	46.8%	18.9%	46.5%	19.9%
$\%CF$	13%	7%	50%	22%
$\%DR$	93%	103%	79%	107%
$-2Cov(CF, DR)$	-6%	-10%	-29%	-30%

predictability of returns increases substantially with the break-adjustment, especially in the full sample. In contrast, predictability of dividend growth rates is not affected much. As with the unadjusted data, the total predictability of returns is captured almost completely by the dp ratio while the predictability of dividend growth rates is severely understated. For the full sample, the Campbell-Shiller decomposition assigns 64% of the variation of the break-adjusted dp ratio to expected returns and 36% to expected growth rates.

In sum, once a low-frequency component in expected returns is removed, it is clear that both returns and dividend growth rates are strongly predictable in both the full sample and the post-war sample. OLS regressions with the dp ratio as predictor severely under-estimate the extent of dividend growth predictability. This problem is so severe that they lead to the wrong conclusion regarding dividend growth predictability in the post-war sample. Dividend growth predictability, however, is more short-lived than that of returns. The higher persistence of expected returns leads the present-value model to assign the bulk of variation in price-dividend ratios (90% before and 64% after removing a low-frequency component in expected returns) to discount rate news.

5 Extensions

This section discusses extensions beyond U.S. stock return and dividend growth predictability by the dividend-price ratio. We discuss repurchase-adjusting returns, other return predictors, evidence from other countries, and evidence from other asset classes.

5.1 Repurchase Adjustment

A branch of the literature considers returns, dividend growth, and dividend-price ratios that are adjusted for stock repurchases. The idea is that dividends maybe too narrow a cash-flow measure because firms can substitute from dividend payments to share repurchases without reducing shareholders' payouts per share. In fact, the fraction of dividend paying firms dropped from 66.5% in 1978 to 20.8% in 1999 (Fama and French, 2001), while repurchases gradually rose. The observed low-frequency decline in the traditional dividend-price ratio in the 1990s may be partially attributable to this change in the composition of payouts.

Boudoukh, Michaely, Richardson, and Roberts (2007) show that the repurchase-adjusted dp ratio more strongly predicts U.S. stock returns. Updating their results to 2009, we find that the (market-reinvested) return predictability coefficient κ_r is 0.188 for the sample 1926-2009, with a t-stat of 2.96 and an R^2 of 6.88%.⁸ Comparing this to the non-repurchase-adjusted results of Table 1, this is indeed higher than the 0.104 slope with t-stat of 2.08 and R^2 of 4.82. In the post-1945

⁸We thank Itamar Drechsler for making his repurchase-adjusted data available to us.

sample, the return predictability coefficient increases from 0.126 (with t-stat of 2.58) to 0.205 (with t-stat of 3.17). Dividend growth remains unpredictable in both samples.

However, we have advocated the use of cash-reinvested dividends. Under that reinvestment assumption, κ_r increases from 0.077 without repurchase adjustment to 0.117, while the t-statistic increases from 1.31 to 1.47 and the R^2 from 2.90 to 3.27. Hence, we continue to find no evidence for return predictability in the full sample using standard predictability regressions, despite the repurchase adjustment. In the post-1945 sample, we estimate $\kappa_r = 0.194$ (with t-stat of 3.23), higher than the benchmark coefficient of 0.130 (with t-stat of 2.56). One important finding is that dividend growth is predictable by the (cash-reinvested) repurchase-adjusted dp ratio in the full sample. The coefficient κ_d is -0.221 with a t-stat of -2.60 and an R^2 of 17.10%. For comparison, without repurchase adjustment, we found that κ_d is -0.078 with a t-stat of -1.48 and an R^2 of 7.64%. Dividend growth in the post-1945 sample continues to be unpredictable.

The reason for these quantitatively and qualitatively different findings is that the repurchase adjustment makes the dp series less persistent. This alleviates some of the statistical issues reviewed on above.⁹ At a mechanical level, the repurchase adjustment works similarly to the break-adjustment we discussed above. Put differently, changes in firms' payout policies may well be an important contributor to the non-stationarity in the traditional dp ratio.¹⁰ Finally, whether a repurchase adjustment is appropriate or not depends on the question at hand: does one consider an investor who always holds one share of the stock or an investor who participates in every stock repurchase. Yet another alternative is to consider an investor who participates both in stock repurchases and in all initial public offerings (see Larrain and Yogo, 2008; Bansal, Fang, and Yaron, 2005). Larrain and Yogo (2008) show that about half of the variation in the price-dividend ratio is due to variation in expected growth rates, once repurchases and issuances are included in the cash flow measure.

5.2 Other predictors

There is a voluminous literature documenting stock return predictability by other variables than the dividend-price ratio. They can be grouped in four categories: other financial ratios, term structure

⁹The first order autocorrelation of the repurchase-adjusted dp ratio is 0.680 under the cash reinvestment assumption and 0.825 under the market return reinvestment assumption. The traditional dp ratio's persistence is 0.858 and 0.927, respectively. The repurchase-adjusted dp series are about one-third less volatile.

¹⁰However, there is evidence that it is not the only or maybe even not the main determinant of the non-stationarity. Lettau and Van Nieuwerburgh (2008) find evidence for structural breaks in the repurchase-adjusted dividend-price ratio. Break-adjusting this series further strengthens return predictability under market-reinvestment of dividends. Under the cash-reinvestment assumption, the return predictability coefficient also increases, but is still not statistically significant at the 10% level. Since repurchase adjustment intensities are higher among NASDAQ firms, another (cruder) way to repurchase-adjust the data is to omit these firms. As we showed above, the results are similar with or without these firms.

variables, macro-economic quantity variables, and corporate decision variables. Among financial ratios, the earnings-price ratio (Campbell and Shiller, 1988b,a, 2003), the book-to-market ratio (Kothari and Shanken, 1997; Pontiff and Schall, 1998), the dividend-payout ratio (Lamont, 1998), the variance risk premium (Bollerslev et al., 2009; Drechsler and Yaron, 2010), a linear combination of dp ratios of various stock portfolios (Kelly and Pruitt, 2010); among interest rate variables, the term and default spreads on bonds and the short-term interest rate (Fama and Schwert, 1977; Campbell, 1987; Fama and French, 1989; Hodrick, 1992; Ang and Bekaert, 2007); among macro-economic variables, the investment rate (Cochrane, 1991a; Lamont, 2000), the consumption-wealth ratio (Lettau and Ludvigson, 2001), the labor income-to-consumption ratio (Menzly et al., 2004), the housing collateral ratio (Lustig and Van Nieuwerburgh, 2005), the housing to non-housing consumption ratio (Piazzesi et al., 2007); and among corporate decision variables, equity share in total new equity and debt issues (Baker and Wurgler, 2000) or IPO activity.¹¹

The multitude of predictors naturally leads to the question whether a few factors effectively summarize the dynamics of expected returns. Using dynamic factor analysis, Ludvigson and Ng (2007) summarize the information of 172 financial series into a volatility factor and a “risk premium” factor. The R^2 of a regression of quarterly excess stock market returns (from 1960-2002) on these two lagged factors is 9%. Adding *cay* increases the R^2 to 16%. The first common factor of 209 macro-economic series is especially useful for predicting the conditional volatility of excess stock returns.¹² A different approach to deal with the multitude of possible predictors is a Bayesian averaging approach, which takes into account the uncertainty that investors may have about the “right” return predictors (e.g., Avramov, 2002; Cremers, 2002). Finally, Pastor and Stambaugh (2009) incorporate various return predictors in a “predictive system.” They specify realized and expected returns as in equation (7), but augment the measurement equation of the system with a vector-autoregression for observables. Compared to the standard predictability regressions, this recognizes that a linear combination of these observables is an imperfect proxy for expected returns. An interesting extension would be to integrate their approach with the present-value model.

5.3 International Evidence

Predictability of stock returns does not only arise for the US. Studies by Ferson and Harvey (1993), Harvey (1995), Campbell (2003), Paye and Timmermann (2006), Ang and Bekaert (2007),

¹¹In addition, there is a large literature predicting the stock returns of individual firms using firm-level variables such as the market-value to replacement cost (Q) or book-to-market ratio, the investment rate, the labor hiring rate, past returns, earnings surprises, accruals, corporate events, etc; see for example Zhang (2005); Bazdrech et al. (2009); Hirshleifer et al. (2009).

¹²Ludvigson and Ng (2007) find a robust positive correlation between expected returns and volatility of returns using this methodology, clarifying a mixed bag of results in the prior literature on the risk-return relationship (see also Brandt and Kang, 2004; Guo and Whitelaw, 2006; Pastor et al., 2008, on the risk-return tradeoff).

and Hjalmarsson (2010) analyze a large cross-section of countries. Generally, they find evidence in favor of predictability by financial ratios in some countries but not others and more robust results for the predictive ability of term structure variables. In recent work, Plazzi (2009) uses the present-value identity in a two-country predictability framework for the U.S. and the U.K. He finds evidence for increasing real and financial integration between the two economies.

5.4 Other Asset Classes

While the focus of this review is on aggregate stock returns, there is a large literature on predictability in the cross-section of stock returns as well as predictability of other asset classes. In the cross-section of stock returns, Cohen et al. (2003) show that the expected return difference between value stocks (with high book-to-market ratios) and growth stocks (with low book-to-market ratios) is especially high when the lagged book-to-market ratio difference between value and growth is high. The latter is high in recessions. Using the structure of a present-value model, they argue that the cross-sectional variance in book-to-market ratios mostly reflects variation in expected cash flow growth rather than variation in expected returns. This result stands in contrast to the typical time series result that most of the variation in valuation ratios reflects expected return variability.

Among other asset classes, government bond returns of various maturities are predictable by the term spread, e.g., the difference between the five year and the one year yield (Campbell and Shiller, 1991), or by a linear combination of forward rates Fama and Bliss (1987); Stambaugh (1988); Cochrane and Piazzesi (2005). In the language of the term structure literature, the latter paper is consistent with a price of level risk that varies with the Cochrane-Piazzesi factor (Cochrane and Piazzesi, 2008). In related work, Joslin, Priebsch, and Singleton (2010) find that the price of both level and slope risk moves around with real economic activity, e.g. industrial production. Kojien, Lustig, and Van Nieuwerburgh (2009) connect bond return predictability to the cross-section of stock returns by showing that value stocks' returns are more sensitive to innovations in the Cochrane-Piazzesi factor than growth stocks' returns. This sensitivity differential can be traced back to value stocks' dividend growth which falls much more in recessions than growth stocks' dividend growth. Such joint treatment of bond and stock returns suggests consistent risk pricing between both markets. The connection with the business cycle suggests that the value premium is consistent with rational asset pricing theories.

Like for bond markets, the expectations hypothesis has been rejected for foreign currency markets. High interest rate currencies do not appreciate as much as predicted by the EH theory, leading to return predictability.¹³ Bekaert and Hodrick (1992) show that dividend-price ratios, forward

¹³Time-varying risk premiums lead to an omitted variables bias in the EH tests, to the extent that the variable moving risk premia is correlated with interest rates. In contrast to the foreign exchange literature, the expectations hypothesis for bonds cannot be rejected for several countries (see Bekaert et al., 1997; Bekaert and Hodrick, 2001).

premiums, and lagged excess stock returns all predict currency returns. More recently, Lustig, Roussanov, and Verdelhan (2010) show that interest rates, exchange rates, and U.S. industrial production strongly predict returns on portfolios of currencies and that currency expected returns are counter-cyclical. Binsbergen, Koijen, and Neiman (2010) relate fluctuations in the net foreign asset position to currency return predictability. Several structural models have been shown to generate predictable (and counter-cyclical) currency returns (e.g., Bansal and Shaliastovich, 2010; Verdelhan, 2010; Gabaix and Farhi, 2009).

Present value models have been used to explore return predictability in residential real estate (Campbell et al., 2009) and commercial real estate (Plazzi et al., 2010). Commodity returns are also predictable by interest rates, the yield spread, the forward premium, and by capital flows into commodity markets measured as open interest (see Hong and Yogo, 2010, and the references therein). Finally, Brandt, Binsbergen, and Koijen (2010) study the risk-return properties of equity dividend strips. They find that the prices of short-term dividend strips are more volatile than their subsequent realization. This implies that the returns on short-term dividend strips are strongly predictable using the price of the dividend strip normalized by the dividend as predictor, which is analogue of the price-dividend ratio for the aggregate stock market.

6 Conclusion

In this paper, we review the literature on return and cash flow growth predictability from the perspective of the present value model. The present-value model allows for both time variation in expected returns and expected growth rates as drivers of stock price-dividend ratios and disentangles both sources of variation. The main conclusions of our paper can be summarized as follows:

1. Stock returns are less and dividend growth are more predictable over the full 1927-2009 sample than commonly believed, when using standard predictability regressions with the dividend-price ratio (dp) as predictor. This conclusion hinges on the dividend reinvestment assumption one makes to construct the dividend-price ratio and dividend growth. We have advocated reinvestment at the risk-free rate, not the stock market return.
2. In the post-1945 sample, these results reverse with no dividend growth and stronger return predictability, again using simple predictability regressions with the dp ratio as predictor.
3. The dp ratio displays near-unit-root behavior. Removing a low-frequency component from the dp ratio results in substantially stronger stock return and dividend growth predictability over the full 1927-2009 sample, using predictive regressions.

4. No predictability using the univariate regression approach does not imply no predictability. While standard dp regressions uncover about the same amount of return predictability as the present-value model, there is a huge discrepancy for dividend growth predictability. Standard regressions strongly understate the amount of dividend growth predictability in both the full sample and the post-war sample. In the latter sample they fail to detect predictability completely.
5. While dividend growth predictability is strong, it is short-lived. In contrast, return predictability is modest, but expected returns are persistent. As a result, about 90% of the variation in price-dividend ratios is due to variation in expected returns. Removing a low-frequency component in expected returns lowers this estimate to 64%.

These results all point to the importance of going beyond simple predictive regressions and considering the present-value identity. The same present-value framework has been used to study predictability in other asset markets. These results have immediate implications for those modeling the risk-return trade-off in asset markets. Such models should feature both time-varying dividend growth (at business cycle frequency) and time-varying expected returns (at generational frequency). While substantial progress has been made developing structural asset pricing models with these properties, no model we are aware of is quantitatively consistent all these conclusions. A further challenge is to understand better the economic reasons for dividend growth and return predictability. This ultimately requires studying firms' investment and financing policies more closely, as well as strengthening the link between asset prices and macroeconomic activity. Finally, future work should aim to integrate the results on predictability in equity, bond, currency, and real estate markets.

References

- Yakov Amihud and Clifford M. Hurwich. Predictive regressions: A reduced-bias estimation method. *Financial and Quantitative Analysis*, 39:813–841, 2004.
- Andrew Ang and Geert Bekaert. Stock return predictability: Is it there? *Review of Financial Studies*, 20 (3):651–707, 2007.
- Doron Avramov. Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64 (3):423–458, 2002.
- Malcolm Baker and Jeffrey Wurgler. The equity share in new issues and aggregate stock returns. *Journal of Finance*, 55:2219–2257, 2000.
- Ravi Bansal and Ivan Shaliastovich. A long-run risks explanation of predictability puzzles in bond and currency markets. Working Paper Duke University and Wharton School, February 2010.
- Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59:1481–1509, August 2004.
- Ravi Bansal, Ed Fang, and Amir Yaron. Equity capital: A puzzle? August 2005. Working Paper Duke University.
- Ravi Bansal, Dana Kiku, and Amir Yaron. An empirical evaluation of the long-run risks model for asset prices. Working Paper Duke University and the Wharton School, August 2009.
- Santiago Bazdreh, Frederico Belo, and Xiaojin Lin. Labor hiring, investment, and stock return predictability in the cross-section. Working Paper University of Minnesota, June 2009.
- Geert Bekaert and Robert Hodrick. Characterizing predictable components in excess returns. *Journal of Finance*, 47 (2):467–509, 1992.
- Geert Bekaert and Robert Hodrick. Expectation hypotheses tests. *Journal of Finance*, 56 (4):1357–1394, 2001.
- Geert Bekaert, Robert Hodrick, and David Marshall. On biases in tests of the expectation hypothesis of the term structure of interest rates. *Journal of Financial Economics*, 44:309–348, 1997.
- Jules van Binsbergen and Ralph Koijen. Predictive regressions: A present-value approach. *Journal of Finance*, 65 (4), July 2010.
- Jules van Binsbergen, Ralph Koijen, and Brent Neiman. What moves real exchange rates? Working Paper Chicago Booth, April 2010.

- Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. *Review of Financial Studies*, 22:4463–4492, 2009.
- Peter Bossaerts and Pierre Hillion. Implementing statistical criteria to select return forecasting models: What do we learn? *Review of Financial Studies*, 12:405–428, 1999.
- Jacob Boudoukh, Roni Michaely, Matthew Richardson, and Michael Roberts. On the importance of measuring payout yield: Implications for empirical asset pricing. *Journal of Finance*, 62:877–915, 2007.
- Jacob Boudoukh, Matthew Richardson, and Robert F. Whitelaw. The myth of long-horizon predictability. *Review of Financial Studies*, 21 (4):1577–1605, 2008. New York University Working Paper.
- Michael Brandt, Jules van Binsbergen, and Ralph Koijen. On the timing and pricing of cash flows. Working Paper Chicago Booth, November 2010.
- Michael W. Brandt and Qiang Kang. On the relation between the conditional mean and volatility of stock returns: A latent var approach. *Journal of Financial Economics*, 72:217–257, 2004.
- Laurent Calvet, Martín Gonzalez-Eiras, and Paolo Sodini. Financial innovation, market participation and asset prices. *Journal of Financial and Quantitative Analysis*, 39 (3):431–459, 2004.
- John Y. Campbell. Stock returns and the term structure. *Journal of Financial Economics*, 18:373–399, 1987.
- John Y. Campbell. A variance decomposition for stock returns. *Economic Journal*, 101:157–179, 1991.
- John Y. Campbell. Consumption-based asset pricing. In George Constantinides, Milton Harris, and Rene Stulz, editors, *Handbook of the Economics of Finance* forthcoming. North-Holland, Amsterdam, 2003.
- John Y. Campbell and John H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999.
- John Y. Campbell and Robert J. Shiller. Stock prices, earnings and expected dividends. *Journal of Finance*, 43:661–76, 1988a.
- John Y. Campbell and Robert J. Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1:195–227, 1988b.
- John Y. Campbell and Robert J. Shiller. Yield spreads and interest rate movements: A bird’s eye view. *Review of Economic Studies*, 58:495–514, 1991.
- John Y. Campbell and Robert J. Shiller. Valuation ratios and the long-run stock market outlook: An update. In Nicholas Barberis and Richard Thaler, editors, *Advances in Behavioral Finance, Volume II* forthcoming. Russell Sage Foundation, New York, NY, 2003.

- John Y. Campbell and Samuel Thompson. Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21 (4):1509–1531, 2008.
- John Y. Campbell and Motihiro Yogo. Efficient tests of stock return predictability. *Journal of Financial Economics*, 81:27–60, 2006.
- John Y. Campbell, Andrew W. Lo, and A. Craig MacKinlay. *The Econometrics of Financial Markets*. Princeton University Press, 1997.
- Sean Campbell, Morris Davis, Joshua Gallin, and Robert Martin. What moves housing markets: A variance decomposition of the rent-price ratio. *Journal of Urban Economics*, 66 (2):90–102, 2009. Working Paper, Federal Reserve Board of Governors.
- Long Chen. On the reversal of return and dividend predictability: A tale of two periods. *Journal of Financial Economics*, 92:128–151, 2009.
- Long Chen, Zhi Da, and Richard Priestley. Dividend smoothing and predictability. November 2010. Working Paper Washington University.
- John H. Cochrane. Production-based asset pricing and the link between stock returns and economic fluctuations. *Journal of Finance*, 46(1):209–37, March 1991a.
- John H. Cochrane. Explaining the variance of price-dividend ratios. *Review of Financial Studies*, 5(2): 243–280, June 1991b.
- John H. Cochrane. *Asset Pricing*. Princeton University Press, Princeton, N.J., revised edition edition, 2001.
- John H. Cochrane. Financial markets and the real economy. In *International Library of Critical Writings in Financial Economics*, volume 18. London: Edward Elgar, March 2006.
- John H. Cochrane. The dog that did not bark: A defense of return predictability. *Review of Financial Studies*, 21 (4):1533–1575, 2008.
- John H. Cochrane and Monika Piazzesi. Bond risk premia. *American Economic Review*, 95:138–160, 2005.
- John H. Cochrane and Monika Piazzesi. Decomposing the yield curve. Working Paper, University of Chicago, 2008.
- Randolph B. Cohen, Christopher Polk, and Tuomo Vuolteenaho. The value spread. *Journal of Finance*, 58 (2):609–641, 2003.
- George Constantinides and Anisha Ghosh. The predictability of returns with regime shifts in consumption and dividend growth. Working Paper Chicago Booth and Tepper School, July 2010.

- Martijn Cremers. Stock return predictability: A bayesian model selection perspective. *Review of Financial Studies*, 15 (4):1223–1249, 2002.
- T. Dangl and Michael Halling. Predictive regressions with time-varying coefficients. September 2009. Working Paper University of Utah.
- Jerome Detemple and Shashidhar Murthy. Intertemporal asset pricing with heterogeneous beliefs. *Journal of Economic Theory*, 62:294–320, 1994.
- Itamar Drechsler and Amir Yaron. What’s vol got to do with it. *Review of Financial Studies*, forthcoming, 2010.
- Pjotr Elias. Optimal median unbiased estimation of coefficients on highly persistent regressors. Unpublished paper, Department of Economics, Princeton University, 2005.
- Eugene F. Fama. The behavior of stock market prices. *Journal of Business*, 38:34–101, 1965.
- Eugene F. Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25:383–417, 1970.
- Eugene F. Fama. Efficient markets ii. *Journal of Finance*, 46(5):1575–1618, 1991.
- Eugene F. Fama and Robert H. Bliss. The information in long-maturity forward rates. *American Economic Review*, 77(4):680–692, September 1987.
- Eugene F. Fama and Kenneth R. French. Dividend yields and expected stock returns. *Journal of Financial Economics*, 22:3–27, 1988.
- Eugene F. Fama and Kenneth R. French. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, 25:23–49., 1989.
- Eugene F. Fama and Kenneth R. French. Disappearing dividends: Changing firm characteristics or lower propensity to pay. *Journal of Financial Economics*, 60(1):3–43, 2001.
- Eugene F. Fama and G. William Schwert. Asset returns and inflation. *Journal of Financial Economics*, 5 (2):115–146, 1977.
- Jack Favilukis, Sydney Ludvigson, and Stijn Van Nieuwerburgh. Macroeconomic implications of housing wealth, housing finance, and limited risk-sharing in general equilibrium. Working Paper New York University, May 2010.
- Wayne E. Ferson and Campbell Harvey. The risk and predictability of international equity returns. *Review of Financial Studies*, 6:527–566, 1993.

- Wayne E. Ferson, Sergei Sarkissian, and Timothy T. Simin. Spurious regressions in financial economics? *Journal of Finance*, 58(4):1393–1413, August 2003.
- Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro finance. Working Paper NYU Stern, March 2009.
- Xavier Gabaix and Emmanuel Farhi. Rare disasters and exchange rates. March 2009. Working Paper NYU and Harvard.
- William N. Goetzmann and Philippe Jorion. Testing the predictive power of dividend yields. *Journal of Finance*, 48:663–679, 1993.
- William N. Goetzmann and Philippe Jorion. A longer look at dividend yields. *Journal of Business*, 68(4):483–508, 1995.
- Amit Goyal and Ivo Welch. Predicting the equity premium with dividend ratios. *Management Science*, 49(5):639–654, May 2003.
- Amit Goyal and Ivo Welch. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008.
- Hui Guo and Robert F. Whitelaw. Uncovering the risk-return relation in the stock market. *Journal of Finance*, 61(3):1433–1463, 2006.
- Lars Peter Hansen, William Roberds, and Thomas J. Sargent. Time-series implications of present value budget balance and martingale models of consumption and taxes. In Lars P. Hansen and Thomas J. Sargent, editors, *Rational Expectations Econometrics*. Westview Press, 1991.
- Campbell R. Harvey. Predictable risk and returns in emerging markets. *Review of Financial Studies*, 8:773–816, 1995.
- David Hirshleifer, Kewei Hou, and Siew Hong Teoh. Accruals, cash flows, and aggregate stock returns. *Journal of Financial Economics*, 91:389–406, 2009.
- Erik Hjalmarsson. Predicting global stock returns. *Journal of Financial and Quantitative Analysis*, 45:49–80, 2010.
- Robert Hodrick. Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies*, 5:357–386, 1992.
- Harrison Hong and Motohiro Yogo. Commodity market capital flow and asset return predictability. Working Paper Princeton University and Minneapolis Federal Reserve, February 2010.
- Atsushi Inoue and Lutz Kilian. In-sample or out-of-sample tests of predictability: Which one should we use? *Econometric Reviews*, 23:371–402, 2004.

- Michael C. Jensen. Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, 6:95–101, 1978.
- Michael Johannes, Arthur Korteweg, and Nick Polson. Sequential learning, predictive regressions, and optimal portfolio returns. October 2008. Working Paper Columbia, Stanford, and Chicago.
- Scott Joslin, Michael Priebsch, and Kenneth Singleton. Risk premiums in dynamic term structure models with unspanned macro risks. April 2010. Working Paper Stanford University and MIT.
- Bryan Kelly and Seth Pruitt. Disaggregate valuation ratios and market expectations. Chicago Booth Working Paper, September 2010.
- Ralph Koijen and Stijn Van Nieuwerburgh. Financial economics, market efficiency and return predictability. In Robert Meyers, editor, *Encyclopedia of Complexity and System Science*, pages 3448–3456. Springer, 2009.
- Ralph Koijen, Hanno Lustig, and Stijn Van Nieuwerburgh. The cross-section and time-series of stock and bond returns. Working Paper New York University, September 2009.
- S. P. Kothari and Jay Shanken. Stock return variation and expected dividends : A time-series and cross-sectional analysis. *Journal of Financial Economics*, 31 (2):177–210, 1992.
- S. P. Kothari and Jay Shanken. Book-to-market, dividend yield, and expected market returns: A time series analysis. *Journal of Financial Economics*, 44:169–203, 1997.
- Owen Lamont. Earnings and expected returns. *Journal of Finance*, 53:1563–1587, October 1998.
- Owen Lamont. Investment plans and stock returns. *Journal of Finance*, 55(6):2719–45, December 2000.
- Borja Larrain and Motohiro Yogo. Does firm value move too much to be justified by subsequent changes in cash flow? *Journal of Financial Economics*, 87 (1):200–226, 2008.
- Martin Lettau and Sydney Ludvigson. Measuring and modeling variation in the risk-return tradeoff. In Yacine At-Sahalia and Lars Peter Hansen, editors, *Handbook of Financial Econometrics*, volume 1, pages 617–690. Elsevier Science B.V., North Holland, Amsterdam, March 2010. New York University Working Paper.
- Martin Lettau and Sydney C. Ludvigson. Expected returns and expected dividend growth. *Journal of Financial Economics*, 76:583–626, 2005.
- Martin Lettau and Sydney C. Ludvigson. Consumption, aggregate wealth and expected stock returns. *The Journal of Finance*, 56(3):815–849, June 2001.
- Martin Lettau and Stijn Van Nieuwerburgh. Reconciling the return predictability evidence. *Review of Financial Studies*, 21 (4):1607–1652, 2008.

- Martin Lettau, Sydney Ludvigson, and Jessica Wachter. The declining equity premium: What role does macroeconomic risk play? *Review of Financial Studies*, 21:1653–1687, 2008.
- Jonathan W. Lewellen. Predicting returns with financial ratios. *Journal of Financial Economics*, 74(2): 209–235, 2004.
- Sydney C. Ludvigson and Serena Ng. The empirical risk-return relation: A factor analysis approach. *Journal of Financial Economics*, 83 (1):171–222, 2007. Unpublished manuscript.
- Hanno Lustig and Stijn Van Nieuwerburgh. Housing collateral, consumption insurance and risk premia: An empirical perspective. *Journal of Finance*, 60(3):1167–1219, 2005.
- Hanno Lustig and Stijn Van Nieuwerburgh. Can housing collateral explain long-run swings in asset returns? Working Paper NYU Stern and UCLA, June 2007.
- Hanno Lustig and Stijn Van Nieuwerburgh. The returns on human capital: Good news on wall street is bad news on main street. *Review of Financial Studies*, 21 (5):2097–2137, 2008.
- Hanno Lustig and Stijn Van Nieuwerburgh. How much does household collateral constrain regional risk sharing. *Review of Economic Dynamics*, Forthcoming, 2009.
- Hanno Lustig, Nick Roussanov, and Adrien Verdelhan. Countercyclical currency risk premia. Working Paper UCLA Anderson, Wharton School, and MIT Sloan, August 2010.
- Ellen R. McGrattan and Edward C. Prescott. Taxes, regulations, and the value of u.s. and u.k. corporations. *Review of Economic Studies*, 72, 2005.
- Lior Menzly, Tano Santos, and Pietro Veronesi. Understanding predictability. *Journal of Political Economy*, 112(1):1–47, February 2004.
- Charles C. Nelson and Myung J. Kim. Predictable stock returns: The role of small sample bias. *Journal of Finance*, 43:641–661, 1993.
- Lubos Pastor and Robert F. Stambaugh. The equity premium and structural breaks. *Journal of Finance*, 56(4):1207–1239, 2001.
- Lubos Pastor and Robert F. Stambaugh. Predictive systems: Living with imperfect predictors. *Journal of Finance*, 64 (4):1583–1628, 2009.
- Lubos Pastor, Meenakshi Sinha, and Bhaskaran Swaminathan. Estimating the intertemporal risk-return tradeoff using the implied cost of capital. *Journal of Finance*, 63:2859–2897, 2008.
- Bradley S. Paye and Allan Timmermann. Instability of return prediction models. *Journal of Empirical Finance*, 13:274–315, 2006. Forthcoming *Journal of Empirical Finance*.

- Monika Piazzesi, Martin Schneider, and Selale Tuzel. Housing, consumption and asset pricing. *Journal of Financial Economics*, 83(3):531–569, 2007.
- Alberto Plazzi. International stock return correlation: real or financial integration? a structural present value approach. Mimeo UCLA Anderson School of Management, 2009.
- Alberto Plazzi, Walter Torous, and Rossen Valkanov. Expected returns and expected growth in rents of commercial real estate. *Review of Financial Studies*, Forthcoming, 2010.
- Jeff Pontiff and L. D. Schall. Book-to-market ratios as predictors of market returns. *Journal of Financial Economics*, 49:141–160, 1998.
- Vincenzo Quadrini and Urban Jermann. Stock market boom and the productivity gains of the 1990s. *Journal of Monetary Economics*, pages 413–432, 2007. Unpublished paper, University of Southern California.
- Michael S. Rozeff. Dividend yields are equity risk premia. *Journal of Portfolio Management*, 49:141–160, 1984.
- Robert J. Shiller. Stock prices and social dynamics. *Brookings Papers on Economic Activity*, 2:457–510, 1984.
- Matthew Spiegel. Forecasting the equity premium: Where we stand today. *Review of Financial Studies*, 21 (4):1453–1454, 2008.
- Robert F. Stambaugh. The information in forward rates: Implications for models of the term structure. *Journal of Financial Economics*, pages 41–70, 1988.
- Robert F. Stambaugh. Predictive regressions. *Journal of Financial Economics*, 54:375–421, 1999.
- Lawrence H. Summers. Does the stock market rationally reflect fundamental values? *Journal of Finance*, 41:591–601, 1986.
- Allan G. Timmermann. How learning in financial markets generates excess volatility and predictability of stock returns. *Quarterly Journal of Economics*, 108:1135–1145, 1993.
- Walter Torous, Rossen Volkanov, and Shu Yan. On predicting returns with nearly integrated explanatory variables. *Journal of Business*, 77:937–966, 2004.
- Rossen Valkanov. Long-horizon regressions: Theoretical results and applications. *Journal of Financial Economics*, 68:201–232, 2003.
- Adrien Verdelhan. A habit-based explanation of the exchange rate risk premium. *Journal of Finance*, 65 (1):123–145, October 2010.

Luis Viceira. Testing for structural change in the predictability of asset returns. Unpublished manuscript, Harvard University, 1996.

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

Jessica Wachter. Asset allocation. *Annual Review of Financial Economics*, 2, 2010.

Lu Zhang. The value premium. *Journal of Finance*, 60 (1):67–103, 2005.