Expected Returns and Dividend Growth Rates Implied in Derivative Markets

Benjamin Golež[*]
Universitat Pompeu Fabra

JOB MARKET PAPER

January, 2011[†]

Abstract

I show that the dividend growth implied in S&P 500 options and futures predicts changes in dividends and thereby improves the forecasts of market returns. Guided by a simple present value model, I use the implied dividend growth to correct the standard dividend–price ratio (DP) for variation in expected dividend growth. I find that the corrected DP predicts S&P 500 returns in the period 1994–2009 significantly better than does the uncorrected DP. This predictive improvement is especially pronounced over the monthly horizon, holds both in-sample and out-of-sample, yields a sizable gain in the Sharpe ratio, and is robust to small sample bias. The results indicate that expected returns and expected dividend growth are highly correlated.

Keywords: present value models, dividend-price ratio, return predictability, options, futures, implied dividend growth

JEL classification: G12, G13, G14, G17, C22, C53

[*]Department of Economics and Business, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain. Tel. +34-93-542-2670. Email: benjamin.golez@upf.edu.

[†]I would like to thank Mascia Bedendo, Geert Bekaert, Francesco Corielli, Martijn Cremers, Jens Carsten Jackwerth, Ralph Koijen, Gueorgui I. Kolev, Peter Koudijs, José M. Marín, Francisco Peñaranda, Christopher Polk, Jesper Rangvid, and Robert Zymek for invaluable comments. I also thank seminar participants at Universitat Pompeu Fabra, Swiss Banking Institute, and Bocconi University as well as participants at 2010 Eastern Finance Association, Fall 2009 Chicago Quantitative Alliance Conference, 2009 European Finance Association Doctoral Tutorial, 12th Conference of the Swiss Society for Financial Market Research, and 2009 Campus for Finance. Previous versions of the paper were circulated under the title "Options Implied Dividend Yield and Market Returns". I gratefully acknowledge financial support from the Spanish Ministry of Education and Innovation.

1 Introduction

The predictability of market returns is of great interest to market practitioners and has important implications for asset pricing. However, there is still no consensus on whether returns are predictable. Although many studies argue that returns can be predicted by price multiples such as the dividend–price ratio (Fama and French, 1988; Lewellen, 2004; Cochrane, 2008a), others document that predictability is subject to statistical biases and is difficult to exploit for purposes of portfolio allocation (Stambaugh, 1999; Goyal and Welch, 2008).

In this paper I reexamine the role of dividend ratios for predicting market returns. I argue that the poor performance of the dividend-price ratio (DP) in predicting returns is largely due to the time-varying nature of the expected dividend growth. I introduce a novel proxy for expected dividend growth, which is extracted from index options and futures, and derive a simple present value model to guide the empirical analysis. Using the dividend growth implied in derivative markets to correct the DP for variation in expected dividend growth, I find that short-term market returns are strongly predictable. Indeed, the corrected DP predicts monthly market returns both in-sample and out-of-sample, and it is also robust to the statistical biases that have been shown to hinder the predictive ability of the uncorrected DP.

The insight that the time-varying expected dividend growth can reduce the ability of the DP to predict returns has long been part of the predictability literature (Campbell and Shiller, 1988; Fama and French, 1988). According to the textbook treatment, the DP varies over time not only because of changes in expected returns but also because of changes in expected dividend growth. Therefore, as pointed out by Fama and French (1988), the DP is only a noisy proxy for expected returns in the presence of time-varying expected dividend growth (see also Cochrane, 2008a; Rytchkov, 2008; Binsbergen and Koijen, 2010). Moreover, since the DP increases with expected returns and decreases with expected dividend growth, the problems caused by time-varying expected dividend growth are pronounced when expected returns and expected dividend growth are positively correlated (Menzly et al., 2004; Lettau and Ludvigson, 2005).[1] This positive correlation offsets the changes in expected returns and those in expected dividend growth, which further reduces the DP's ability to predict returns.

Thus, if our task is to predict returns, then the DP is insufficient: We must also account for the time-varying value of expected dividend growth. Yet this value is difficult to estimate because it aggregates investors' expectations about future growth opportunities. Recent studies on return predictability typically assume that the future will be similar to the past and then go on to extract expected dividend growth from historical data. For example, Binsbergen and Koijen (2010) take a latent variable approach within the present value model to filter out both expected returns and expected dividend growth from the history of dividends and prices (see also Rytchkov, 2008). Lacerda and Santa-Clara (2010) use a simple average of historical dividend growth as a proxy for expected dividend growth. These authors all conclude that improved prediction of dividend growth will, in turn, improve the predictability of longer-term (i.e., annual) returns. Nevertheless, their methods exploit only the information that can be derived from past dividends and prices. In contrast, investors base expectations about future cash flows on a much richer—and forward-looking—information set.

This paper takes a different approach to estimating expected dividend growth. Instead of relying on historical data, I extract a proxy for investors' expected dividend growth from derivative markets (index options and index futures). Prices of options and futures depend on, inter alia, the dividends that the underlying asset pays until the expiration of the contracts. Therefore, derivative markets provide us with a unique laboratory for estimating the dividends that investors expect to realize in the near future. Because index derivatives are highly liquid, new information about future cash flows is rapidly incorporated into the estimated implied dividends. For this reason, implied dividends are particularly well suited for revealing expectations over short horizons, where the constant flow of information causes rapid changes in investors' expectations regarding future dividends and returns.

[1] Menzly et al. (2004) show that a positive correlation between expected returns and expected dividend growth arises (in a general equilibrium model) as a natural consequence of dividend growth predictability.

To provide an analytical framework for the empirical analysis, I first derive a simple present value model. Like Binsbergen and Koijen (2010), I combine the Campbell and Shiller (1988) present value identity with a simple, first-order autoregressive process for the expected return and the expected dividend growth. In this environment, the future return is a function of the DP and the expected dividend growth, where both terms enter linearly. We can therefore consider predicting returns through a multivariate regression of returns on the DP and an estimate for the expected dividend growth, or we can combine them in a single predictor—the so-called corrected DP. The corrected DP can be interpreted as the dividend–price ratio adjusted for variation in expected dividend growth.

Following the implications of the present value model, I proceed with estimating the proposed proxy for the expected dividend growth. To extract the dividend growth implied in index options and index futures, I first estimate an implied dividend yield. By combining the no-arbitrage, cost-of-carry formula for index futures and the put–call parity condition for index options, I derive an expression that enables estimation of the implied dividend yield in a model-free way, and solely in terms of the observed prices of derivatives and their underlying asset. Once estimated, I combine the implied dividend yield with the realized DP to calculate the implied dividend growth and the corrected DP.

I apply the empirical analysis to the S&P 500 index. Given the requirement for data on both options and futures, the analysis is restricted to the period from January 1994 through December 2009.[2] The main results can be summarized as follows. Consistent with previous studies, I find that the standard DP is a rather poor predictor of both future returns and dividend growth. The predictive coefficients on the DP are insignificant in all the forecasting regressions for horizons ranging from one to six months. In contrast, the implied dividend growth reliably predicts dividend growth for all the considered horizons. In line with this

[2]Notice that the post 1994 period is not affected by the breaks in the mean of the DP, which have been shown to affect the forecasting relationship of returns and the DP over longer periods of time (Lettau and Nieuwerburgh, 2008; Favero et al., 2010).

observation, the ability to predict market returns improves considerably when implied dividend growth is included as an additional regressor in the standard DP regression for predicting returns. Furthermore, the results confirm that the DP and the implied dividend growth can be replaced by a single predictor: the corrected DP. The predictive coefficient on the corrected DP is statistically significant for all the considered return horizons. The improvement in the predictability is especially strong for short time horizons. In the predictive regressions with monthly returns, the corrected DP exhibits an in-sample adjusted R^2 of 4.61% and an out-of-sample R^2_{OS} of 6.06%, as compared with 0.33% and -0.15% (respectively) for the uncorrected DP. For a mean-variance investor, the documented improvement in predicting returns translates into a gain of 0.32 in terms of the Sharpe ratio. Since the corrected DP is less persistent than the uncorrected DP and since innovations to the corrected DP are only weakly related to returns, the corrected DP has the additional advantage of being robust to small sample bias that that has been shown to hinder the predictive ability of the uncorrected DP. Furthermore, the documented improvement in predictive accuracy is not due to duplication by implied dividend growth of information embedded within other options-implied predictors such as variance risk premia (Bollerslev et al., 2009) and cannot be replicated by using historical dividend growth in place of implied dividend growth.

Consistent with the empirical results, a variance decomposition of the DP reveals considerable variation in both expected returns and expected dividend growth. However, like Lettau and Ludvigson (2005), I find that expected returns and expected dividend growth are highly correlated (0.88). This high correlation means that movements in expected returns and expected dividend growth offset each other's effect in the DP, which renders the DP relatively smooth. Correcting the DP for the implied dividend growth restores the variation that is offset by this strong comovement, and thus implies that expected returns vary significantly more than is suggested by variation in the uncorrected dividend–price ratio.

The paper draws upon a large number of studies in the predictability literature and is also related to other papers using implied dividends. Dividends implied in derivative markets have been used as an input in the calculation of risk-neutral densities (Ait-Sahalia and Lo, 1998), and to study empirical properties of dividend strips (Binsbergen, Brandt and Koijen, 2010). However, this paper is the first to employ implied dividends for the purpose of predicting market returns. I also use a new techinique which enables me to extract dividends from derivative prices without resorting to the use of proxies for the implied interest rate. This is important as interest rates implied in derivative markets may differ from observable interest rates (Naranjo, 2009).

The rest of the paper is organized as follows. Section 2 derives the present value model. Section 3 details the technique proposed to extract the dividend growth that is implied in the market for derivatives. Section 4 presents the data, and Section 5 reports on the results of predictive regressions involving dividend growth and market returns. Section 6 considers additional statistical tests and compares the documented predictability with alternative predictors. Section 7 presents a variance decomposition of the dividend–price ratio, and Section 8 is devoted to robustness checks. Section 9 concludes the paper.

2 Present value model

To provide an analytical framework for the empirical analysis, this section derives a simple log-linear present value model. The model combines the Campbell and Shiller (1988) present value identity with AR(1) processes for expected returns and expected dividend growth rates. A similar approach is used in Binsbergen and Koijen (2010) and Rytchkov (2008).[3,4] The main innovation of this study lies in the empirical estimation of this setup. I use the present value model mainly to motivate the return predictive regressions.

[3]The AR(1) structure is motivated by growing evidence that both expected returns and expected dividend growth rates are time-varying and persistent (Menzly et al., 2004; Lettau and Ludvigson, 2005; Bansal and Yaron, 2004).

[4]Present value models with different processes for expected returns and expected dividend growth are extensively analyzed in Cochrane (2008b).

Define log return r_{t+1}, log dividend growth Δd_{t+1}, and log dividend-price ratio dp_t as:

$$r_{t+1} = \log\left[\frac{P_{t+1}+D_{t+1}}{P_t}\right], \quad \Delta d_{t+1} = \log\left[\frac{D_{t+1}}{D_t}\right], \quad dp_t = \log\left[\frac{D_t}{P_t}\right] \tag{1}$$

Rewrite returns as in Campbell and Shiller (1988):

$$r_{t+1} \simeq \kappa + dp_t + \Delta d_{t+1} - \rho dp_{t+1} \tag{2}$$

where $\rho = \frac{\exp(-\overline{dp})}{1+\exp(-\overline{dp})}$ and $\kappa = \log\left[1+\exp(-\overline{dp})\right] + \rho\overline{dp}$ are constants related to the long-run average of the dividend-price ratio, $\overline{dp}$. Iterate (2) forward to obtain the Campbell and Shiller (1988) present value identity:

$$dp_t \simeq -\frac{\kappa}{1-\rho} + E_t\sum_{j=0}^{\infty}\rho^j(r_{t+1+j}) - E_t\sum_{j=0}^{\infty}\rho^j(\Delta d_{t+1+j}) \tag{3}$$

Let $\mu_t = E_t(r_{t+1})$ be the conditional expected return and let $g_t = E_t(\Delta d_{t+1})$ be the conditional expected dividend growth. Suppose that μ_t and g_t follow $AR(1)$ processes:

$$\mu_{t+1} = \delta_0 + \delta_1(\mu_t) + \varepsilon^{\mu}_{t+1} \tag{4}$$

$$g_{t+1} = \gamma_0 + \gamma_1(g_t) + \varepsilon^{g}_{t+1} \tag{5}$$

$$\Delta d_{t+1} = g_t + \varepsilon^{d}_{t+1} \tag{6}$$

where ε^{μ}_{t+1}, ε^{g}_{t+1} and ε^{d}_{t+1} are zero mean errors. Combine the present value identity in (3) with the $AR(1)$ assumptions to find the dividend-price ratio:

$$dp_t \simeq \varphi + \left(\frac{1}{1-\rho\delta_1}\right)\mu_t - \left(\frac{1}{1-\rho\gamma_1}\right)g_t \tag{7}$$

where φ is a constant related to $\kappa, \rho, \delta_0, \delta_1, \gamma_0, \gamma_1$ (details are provided in Appendix).

Equation (7) states that the log dividend-price ratio is related to expected returns and is therefore a good candidate for predicting future returns. However, according to (7), dp_t also contains information about expected dividend growth. Hence, if expected dividend growth varies over time, the dp_t is only a noisy proxy for expected returns and an imperfect predictor for future returns (Fama and French, 1988; Binsbergen and Koijen, 2010; Rytchkov, 2008; Lacerda and Santa-Clara, 2010). Since the dp_t increases with expected returns and decreases with expected dividend growth, the problem is pronounced when expected returns and expected dividend growth are positively correlated (Menzly et al., 2004; Lettau and Ludvigson, 2005). This positive correlation offsets the changes in expected returns and those in expected dividend growth, which further reduces the ability of the dp_t to predict returns.

Thus, if our task is to predict returns, then the dp_t is insufficient: We must also account for the time-varying value of expected dividend growth. To see this formally, combine (2), (6) and (7) to obtain a return forecasting equation:

$$
\begin{aligned}
r_{t+1} &\simeq \kappa + dp_t + \Delta d_{t+1} - \rho dp_{t+1} && (8) \\
&\simeq \psi + (1-\rho\delta_1)dp_t + \left(\frac{1-\rho\delta_1}{1-\rho\gamma_1}\right) g_t + v^r_{t+1} && (9)
\end{aligned}
$$

where $v^r_{t+1} = \varepsilon^d_{t+1} - \rho\left(\frac{\varepsilon^\mu_{t+1}1}{1-\rho\delta_1} - \frac{\varepsilon^g_{t+1}}{1-\rho\gamma_1}\right)$ and ψ is a constant related to $\kappa, \rho, \delta_0, \delta_1, \gamma_0, \gamma_1$.

In line with the above argument, equation (8) reveals that, if our task is to predict returns, we need both dp_t and an estimate for expected dividend growth.

Since dp_t and the expected dividend growth are linearly related to future returns, we can also replace them by a single predictor:

$$r_{t+1} \simeq \psi + (1-\rho\delta_1)dp_t + \left(\frac{1-\rho\delta_1}{1-\rho\gamma_1}\right) g_t + v^r_{t+1} \tag{10}$$

$$\simeq \psi + (1-\rho\delta_1)\left[dp_t + g_t\left(\frac{1}{1-\rho\gamma_1}\right)\right] + v^r_{t+1} \tag{11}$$

$$\simeq \psi + (1-\rho\delta_1)dp^{Corr}_t + v^r_{t+1} \tag{12}$$

where $dp^{Corr}_t = dp_t + g_t\left(\frac{1}{1-\rho\gamma_1}\right)$ is the corrected dividend-price ratio and can be interpreted as the dividend-price ratio that is adjusted for variation in the expected dividend growth. The corrected dividend-price ratio depends on the dp_t, the expected dividend growth, the linearization constant and the persistence of the expected dividend growth.[5,6]

3 Estimating implied dividend growth

The present value model outlined in the previous section implies that the dividend-price ratio is not enough to capture variation in expected returns. Additionally, we need an estimate for the expected dividend growth.

In this study, I propose extracting investors' expected dividend growth from derivative markets (index options and index futures). Prices of options and futures depend on, inter alia, the dividends that the underlying asset pays until the expiration of the contracts. Therefore, we can invert the pricing relations to extract a proxy for expected dividend growth from the

[5] The fact that correction depends on the persistence of the expected dividend growth is an interesting insight since persistence of the expected dividend growth is one of the driving forces of the return predictability in the long-run risk models pioneered by Bansal and Yaron (2004).

[6] Lacerda and Santa-Clara (2010) derive a similar correction for the adjusted dividend-price ratio:

$$dp^{Adj.}_t = dp_t + \overline{g}_t\left(\frac{1}{1-\rho_t}\right)$$

In their version, the adjusted dividend-price ratio ($dp^{Adj.}_t$) does not depend on the persistence of the expected dividend growth because they assume that expected dividend growth is equal to the average historical dividend growth ($\overline{g}_t$).

observable prices of derivatives.

I employ a two step approach to estimating the implied dividend growth. In the first step, I extract an implied dividend yield embedded in derivative markets. In the second step, I combine the estimated implied dividend yield with the realized dividend-price ratio to calculate the implied dividend growth.

Below, I describe the proposed method for the estimation of the implied dividend yield. Transition from the implied dividend yield to the implied dividend growth is presented along with the estimation of the realized dividend-price ratio in the next section.

Implied dividend yield. To express the implied dividend yield in terms of the observable prices of derivatives, I combine two well-known no-arbitrage conditions, the cost-of-carry formula for index futures and the put-call parity condition for index options.

Under a standard assumption that the index pays a continuously compounded dividend yield (λ), the cost-of-carry formula for the future price is:[7]

$$F_t(\tau) = S_t \exp\left[(r_t(\tau) - \lambda_t(\tau))\tau\right] \tag{13}$$

where F_t is the future's price, S_t is the price of the underlying, $\lambda_t(\tau)$ is the annualized continuously compounded dividend yield between t and $t+\tau$ and $r_t(\tau)$ is the annualized continuously compounded interest rate from t to $t+\tau$.

Similarly, by no-arbitrage, the difference between a European call and a European put written on the index can be expressed as:

$$C_t(K,\tau) - P_t(K,\tau) = S_t \exp\left[-\lambda_t(\tau)\tau\right] - K \exp\left[-r_t(\tau)\tau\right] \tag{14}$$

where $C_t(K,\tau)$ and $P_t(K,\tau)$ are the prices of a European call and a European put option with

[7] For simplicity I do not consider any convexity adjustment for the stochastic dividend yield. See Lioui (2006) for the derivation of the put-call parity under the stochastic dividend yield.

the same maturity τ and the same strike price K.

Both no-arbitrage conditions relate prices of derivatives to the future dividend yield and the risk-free rate. Hence, we can combine them to first solve for the interest rate implied in the derivative markets:

$$r_t(\tau) = \frac{1}{\tau} \ln \left[\frac{F_t(\tau) - K}{C_t(K,\tau) - P_t(K,\tau)} \right] \tag{15}$$

Once we have an expression for the implied interest rate, we can plug it back in (14) to obtain an expression for the implied dividend yield:

$$\lambda_t(\tau) = -\frac{1}{\tau} \ln \left[\left(\frac{C_t(K,\tau) - P_t(K,\tau)}{S_t} \right) + \frac{K}{S_t} \left(\frac{C_t(K,\tau) - P_t(K,\tau)}{F_t(\tau) - K} \right) \right] \tag{16}$$

Equation (16) relates implied dividend yield to the observable market prices and enables us to estimate the implied dividend yield using only information that is available at time t. All we need is a European call option and a European put option with the same strike and the same maturity, the future price with the same expiration date as the options, and the price of the underlying.

It is important to note that the expression for the implied dividend yield is derived from no-arbitrage conditions. As such, it is free of any parametric options (and futures) pricing models and enables us to estimate the implied dividend yield in a model-free way. Also, the combination of two no-arbitrage conditions allows us to substitute the interest rate and estimate the implied dividend yield without resorting to the use of proxies for the implied interest rate. This is important because the implied interest rate may deviate from the observable proxies for the interest rate (Naranjo, 2009).

4 Data

I use the S&P 500 index as a proxy for the aggregate market. The S&P 500 price index and total return index (dividends reinvested) are downloaded from Datastream. The S&P 500 futures data comes from Chicago Merchandile Exchange and the S&P 500 options data is obtained from Market Data Express.

Futures on S&P 500 have been traded since April 1982 and European options on S&P 500 have existed since April 1986. However, Market Data Express options data only goes back to January 1990. Also, until 1994, the settlement procedure for S&P 500 options and futures differed. While futures are settled in the opening value of the index since June 1987, the most liquid S&P 500 options expired in the closing value of the index until December 1993.[8] Since liquid options and futures with matching expiration times are needed to estimate the implied dividend growth, I further restrict the analysis to the period from January 1994 through December 2009. The analysis is based on end-of-month observations.

In some parts of the paper I also make use of other variables. In particular, I download constant maturity 3-month and 6-month Treasury yield from the Federal Reserve Bank of St. Louis and I obtain the S&P 500 earnings-price ratio and the 6-month LIBOR rate from Datastream. Additionally, I obtain the implied variance index (VIX) and the variance risk premia from Hao Zhou's homepage. Finally, I download the consumption-to-wealth ratio from Sydney C. Ludvigson's website.

[8] When S&P 500 futures and S&P 500 options were introduced, they initially expired in the closing value of the index (P.M. settlement). In 1987, the Chicago Merchandile Exchange (CME) changed the expiration procedure of S&P 500 futures from the P.M. settlement to the A.M. settlement (A.M. settlement value is based on the opening prices of the index constituents on the expiration date). As a response, the Chicago Board of Options Exchange (CBOE) introduced a new version of its S&P 500 options that also settle A.M. However, the P.M. settled options remained the most liquid and the A.M. settled options were initially hardly traded. In 1992, CBOE decided that all the S&P 500 options should expire A.M. Since long dated P.M. settled options were already traded on the market, it took until December 1993 before all the traded S&P 500 options became A.M. settled.

4.1 Empirical estimation

Implied dividend yield. I estimate the implied dividend yield at the end of each month according to (16). I use daily settlement prices for futures, mid-point between the last bid and the last ask price for options and closing values for the S&P 500 price index.[9]

It is well-known that no-arbitrage conditions hold well for the S&P 500 index (Kamara and Miller, 1995). Still, due to market frictions (transaction costs and demand imbalances), particular pairs of options and futures may violate no-arbitrage conditions. To take this into account, I calculate the implied dividend yield by aggregating information from a wide set of options and futures.

For each end of the month, I use 10 days of backward-looking data and I construct option pairs (put-call pairs with the same strike and the same maturity) from all the reliable options (options with positive volume or open interest greater than 200 contracts).[10] Then I combine option pairs with the futures of matching maturity and the current value of the underlying index. To eliminate some extreme observations, I discard observations where $\frac{C_t(K,\tau)-P_t(K,\tau)}{F_t(\tau)-K}$ is smaller than 0.5 or greater than 1.5 (and where $F_t(\tau) = K$).[11] Using this data, I obtain several estimates for the implied dividend yields at the end of each month, which I aggregate into a single market's implied dividend yield by taking the median across all the implied dividend yields with the same maturity.

Since within year dividends exhibit seasonality, the common approach in the predictability literature is to calculate the dividend-price ratio by aggregating dividends over one year. In line with this literature, the implied dividend yield should ideally be estimated using options and futures with one year to expiration. However, long maturity derivatives are illiquid. As

[9] Market Data Express end-of-day data covers all the options written on the S&P 500 index, including mini options, quarterlies, weeklies and long-dated options. With the kind help of Market Data Express support team, I first eliminated all but standard S&P 500 options. Additionally, I imposed the standard filters to eliminate missing observations and options that violate the basic no-arbitrage bounds.

[10] Note: the formula for the implied dividend yield holds for all the moneyness levels. Unreported results show that there is no strike price effect, i.e. the implied dividend yield does not depend on the moneyness level.

[11] This filter eliminates a bit less than 2% of observations.

illustrated in Figure 1, open interest concentrates strongly on near to maturity options and futures. The tilt towards short maturities is especially pronounced for futures, for which there is almost no open interest for maturities above 9 months. For this reason, we cannot reliably estimate the implied dividend yield with the maturity of one year and we have to resort to the use of options and futures with shorter expiration dates. This may, nevertheless, introduce some seasonality into the estimated implied dividend yield.

[Insert Figure 1 about here]

To examine the effect of the seasonality in dividend payments on the implied dividend yield, I first estimate the whole term structure of the implied dividend yields. Since there are only four dates per year when options and futures expire simultaneously (third Friday in March, June, September and December),[12] I proceed as follows. In January, April, July and October, I extract the implied dividend yield for the maturities of 2, 5, or 8 months. In February, May, August and November, I extract the implied dividend yield for the maturities of 1, 4 or 7 months. Finally, in March, June, September and December, I estimate the implied dividend yield for the maturities of 3, 6 or 9 months. Then I linearly interpolate the estimated yields to obtain the term structure of the implied dividend yields with constant maturities (between 3 and 7 months).

Table I presents the summary statistics for the implied dividend yields with different maturities. All the yields have approximately the same mean, but differ with respect to their volatility. As expected, due to the seasonality in dividend payments, implied dividend yields with short maturities (3 and 4 months) are the most volatile. With the increase of the maturity, the volatility of the implied dividend yields first decays and then stabilizes, so that implied dividend yields with 6 and 7 months to maturity exhibit approximately the same volatility (see

[12] Options expire on a monthly cycle (third Friday in a month) and futures expire on a quarterly cycle (third Friday in March, June, September and December).

also Figure 2). This suggests that the problem of seasonality in dividend payments is largely diminished for the implied dividend yield with a maturity of at least 6 months. Given these results, I choose to conduct the main analysis using the implied dividend yield with maturity of 6 months.

By construction, the estimated implied dividend yield is continuously compounded. To make it comparable with the realized dividend-price ratio, I transform it into a raw (effective) implied dividend yield, $IDY_t = \exp(\widehat{\lambda_t}) - 1$. The log implied dividend yield is simply $idy_t = log(IDY_t)$.

[Insert Table I about here]

[Insert Figure 2 about here]

Market returns, dividend growth and dividend-price ratio. I follow the standard definitions for the realized variables. Monthly returns are defined as:

$$r_t^M = \log\left[\frac{P_t + D_t}{P_{t-1}}\right] \tag{17}$$

where P_t and D_t denote the price and dividends in month t. The dividend-price ratio is calculated by aggregating dividends over one year:

$$dp_t = \log\left[DP_t\right] = \log\left[\frac{D_t^{12}}{P_t}\right] \tag{18}$$

where D_t^{12} is the sum of dividends over the last 12 months. Monthly dividend growth is defined as in Ang and Bekaert (2007):

$$\Delta d_t^M = \log\left[\frac{D_t^{12}}{D_{t-1}^{12}}\right] \tag{19}$$

All the ratios are calculated from the S&P 500 price index and the total return index downloaded from Datastream. Since Datastream calculates the total return index by reinvesting dividends daily, I first extract the daily amount of dividends. Then I calculate D_t and D_t^{12} by summing dividends over the past month and year, respectively.

Implied dividend growth and the corrected dividend-price ratio. Based on the implied dividend yield and the dividend-price ratio, I calculate the implied dividend growth (idg) and the corrected dividend-price ratio (dp_t^{Corr}) as:

$$idg_t = \log\left[\frac{IDY_t}{DP_t}\right] = idy_t - dp_t \tag{20}$$

$$dp_t^{Corr} = dp_t + \left(\frac{1}{1-\widehat{\rho}\widehat{\gamma}_1}\right) idg_t \tag{21}$$

where $\widehat{\rho}$ is the estimated linearization constant and $\widehat{\gamma}_1$ is the $AR(1)$ coefficient of the implied dividend growth.

4.2 Data description

Table II reports the summary statistics for the variables sampled monthly. All the variables are annualized and expressed in logs. Returns and dividend growth rates are on average 7.33% and 3.61%, respectively.

The proxy for the expected dividend growth (implied dividend growth) is on average somewhat higher than the realized dividend growth rate (6.02%) and it nicely reflects market conditions. As shown in Figure 3, the implied dividend growth is positive during the market booms (1994-97 and 2002-2007), when investors were optimistic about future growth opportunities, and it is negative in times of stock market busts, such as in 1998 (Asian-Russian-LTCM crisis), in 2001 (dot.com bubble burst), and in 2008/2009 (the recent financial crisis), when investors were rather pessimistic about growth opportunities. The implied dividend growth is also relatively persistent. It exhibits a first order autocorrelation coefficient of 0.53 and it thereby justifies modeling expected dividend growth rate as a persistent process.

[Insert Table II about here]

[Insert Figure 3 about here]

The corrected dividend-price ratio is calculated as:

$$
\begin{aligned}
dp_t^{Corr} &= dp_t + \left(\frac{1}{1-\widehat{\rho}\widehat{\gamma_1}}\right) idg_t = dp_t + \left(\frac{1}{1-(0.98 * 0.53)}\right) idg_t && (22) \\
&= dp_t + 2.08 * idg_t. && (23)
\end{aligned}
$$

where $\widehat{\rho} = \frac{\exp(-\overline{dp})}{1+\exp(-\overline{dp})} = \frac{\exp(-4.03)}{1+\exp(-4.03)} = 0.98.$[13]

Figure 4 plots dp_t^{Corr} along with the dp_t. Both dividend ratios exhibit strong comovement (pairwise correlation coefficient of 0.72), but they differ in three important aspects.

[Insert Figure 4 about here]

First, in line with the patterns revealed by the expected dividend growth, the dp_t^{Corr} is on average higher than the dp_t in the boom periods and it is lower than the dp_t in the bust periods. This means that the dp_t tends to predict returns that are too low to be justified with the market's optimism about growth opportunities during the boom periods. Simultaneously, the dp_t tends to forecast returns that are too high during the crisis periods. This is especially apparent at the end of the sample when the market experienced one of the largest drops in the

[13] Note: the construction of the corrected dividend-price ratio introduces a look-ahead bias because $\widehat{\rho}$ and $\widehat{\gamma_1}$ are estimated using the data of the whole sample and are therefore not available at time t. However, the out-of-sample predictability results in the Section 6 show that the look-ahead bias plays only a minor role when predicting returns with the corrected dividend-price ratio.

history of the U.S. market, but the uncorrected dividend-price ratio rose and therefore implied unrealistically high returns.

Second, the corrected dividend-price ratio is notably more volatile than the uncorrected dividend-price ratio. The standard deviation is 0.27 for the dp_t and 0.54 for the dp_t^{Corr}. In the context of the present value model, this increase in volatility implies that expected returns and expected dividend growth are highly correlated. To see this formally, expand the variance of the corrected dividend-price ratio as:

$$var(dp_t^{Corr}) = var(dp_t) + 2\left(\frac{1}{1-\rho\gamma_1}\right)\left(\frac{1}{1-\rho\delta_1}\right)cov(\mu_t, g_t) - \left(\frac{1}{1-\rho\gamma_1}\right)^2 var(g_t) \qquad (24)$$

Equation (24) says that the variance of the dp_t^{Corr} can be higher than the variance of the dp_t only if expected returns and expected dividend growth rates covary and the covariation is big enough $\left(2\left(\frac{1}{1-\rho\delta_1}\right)cov(\mu_t, g_t) > \left(\frac{1}{1-\rho\gamma_1}\right)var(g_t)\right)$. Furthermore, since dp_t increases with expected returns and decreases with expected dividend growth, this positive covariation also affects the uncorrected dividend-price ratio. It offsets shocks to expected returns and expected dividend growth and reduces the volatility of the dp_t (Lettau and Ludvigson, 2005, Rytchkov, 2008; Van Binsbergen and Koijen, 2010). Thus, correcting the dp_t for the implied dividend growth restores the variation, which is otherwise offset by the comovement of the expected return and the expected dividend growth (see also Lacerda and Santa-Clara, 2010).

Last, consistent with the increase in volatility of the dp_t^{Corr}, the dp_t^{Corr} is also less persistent than the dp_t. While dp_t exhibits first order autocorrelation coefficient of 0.98, the $AR(1)$ for dp_t^{Corr} is notably lower and amounts to 0.74. This decrease in persistence is important because highly autocorrelated predictors are typically subject to small sample bias (Stambaugh, 1999) and produce inaccurate inference results in the case of overlapping observations (Boudoukh et al., 2008). Given its lower persistence, the dp_t^{Corr} is therefore largely free of the common concern related to the use of highly persistent variables for predicting returns.

By applying equation (15) and following the same estimation procedure as for the implied dividend yield, I additionally estimate the implied interest rate (IIR_t). Although IIR_t is not of special interest for this study, it is important to note that the IIR_t behaves as we would expect. As shown in Figure 5, IIR_t strongly covaries with the T-bill rate and the LIBOR rate and it is on average closer to the LIBOR rate (see also Naranjo, 2009). Still, IIR_t is more volatile than the T-bill rate and the LIBOR rate at the beginning of the analyzed period and it deviates from both proxies for the interest rate during the recent financial crisis, when it is notably lower than the LIBOR rate. This shows that the implied interest rate may deviate from the observable proxies for the interest rate and it therefore points at the importance of isolating the effect of the interest rate when estimating the implied dividend yield.

[Insert Figure 5 about here]

5 Empirical results

This section presents dividend growth and market return predictability results. Since derivative markets subsume market expectations about the near future, the implied dividend ratios should be especially suitable for tracking short term variations in future dividends and returns as opposed to long term tendencies in asset markets. To investigate this, I consider predicting dividend growth rates and market returns at the horizons ranging from one to six months.

I use standard predictive regressions, in which returns or dividend growth rates are regressed on the lagged predictors. I report OLS t-statistics for the case of non-overlapping monthly observations and Hodrick (1992) t-statistics for the case of longer horizon regressions with overlapping observations.[14] Additionally, I report the adjusted R^2. Note however that the R^2

[14] Ang and Bekaert (2007) show that the performance of Hodrick (1992) standard errors, which are based on

in the context of overlapping observations needs to be interpreted with caution because it tends to increase with the length of the overlap even in the absence of true predictability (Valkanov, 2003; Boudoukh et al., 2008).

5.1 Predicting dividend growth

Figure 3 shows that the implied dividend growth tracks general market conditions and it therefore seems to be a good proxy for the expected dividend growth. In this subsection, I complement this argument by showing that the implied dividend growth also uncovers part of the variation in the future dividend growth.

For a comparison with the implied dividend growth, I consider whether the dividend-price ratio predicts future dividend growth. I use dp_t as a competing predictor for two reasons. Firstly, dp_t is itself a function of the expected dividend growth and could therefore predict future dividend growth as opposed to future returns. Secondly, implied dividend growth is defined as the difference between the implied dividend yield and the dividend-price ratio. Therefore, it is necessary to show that the implied dividend growth does not predict future dividend growth simply because it is duplicating information contained in the dp_t.

The main regression takes the following form:

$$\Delta d_{t+h} = a_0 + a_1(X_t) + \varepsilon_{t+1} \tag{25}$$

where $\Delta d_{t+h} = (12/h)\sum_{i=1}^{h} \Delta d_{t+i}^{M}$ is the annualized dividend growth with $h = 1, 2, 3$ or 6 months and X_t is either idg_t, or dp_t, or both. For $h = 1$, t-statistics are based on the simple OLS. For $h = 2, 3$ or 6, t-statistics are computed according to Hodrick (1992).

Table III presents results. I start by analyzing regression results with the dividend-price

summing the predictors in the past, is superior to other standards errors that are frequently employed in the literature, such as the Newey-West (1987) standard errors, or the Hansen and Hodrick (1980) standard errors.

ratio. The estimated parameter on the dp_t is negative, just as the theory suggests, but the associated t-statistics are insignificant at the conventional 5% level and range between 1.41 and 1.96. Also, the *adj.* R^2 is low and ranges from 1.47% for monthly dividend growth to 3.28% for half-annual dividend growth. In comparison, the implied dividend growth is positively related to future dividend growth and explains 4.79% of the variation in the monthly dividend growth and 18.42% of the variation in the half-annual dividend growth. Furthermore, all the estimated coefficients on the implied dividend growth are statistically significant and range between 3.25 and 4.76.

As reported in the last panel of Table III, adding dividend-price ratio as an additional predictor to the implied dividend growth boosts statistical significance of the implied dividend growth and leads to further increase in the *adj.* R^2. The *adj.* R^2 in a bivariate predictive regression amounts to 9.31% for monthly dividend growth and to 30.81% for half annual dividend growth. Since this is more than the sum of the *adj.* $R^{2\prime}s$ in the univariate regressions, it clearly indicates that the implied dividend growth is not duplicating information about future returns that is already captured in the dividend-price ratio.

[Insert Table III about here]

5.2 Predicting market returns

I employ three specifications for the return predictive regressions. The first is the standard predictive regression, in which returns are regressed on the lagged dividend-price ratio:

$$r_{t+h} = b_0 + b_1(dp_t) + \varepsilon_{t+1} \tag{26}$$

The second regression augments the first by using the proxy for the expected dividend growth (implied dividend growth):

$$r_{t+h} = c_0 + c_1(dp_t) + c_2(idg_t) + \varepsilon_{t+1} \tag{27}$$

The last return regression replaces the dividend-price ratio and the implied dividend growth by the corrected dividend-price ratio:

$$r_{t+h} = d_0 + d_1(dp_t^{Corr}) + \varepsilon_{t+1} \tag{28}$$

In all the regressions, $r_{t+h} = (12/h)\sum_{i=1}^{h} r_{t+i}^{M}$ is the annualized market return with $h = 1, 2, 3$ or 6 months. For $h = 1$, t-statistics are based on the simple OLS. For $h = 2, 3$ or 6, t-statistics are computed according to Hodrick (1992).

Table IV presents the regression results. I start by analyzing univariate regression results of returns on the lagged dp_t. The estimated coefficient on the dp_t is positive, as suggested by the theory, but the t-statistics are insignificant at the 5% level of statistical significance and range between 1.28 and 1.67. Also, the associated *adj.* R^2 is relatively low and ranges from 0.33% for monthly returns to 7.02% for half-annual returns.

When implied dividend growth is added as an additional regressor to the dividend-price ratio, the return predictability improves for all the considered horizons. The *adj.* R^2 increases from 0.33% to 5.20% in the regression with monthly returns and from 7.02% to 8.71% in the regression with half-annual returns. This result is directly in line with the observation that the implied dividend growth predicts future dividend growth and thereby implies that variation in the expected dividend growth plays an important role for uncovering variation in the future returns.

As suggested by the present value model and confirmed by the last regression, the dp_t and the implied dividend growth can also be replaced by a single predictor, the corrected dividend-price ratio. The corrected dividend-price ratio predicts returns approximately as well as the dividend-price ratio and the implied dividend growth together. The *adj.* R^2 amounts to 4.61%

at the monthly horizon and to 8.56% at the half-annual horizon. Also, the estimated parameter on the dp_t^{Corr} is always statistically significant with the t-statistics ranging from 3.19 at the monthly horizon to 2.33 at the half-annual horizon.

[Insert Table IV about here]

6 Additional tests

The results imply that the corrected dividend-price ratio predicts returns significantly better than the realized dividend-price ratio, and that the improvement in the predictability is especially pronounced over the monthly horizon. However, all the results so far are based on the in-sample predictive regressions, which have been criticized on the grounds that they may be subject to the small sample bias (Stambaugh, 1999), and they may not necessarily imply that the documented predictability can be exploited in real time (Goyal and Welch, 2008).

To address these issues, this section considers small sample bias correction, out-of-sample predictability and a simple out-of-sample trading strategy. Additionally, I compare the return predictive ability of the corrected dividend-price ratio to the alternative corrections for the dividend-price ratio and to other popular predictors. To avoid the statistical problems inherent in the use of overlapping observations (Boudoukh et al., 2008), the analysis is restricted to predicting non-overlapping monthly returns.

6.1 Is there a small sample bias?

Dividend ratios are very persistent and an extensive literature argues that the standard OLS predictive regressions applied to highly persistent variables may lead to severe biases in small samples (Stambaugh, 1999; Amihud and Hurvich, 2004).

To analyze the source of the bias, consider a model where returns are predicted by a variable (X_t) that follows first-order autoregressive process:

$$r_{t+1} = \alpha + \beta X_t + u_{t+1} \tag{29}$$

$$X_{t+1} = \delta + \rho X_t + v_{t+1} \tag{30}$$

where $|\rho| < 1$ and the errors ($u_{t+1,}v_{t+1}$) are distributed as:

$$\begin{pmatrix} u_{t+1} \\ v_{t+1} \end{pmatrix} \sim_{iid} N(0, \Sigma), \Sigma = \begin{pmatrix} \sigma_u^2 & \sigma_{uv} \\ \sigma_{uv} & \sigma_v^2 \end{pmatrix} \tag{31}$$

If errors are correlated ($\sigma_{uv} \neq 0$), OLS produces a biased estimate of β in small samples (Stambaugh, 1999). The larger the ρ, i.e. the persistence of shocks to the predictor variable, the larger the bias. For dividend-price ratios, σ_{uv} is negative and ρ is close to one. This results in upward biased estimates of β and the corresponding t-statistics.

To correct for the small sample bias, I follow the correction methodology proposed by Amihud and Hurvich (2004) and employed in several recent studies (Boudoukh et al. 2007; Kolev, 2008; Lioui and Rangvid, 2009). First I estimate (29) to obtain an OLS estimate $\widehat{\rho}$. Then, I calculate the bias corrected estimator for $\widehat{\rho}$:

$$\widehat{\rho}^c = \widehat{\rho} + (1 + 3\widehat{\rho})/n + 3(1 + 3\widehat{\rho})/n^2 \tag{32}$$

where n is the length of the time series. The estimator $\widehat{\rho}^c$ is then used to calculate the bias corrected errors:

$$v^c_{t+1} = X_{t+1} - [(1 - \widehat{\rho}^c)\Sigma^n_{t=1}(X_{t+1}/n) + \widehat{\rho}^c X_t] \tag{33}$$

Finally, I run an OLS regression of returns on the predictor variable X_t and the v_{t+1}^c :

$$r_{t+1}= \alpha + \beta^c X_t + \phi^c v_{t+1}^c + \varepsilon_{t+1} \tag{34}$$

The estimate of β^c gives us the bias corrected estimator of β. The corresponding bias corrected t-statistic is calculated as:

$$t^c = \widehat{\beta}^c / \sqrt{\left(\widehat{\phi}^c\right)^2 \left(\widehat{SE}(\widehat{\rho})\right)^2 (1 + 3/n + 9/n^2)^2 + \left(\widehat{SE}(\widehat{\beta}^c)\right)^2} \tag{35}$$

I apply the bias correction to the realized dividend-price ratio and to the corrected dividend-price ratio. Table V compares and contrasts the slope estimates and the t-statistics based on the standard OLS with those obtained after correcting for the small sample bias.

The realized dividend-price ratio is an insignificant predictor for monthly returns even before correcting for the small sample bias. After correction, the estimated predictive coefficient even changes its sign and becomes negatively related to future returns. Unlike the realized dividend-price ratio, the corrected dividend-price ratio is largely unaffected by the small sample bias correction. The adjusted slope coefficient is almost identical to the OLS slope coefficient (0.22 in comparison to 0.23) and the adjusted t-statistic is only marginally smaller than the OLS t-statistic (3.10 in comparison to 3.19).

[Insert Table V about here]

A rather small effect of the small sample bias correction on the inference of the dp_t^{Corr} is due to a combination of two effects. First, the corrected dividend-price ratio is less persistent than the realized dividend-price ratio (0.74, in comparison to 0.98). Second, the innovations to the predictor variable and to the returns are only weakly correlated (-0.24 for the dp_t^{Corr}

in comparison to -0.97 for the realized dividend-price ratio). The combination of both effects enables the dp_t^{Corr} to remain statistically significant predictor for monthly returns and hence, implies that the corrected dividend-price ratio is by and large robust to small sample bias.

6.2 Out-of-sample predictability

Goyal and Welch (2008) demonstrate that variables with in-sample predictive power may not necessarily predict returns out-of-sample. I follow their approach to test whether the corrected dividend-price ratio predicts returns out-of-sample better than the realized dividend-price ratio.

I calculate the out-of-sample R^2 as in Campbell and Thompson (2008) and Goyal and Welch (2008):

$$R_{OS}^2 = 1 - \frac{\sum_{t=1}^{T} (r_{t+1} - \widehat{\mu}_t)^2}{\sum_{t=1}^{T} (r_{t+1} - \overline{r_t})^2} \tag{36}$$

where $\widehat{\mu}_t$ is the fitted value from a predictive regression estimated through period t and $\overline{r_t}$ is the historical average return estimated through period t. A positive out-of-sample R^2 indicates that the predictive regression has a lower mean-squared prediction error than the historical average return.

To make out-of-sample forecasts, I split the sample in two subperiods. I use the period from January 1994 through December 1999 for the estimation of the initial parameters and the period from January 2000 through December 2009 for the calculation of the R_{OS}^2. All out-of-sample forecasts are based on a recursive scheme using all the available information up to time t. I calculate R_{OS}^2 for the realized and the corrected dividend-price ratio.

Recall that the corrected dividend-price ratio is defined as:

$$dp_t^{Corr} = dp_t + \left(\frac{1}{1-\widehat{\rho}\widehat{\gamma}_1}\right) idg_t \tag{37}$$

where ρ (linearization constant) and γ_1 (AR(1) coefficient of the implied dividend growth) are estimated using the whole sample period and therefore introduce a slight look-ahead bias in the construction of the corrected dividend-price ratio. To alleviate the concern that the look-ahead bias may be influencing the results, I additionally estimate the so called No-Look-Ahead-Bias corrected dividend-price ratio:

$$dp_t^{NLAB\_Corr} = dp_t + \left(\frac{1}{1-\widehat{\rho}_t\widehat{\gamma}_t}\right) idg_t \tag{38}$$

where $\widehat{\rho}_t$ and $\widehat{\gamma}_t$ are time-varying and estimated using the same recursive scheme as in the calculation of the out-of-sample R^2_{OS}.

Table VI reports results. The dp_t that exhibits poor ability to predict returns in-sample also fails to predict returns out-of-sample. The R^2_{OS} for the dp_t is -0.15%. In comparison, the out-of-sample R^2 for the dp_t^{Corr} is as high as 6.06%. Thus, the dp_t^{Corr} does not predict returns only in-sample, but it also delivers superior out-of-sample forecasts of the monthly returns relative to the forecasts based on the historical average. Furthermore, approximately the same R^2_{OS}, if not even slightly higher, is also obtained with the corrected dividend-price ratio that is adjusted for the look-ahead bias (R^2_{OS} 6.09%). Hence, the look-ahead bias is not a concern and the dp_t^{Corr} can be effectively used in real time for the portfolio allocation decisions.[15]

[Insert Table VI about here]

[15] The rather small difference in the R^2_{OS} between the dp_t^{Corr} and the $dp_t^{NLAB\_Corr}$ is driven by the fact that the persistence of the implied dividend growth $\widehat{\gamma}_t$ and the linearization constant $\widehat{\rho}_t$ are very stable ($\widehat{\gamma}_t$ ranges from 0.47 to 0.61 and $\widehat{\rho}_t$ is always between 0.98 and 0.99). This makes dp_t^{Corr} and $dp_t^{NLAB\_Corr}$ highly correlated (0.99) and almost indistinguishable from each other.

To illustrate the relative success of the dp_t^{Corr} in predicting returns out-of-sample, Figure 6 plots out-of-sample forecasts along with the realized returns. Although realized returns are significantly more volatile than any of the forecasted returns, there are considerable differences between the forecasts. The forecasts based on the realized dividend-price ratio and the forecasts based on the historical average return are both very smooth and almost indistinguishable from each other. In comparison, the forecasts based on the corrected dividend-price ratio vary significantly more and the changes of the forecasts are typically of the same sign as the changes of the realized returns.

6.3 Economic value of the corrected DP

To assess the economic value of the documented improvement in predicting returns, I run a simple out-of-sample trading strategy. I consider a mean-variance investor who invests in the stock market and the risk-free rate. Each period the investor uses different predictor variables to estimate one period ahead expected return $\widehat{\mu}_t$. Based on these estimates, the investor's portfolio weight on the stock market at time t is given by:

$$w_t = \frac{\widehat{\mu}_t - rf_{t+1}}{\gamma \widehat{\sigma}^2} \tag{39}$$

where rf_{t+1} is the one period ahead risk-free rate, γ is the risk-aversion coefficient and $\widehat{\sigma}^2$ is the variance of the stock market. I set γ equal to 3 and I proxy the variance of the market by the variance as implied in the options on the S&P 500 (VIX). The time-series of portfolio returns is then given by:

$$Rp_{t+1} = w_t r_{r+1} + (1 - w_t) rf_{t+1} \tag{40}$$

I assess economic value of predictors by the certainty equivalent return CE and the Sharpe ratio SR :

$$CE = \overline{Rp} - \frac{\gamma}{2}\widehat{\sigma}^2(Rp) \tag{41}$$

$$SR = \frac{\overline{Rp^e}}{\widehat{\sigma}(Rp^e)} \tag{42}$$

where $\overline{Rp}$ and $\widehat{\sigma}^2(Rp)$ are the mean and the variance of the portfolio return, and the superscript e stands for returns in excess of the risk-free rate. As in the calculation of the out-of-sample R^2_{OS}, I use the period from January 1994 through December 1999 for the estimation of the initial parameters and the period from January 2000 through 2009 for the calculation of the $CE's$ and the $SR's$. All out-of-sample forecasts are based on a recursive scheme using all the available information up to time t.

Table VII reports results. The first column reports certainty equivalents and the second column reports Sharpe ratios. All the values are annualized. Note that the average excess return on the S&P 500 in the period from 2000 to 2009 is negative, which points at the difficulty of building trading strategies with positive Sharpe ratios. Indeed, a trading strategy based on the historical average return delivers a CE of 1.75%[16] and a negative Sharpe ratio (-0.10). Using dividend-price ratio to time the market yields slightly better results. The CE amounts to 2.78% and the Sharpe ratio becomes positive, but remains at the low level of 0.09.

In comparison, the corrected dividend-price ratio yields a CE as high as 5.07% and a Sharpe ratio of 0.41. This is a 0.32 gain in terms of the Sharpe ratio and a 2.29% gain in terms of the CE. In other words, an investor who is timing the market with the dp_t would be willing to pay as much as 2.29% of the invested wealth to get the access to the dp_t^{Corr}.

[Insert Table VII about here]

[16] A comparable number for the period from 1947 to 2007 is 7.4% (Ferreira and Santa-Clara, 2010).

6.4 Alternative predictors

To further assess the return-predictive ability of the corrected dividend-price ratio, I consider a set of alternative return predictors.

Lacerda and Santa-Clara (2010) show that correcting the dividend-price ratio for the 10 year moving average of dividend growth improves predictability of longer horizon (i.e. annual) returns. Following their approach, I construct the dividend-price ratio corrected for the changes in the average historical dividend growth. Since the focus of this study lies on the short horizon predictability and the 10 year moving average of the dividend growth rate is rather slowly evolving, I calculate average historical dividend growth as a moving average of one year of annualized monthly dividend growth rates. Furthermore, to foster comparability with the dividend-price ratio corrected for the implied dividend growth, I assume that the persistence of the historical dividend growth is the same as the persistence of the implied dividend growth. The dividend-price ratio adjusted for the historical dividend growth is then defined as dp_t^{HIST} :

$$dp_t^{HIST} = dp_t + 2.08 * dg_t^{\overline{M}} \tag{43}$$

where $dg_t^{\overline{M}}$ is the moving average of annualized monthly dividend growth rates over the past year.

In addition, I use the variance risk premia (vrp_t) as implied in the S&P 500 (Bollerslev et al., 2009). The variance risk premia is arguably one of the strongest predictors for short horizon returns and is also estimated from the S&P 500 derivatives. Therefore, it is instructive to compare the return predictive ability of the corrected dividend-price ratio to the variance risk premia. Furthermore, I employ two other standard predictors, the earnings-price ratio (ep_t) and the consumption-to-wealth ratio (cay_t) proposed by Lettau and Ludvigson (2001).

Summary statistics. Table VIII reports the basic summary statistics and the unconditional correlation structure for the predictors sampled monthly.[17] The numbers are in line with previous studies. Except for the variance risk premia (vrp_t) and the corrected dividend-price ratio (dp_t^{Corr}), all the predictors are highly persistent with a first-order autocorrelation of more than 0.9. The persistence of the dividend-price ratio corrected for the average historical dividend growth (dp_t^{HIST}) is slightly lower than the persistence of the realized dividend-price ratio, but it is still of the similar magnitude (0.97).

The realized dividend-price ratio (dp_t), dividend-price ratio corrected for the historical dividend growth (dp_t^{HIST}) and the dividend-price ratio corrected for the implied dividend growth (dp_t^{Corr}) are all highly correlated and they exhibit similar relationships with respect to alternative predictors. They are all positively correlated with the earnings-price ratio and the consumption-to-wealth ratio and they are negatively related to the variance risk premia.

[Insert Table VIII about here]

Predicting market returns. Table IX reports results for predicting monthly returns. Since in-sample and out-of-sample results are largely consistent, I evaluate predictors mainly on the in-sample evidence. The traditional predictors based on the realized data explain only a small part of the variation in the future monthly returns. The earnings-price ratio exhibits a slightly negative *adj.* R^2. The dividend-price ratio, as already documented, explains 0.33% of the variation in the future monthly returns. The consumption-to-wealth ratio exhibits *adj.* R^2 of around one percent. Furthermore, correcting dividend-price ratio for the variation in the average historical dividend growth does not seem to improve predictability of monthly market returns in the analyzed period. The *adj.* R^2 in a univariate regression with the dp_t^{HIST} is

[17] Since cay_t is available only at the quarterly frequency, monthly observations of cay_t are defined by the most recently available quarterly observation.

approximately the same, if not slightly lower, as the *adj.* R^2 in the regression with the realized dividend-price ratio.[18]

In comparison, the dividend-price ratio corrected for the implied dividend growth (dp_t^{Corr}) and the variance risk premia (vrp_t) explain a significantly higher portion of the variation in the future monthly returns. The vrp_t exhibits an *adj.* R^2 of 4.06%. The dp_t^{Corr}, as already documented, exhibits an *adj.* R^2 of 4.61%. The variance risk premia and the corrected dividend-price ratio are also the only predictors that are significant at the conventional levels of statistical significance.[19]

Since the corrected dividend-price ratio and the variance risk premia are both based on variables that are extracted from derivative markets, the relative success of the dp_t^{Corr} in predicting future returns could be driven by the fact that the dp_t^{Corr} is simply duplicating information contained in the vrp_t. To address this concern, I additionally consider a bivariate regression with the dp_t^{Corr} and the vrp_t. Quite interestingly, adding vrp_t as an additional predictor to the dp_t^{Corr} boosts statistical significance of both predictors and the *adj.* R^2 increases to as much as 9.99%. This result is even more remarkable because both predictors also predict returns out-of-sample with the R^2_{OS} of 11.55%. Thus, the corrected dividend-price ratio is not duplicating information about future returns that is already captured in the variance risk premia. Additional support for this interpretation is provided by the fact that a bivariate regression with the variance risk premia and, either the realized dividend-price ratio (dp_t), or the dividend-price

[18] Unreported results show that this result is robust to using alternative proxies for the historical dividend growth, such as lagged monthly dividend growth, lagged annual dividend growth or the moving average of 10 years of monthly dividend growth rates.

Furthermore, approximately the same results are also obtained using the correction for the dividend-price ratio proposed by Lacera and Santa-Clara (2010):

$$dp_t^{LSC} = dp_t + dg_t^{\overline{M}} \left(\frac{1}{1-\rho_t} \right)$$

where ρ_t is a time-varying linearization constant and the correction does not depend on the persistence of the dividend growth. The only difference between the dp_t^{LSC} and the dp_t^{HIST} is that the dp_t^{LSC} is more volatile and less persistent because ρ_t is close to 0.98, implying that the typical correction for the dividend-price ratio is around $50 * dg_t^{\overline{M}}$ as opposed to the $2.08 * dg_t^{\overline{M}}$ used in this paper.

[19] Note that the results with the vrp_t are sensitive to the definition of the variance risk premia. A predictive regression with the variance risk premia defined as the difference between the VIX and the objective expectations of the realized variance (as opposed to the difference between the VIX and the actual realized variance) exhibits a slightly negative *adj.* R^2.

ratio adjusted for the historical dividend growth (dp_t^{HIST}) results in a considerably smaller *adj.* R^2 (approximately 5%).

[Insert Table IX about here]

7 Variance decomposition of the DP

Until now I used the present value model merely to motivate the predictive regressions. In this Section, I employ the model to decompose the variance of the dividend-price ratio and to provide further insights for the interpretation of the results. As before, I treat the implied dividend growth as a true proxy for the expected dividend growth (no measurement error) and I use annualized variables taken at the monthly frequency.[20]

Within the framework of the present value model, the variance of the dividend-price ratio can be decomposed as:

$$\begin{aligned} var(dp_t) \simeq & \left(\frac{1}{1-\rho\delta_1}\right)^2 var(\mu_t) + \left(\frac{1}{1-\rho\gamma_1}\right)^2 var(g_t) \\ & -2 * \left(\frac{1}{1-\rho\delta_1}\right)\left(\frac{1}{1-\rho\gamma_1}\right) cov(\mu_t, g_t) \end{aligned} \tag{44}$$

where the first term on the right hand-side presents the contribution of the expected return, the second term presents the contribution of the expected dividend growth and the last term presents the contribution of the covariation between the expected return and the expected dividend growth.

[20]Unreported results show that repeating the exercise with non-overlapping annual data does not change results qualitatively.

I set $(1 - \widehat{\rho}_t \widehat{\delta_1})$ equal to the estimated parameter on the corrected dividend-price ratio (see equation (12)) and I calculate the variance of the expected return by inverting equation (7). I standardize all terms on the right-hand side of (44) by the left-hand side, so that the terms sum up to 100%.

The results imply that 400% of the variance of the dp_t is driven by the variation in the expected returns and 210% of the variance of the dp_t is driven by the time-varying expected dividend growth rate. This means that the covariance term account for as much as 510% of the variation in the dividend-price ratio, which further implies that the correlation between the expected return and the expected dividend growth is as high as 0.88.

Thus, contrary to the standard result that virtually all the variation in the dividend-price ratio is driven by the time-varying expected returns (Campbell 1991; Cochrane 2005), the results show that there is a lot of variation in both expected returns and expected dividend growth. However, the positive correlation between them offsets each other within the dividend-price ratio. This dampens the volatility of the dividend-price ratio and explains why the dp_t fails to predict returns (and dividend growth rates) (see also Menzly et al., 2004 and Lettau and Ludvigson, 2005).

Furthermore, as discussed in Section 3, correcting the dividend-price ratio for changes in the expected dividend growth restores the variation that is offset by the positive correlation between the expected returns and the expected dividend growth. This makes the corrected dividend-price ratio more volatile than the uncorrected dividend-price ratio, and thus implies that expected returns vary significantly more than is suggested by the uncorrected dividend-price ratio.

8 Robustness checks

To validate the documented improvement in predicting market returns, I show that results are robust to several methodological changes in the calculation of the corrected dividend-price ratio.

8.1 Maturity of the implied dividend yield

The corrected dividend-price ratio analyzed throughout the paper is defined as:

$$dp_t^{Corr} = dp_t + 2.08 * idg_t = dp_t + 2.08 * (idy_t - dp_t). \tag{45}$$

where idg_t is the log implied dividend growth rate calculated as the difference between the log implied dividend yield (idy_t) and the log realized dividend-price ratio (dp_t). Given the trade-off between the seasonality in dividend payments and the liquidity of the derivatives, the results in the main analysis are based on the annualized implied dividend yield with 6 months to maturity and the realized dividend-price ratio estimated in a standard way by summing dividends over the past 12 months. To address the concern that the maturity mismatch between the implied dividend yield and the realized dividend-price ratio could be a source of seasonality driving the documented improvement in predicting returns, I consider two robustness checks.

In the first robustness check, I re-estimate the corrected dividend-price ratio using implied dividend yields with maturities between 3 and 7 months. In the second robustness check, I repeat the same exercise, but instead of the standard dividend-price ratio with dividends summed over the past 12 months, I use a dividend-price ratio based on dividends summed over the past 6 months, $dp_t^{6m} = \log\left[\frac{D_t^6}{P_t}\right]$,where D_t^6 is the annualized sum of dividends over the past 6 months.

To foster comparability between the different versions of the corrected dividend-price ratios, I impose that the persistence of all the implied dividend growth rates is the same and equals the persistence of the implied dividend growth rate used in the main analysis. In other words, the corrected dividend-price ratio is always calculated as the dividend-price ratio plus 2.08 times the implied dividend growth.

Table X and Table XI report results for the first and the second robustness check, respectively.[21] I start by analyzing results in Table X. The results seem to offer two general conclusions. First, irrespective of the maturity of the implied dividend yield, the corrected dividend-price ratio exhibits statistically significant predictive coefficients. Second, the maturity of the implied dividend yield seems to matter. The corrected dividend-price ratio based on the implied dividend yields with longer maturities (5, 6 and 7 months) predicts returns better than the dp_t^{Corr} based on the implied dividend yields with maturities of 3 or 4 months. Since implied dividend yields with longer maturities are less prone to the seasonality in dividend payments, the results suggest that the documented improvement in the predictability of monthly returns is unlikely to be driven by the seasonality in dividend payments.

[Insert Table X about here]

[Insert Table XI about here]

Results reported in Table XI further reveal that the documented predictability is robust to the alternative way of constructing the realized dividend-price ratio. Specifically, the comparison of the results reported in Table X and Table XI shows that the dividend-price ratio based on either half-annual or annual dividends exhibit identical $adj.\ R^2$ and the predictive coefficients differ only marginally. Also, the ability of the corrected dividend-price ratio to predict returns is largely unaffected by the alternative way of constructing the realized dividend-price ratio.

[21] Due to the high volatility of the implied dividend yields with short maturity (see Table I), the implied dividend yield with maturity of 3 months takes a negative value on three occasions (December 1999, May 2000 and July 2000). I replace these observations with the implied dividend yield with maturity of 4 months.

All in all, results show that the documented improvement in predicting market returns cannot be explained by the maturity mismatch in the calculation of the implied dividend growth.

8.2 Options moneyness and backward-looking data

The implied dividend yield used in the calculation of the corrected dividend-price ratio is estimated from no-arbitrage relations spanning the prices of index derivatives. Since no-arbitrage relations can be violated for particular pairs of options and futures, but hold well in general (Kamara and Miller, 1995), I calculate implied dividend yield by aggregating information from a wide set of options and futures. Each end of month, I use 10 days of backward-looking data and options across all the moneyness levels.

The use of such a wide set of data is necessary to smooth dividend yield estimates, but it may lead to inclusion of unreliable data. For example, the wider the moneyness level, the more observations we have for calculation of the implied dividend yield. Nevertheless, deep out-of-the money options are less liquid and therefore deemed unreliable. In the main run, I use all the options across all the moneyness levels. Now, I consider filtering out options with moneyness levels below 0.8 (0.9) and above 1.2 (1.1), respectively. Similar argument applies to the use of backward-looking data. We should expect that the most recent data is most important for forecasting purposes. However, more data may be needed to smooth the implied dividend yield estimates. In the main run, I use 10 days of backward-looking data. Now, I consider using either 5 or 15 days of backward-looking data. As before, I always calculate corrected dividend-price ratio as the dividend-price ratio plus 2.08 times the implied dividend growth.

Results are reported in Table XII and confirm the above conjectures. Filtering out unreliable deep out-of-the money options improves the ability of the corrected dividend-price ratio to predict monthly returns. For example, the corrected dividend-price ratio based on options with moneyness levels between 0.9 and 1.1 explains as much as 5.82 percent of the variation in the future monthly market returns. Imposing even tighter restrictions on the moneyness levels should lead to even better results, but is unfortunately limited by the relatively low level of options liquidity at the beginning of the sample period. The liquidity issues are even more pronounced in the second exercise. Using 15 days as opposed to 10 days of backward-looking data does not seem to influence considerably either volatility of the corrected dividend-price ratio or its ability to predict future returns. In comparison, using 5 days of backward-looking data makes the corrected dividend-price ratio more prone to violations of no-arbitrage relations and significantly more volatile. This also reduces its ability to predict returns. Nevertheless, the estimated parameter on the dp_t^{Corr} remains significant and it still predicts returns significantly better than the uncorrected dividend-price ratio.

[Insert Table XII about here]

9 Conclusions

That variation in the expected dividend growth reduces the ability of the dividend-price ratio to predict returns is a long-standing notion in the predictability literature (Fama and French, 1988). However, empirical analysis of this issue is complicated because the expected dividend growth is an aggregate of investors' expectations about future growth opportunities and is therefore difficult to estimate.

In this paper I propose extracting the expected dividend growth from derivative markets (index options and futures). Because prices of derivatives depend on, inter alia, the dividends that the underlying asset pays until the expiration of the contracts, they provide a unique laboratory for estimating the dividend growth that investors expect to realize in the near future. Indeed, I find that the implied dividend growth uncovers variation in future dividend growth and thereby allows for improvements in predicting market returns. Using implied dividend growth as an additional regressor in the standard dividend-price ratio return predictive regression—or correcting the dividend-price ratio for variation in the implied dividend growth—significantly improves the predictability of short-run S&P 500 returns over the past 16 years.

This predictive improvement is especially strong over a short horizon (i.e. monthly returns), holds both in-sample and out-of-sample, yields a sizable gain in the Sharpe ratio, and is robust to small sample bias. Furthermore, these results are not driven by the fact that implied dividend growth duplicates information in other, well-known options-implied predictors (e.g., the variance risk premia), and neither can they be replicated using historical rather than implied dividend growth.

Importantly, the results show that the expected return and expected dividend growth are highly correlated. This high correlation means that movements in expected returns and expected dividend growth offset each other's effect in the dividend-price ratio, which renders the dividend-price ratio relatively smooth. Correcting the dividend-price ratio for the implied dividend growth restores the variation which is otherwise obscured by this strong comovement, and hence implies that expected returns vary significantly more than is apparent from the observed variation in the uncorrected dividend-price ratio.

References

[1] Ait-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset markets, *Journal of Finance* 53, 499-547.

[2] Ang, Andrew, and Geert Bekaert, 2007, Stock return predictability: Is it there?, *Review of Financial Studies* 20, 651-707.

[3] Amihud, Yakov, and Clifford M. Hurvich, 2004, Predictive regressions: A reduced-bias estimation method, *Journal of Financial and Quantitative Analysis* 39, 813-841.

[4] Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481-1509.

[5] Binsbergen, Jules H. van, and Ralph Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance*, forthcoming.

[6] Binsbergen, Jules H. van, Michael W. Brandt, and Ralph Koijen, 2010, On the timing and pricing of dividends, *Unpublished manuscript.*

[7] Boudoukh, Jacob, Roni Michaely, and Matthew Richardson, 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* 62, 877-915.

[8] Boudoukh, Jacob, Matthew Richardson, and Robert F., Whitelaw, 2008, The myth of long-horizon predictability, *Review of Financial Studies* 21, 1577-1605.

[9] Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009. Expected stock returns and variance risk premia, *Review of Financial Studies 22*, 4463-4492.

[10] Campbell, John. Y., 1991, A variance decomposition for stock returns, *Economic Journal 101*, 157-179.

[11] Campbell, John.Y., Samuel B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509-1531.

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

[13] Cochrane, John. H., 2005, *Asset pricing, revised ed.* Princeton, NJ: Princeton University Press.

[14] Cochrane, John H., 2008a, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533-1575.

[15] Cochrane, John H., 2008b, State-space vs. VAR models for stock returns, Unpublished manuscript.

[16] Fama, Eugene F., and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3-27.

[17] Favero, Carlo A., Arie Gozluklu, and Andrea Tamoni, 2010, Demographic trends, the dividend-price ratio and the predictability of long-run stock market returns, *Journal of Financial and Quantitative Analysis*, forthcoming.

[18] Ferreira, Miguel A., Santa-Clara, Pedro, 2010, Forecasting stock market returns: The sum of the parts is more than the whole, *Journal of Financial Economics*, forthcoming.

[19] Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455-1508.

[20] Hansen, Lars P., and Robert J. Hodrick, 1980, Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy* 88, 829-853.

[21] Hodrick, Robert. J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357-386.

[22] Kamara, Avraham, and Thomas Miller Jr., 1995, Daily and intradaily tests of European put-call parity, *Journal of Financial and Quantitative Analysis* 30, 519-539.

[23] Kolev, Gueorgui I., 2008, Forecasting aggregate stock returns using the number of initial public offerings as a predictor, *Economics Bulletin* 7, 1-8.

[24] Lacerda, Filipe, and Pedro Santa-Clara, 2010, Forecasting dividend growth to better predict returns, Unpublished manuscript.

[25] Lettau, Martin and Sydney Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815-849.

[26] Lettau, Martin and Sydney Ludvigson, 2005, Expected returns and expected dividend growth, *Journal of Financial Economics*, 76, 583-626.

[27] Lettau, Martin, and Stijn Van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607-1652.

[28] Lewellen, Jonathan W., 2004, Predicting returns with financial ratios, *Journal of Financial Economics* 74, 209-235.

[29] Lioui, Abraham, 2006, Black-Scholes-Merton revisited under stochastic dividend yields, *Journal of Futures Markets* 26, 703-732.

[30] Lioui, Abraham, and Jesper Rangvid, 2009, Stock return predictability in a monetary economy, Unpublished manuscript.

[31] Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding predicatability, *Journal of Political Economy* 112, 1-47.

[32] Naranjo, Lorenzo, 2009, Implied interest rates in a market with frictions. Unpublished manuscript.

[33] Newey, Whitney K., and Kenneth D. West, 1987, A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703-708.

[34] Rytchkov, Oleg, 2008, Filtering out expected dividends and expected returns, Unpublished manuscript.

[35] Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375-42.

[36] Valkanov, Rossen, 2003, Long-horizon regressions: Theoretical results and applications, *Journal of Financial Economics* 68, 201-32.

[37] Zhou, Hao, 2010, Variance risk premia, asset predictability puzzles, and macroeconomic uncertainty, Unpublished manuscript.

Appendix: *Derivation of the present value model*

Define the variables and specify the environment:

Define log return r_{t+1}, log dividend growth Δd_{t+1}, and log dividend-price ratio dp_t as:

$$r_{t+1} = \log\left[\frac{P_{t+1}+D_{t+1}}{P_t}\right], \quad \Delta d_{t+1} = \log\left[\frac{D_{t+1}}{D_t}\right], \quad dp_t = \log\left[\frac{D_t}{P_t}\right] \tag{46}$$

Let $\mu_t = E_t(r_{t+1})$ be conditional expected return and let $g_t = E_t(\Delta d_{t+1})$ be conditional expected dividend growth. Assume that μ_t and g_t follow $AR(1)$ processes:

$$\mu_{t+1} = \delta_0 + \delta_1(\mu_t) + \varepsilon^{\mu}_{t+1} \tag{47}$$

$$g_{t+1} = \gamma_0 + \gamma_1(g_t) + \varepsilon^{g}_{t+1} \tag{48}$$

$$\Delta d_{t+1} = g_t + \varepsilon^{d}_{t+1} \tag{49}$$

where ε^{μ}_{t+1}, ε^{g}_{t+1} and ε^{d}_{t+1} are zero mean errors.

Derive Campbell and Shiller (1988) present value identity:

Rewrite log returns as:

$$r_{t+1} = dp_t + \Delta d_{t+1} + \log\left[1+\exp(-dp_{t+1})\right] \tag{50}$$

Use first order Taylor expansion to linearize $\log\left[1+\exp(-dp_{t+1})\right]$ around $\overline{dp} = E(dp_t)$:

$$\log\left[1+\exp(-dp_{t+1})\right] \simeq \log\left[1+\exp(-\overline{dp})\right] + \frac{\exp(-\overline{dp})}{1+\exp(-\overline{dp})}\left[-dp_{t+1}+\overline{dp}\right] \tag{51}$$

Define $\rho = \frac{\exp(-\overline{dp})}{1+\exp(-\overline{dp})}$ and $\kappa = \log\left[1+\exp(-\overline{dp})\right] + \rho\overline{dp}$, such that:

$$\log\left[1+\exp(-dp_{t+1})\right] \simeq \kappa - \rho dp_{t+1} \tag{52}$$

Plug (7) into (5) to get an expression for the one-period return:

$$r_{t+1} \simeq \kappa + dp_t + \Delta d_{t+1} - \rho dp_{t+1} \tag{53}$$

Iterate equation (8) forward:

$$dp_t \simeq -\kappa + \rho dp_{t+1} + r_{t+1} - \Delta d_{t+1} \tag{54}$$

$$dp_t \simeq -\kappa + \rho\left(-\kappa + \rho dp_{t+2} + r_{t+2} - \Delta d_{t+2}\right) + r_{t+1} - \Delta d_{t+1} \tag{55}$$

$$dp_t \simeq -\kappa - \kappa\rho + \rho^2 dp_{t+2} + r_{t+1} + \rho r_{t+2} - \Delta d_{t+1} - \rho \Delta d_{t+2} \tag{56}$$

$$dp_t \simeq -\frac{\kappa}{1-\rho} + \rho^{\infty} dp_{t+\infty} + \sum_{j=0}^{\infty} \rho^j \left(r_{t+1+j} - \Delta d_{t+1+j}\right) \tag{57}$$

Assume that $\lim_{j\to\infty} \rho^j dp_{t+j} = 0$ to obtain the Campbell and Shiller (1988) approximation for the log dividend-price ratio (since the relationship holds ex-ante and ex-post, an expectation operator can be added to the right hand sight):

$$dp_t \simeq -\frac{\kappa}{1-\rho} + E_t \sum_{j=0}^{\infty} \rho^j (r_{t+1+j}) - E_t \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j}) \tag{58}$$

Combine Campbell and Shiller (1988) present value identity with the AR(1) processes for the expected return and the expected dividend growth to solve for the dividend-price ratio:

Iterate equations (2) and (3) forward to obtain:

$$E_t(r_{t+1+j}) = \delta_0 \frac{1-\delta_1^j}{1-\delta_1} + \delta_1^j \mu_t \tag{59}$$

$$E_t(\Delta d_{t+1+j}) = \gamma_0 \frac{1-\gamma_1^j}{1-\gamma_1} + \gamma_1^j g_t \tag{60}$$

Work out the expectations of $E_t \sum_{j=0}^{\infty} \rho^j (r_{t+1+j})$ and the $E_t \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j})$:

$$
\begin{aligned}
E_t \sum_{j=0}^{\infty} \rho^j (r_{t+1+j}) &= \sum_{j=0}^{\infty} \rho^j (\delta_0 \frac{1-\delta_1^j}{1-\delta_1} + \delta_1^j \mu_t) && (61)\\
&= \frac{\delta_0}{1-\delta_1} \sum_{j=0}^{\infty} \rho^j - \frac{\delta_0}{1-\delta_1} \sum_{j=0}^{\infty} \rho^j \delta_1^j + \mu_t \sum_{j=0}^{\infty} \rho^j \delta_1^j && (62)\\
&= \frac{\delta_0}{(1-\delta_1)(1-\rho)} - \frac{\delta_0}{(1-\delta_1)(1-\rho\delta_1)} + \mu_t \left(\frac{1}{1-\rho\delta_1} \right) && (63)
\end{aligned}
$$

$$
\begin{aligned}
E_t \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j}) &= \sum_{j=0}^{\infty} \rho^j (\gamma_0 \frac{1-\gamma_1^j}{1-\gamma_1} + \gamma_1^j g_t) && (64)\\
&= \frac{\gamma_0}{1-\gamma_1} \sum_{j=0}^{\infty} \rho^j - \frac{\gamma_0}{1-\gamma_1} \sum_{j=0}^{\infty} \rho^j \delta_1^j + g_t \sum_{j=0}^{\infty} \rho_1^j \gamma_1^j && (65)\\
&= \frac{\gamma_0}{(1-\gamma_1)(1-\rho)} - \frac{\gamma_0}{(1-\gamma_1)(1-\rho\gamma_1)} + g_t \left(\frac{1}{1-\rho\gamma_1} \right) && (66)
\end{aligned}
$$

Finally, insert (18) and (21) in the Campbell and Shiller (1988) present value identity to find the dividend-price ratio:

$$
dp_t \simeq \varphi + \mu_t \left(\frac{1}{1-\rho\delta_1} \right) - g_t \left(\frac{1}{1-\rho\gamma_1} \right) \tag{67}
$$

where $\varphi = -\frac{\kappa}{1-\rho} + \frac{\delta_0}{(1-\delta_1)(1-\rho)} - \frac{\delta_0}{(1-\delta_1)(1-\rho\delta_1)} - \frac{\gamma_0}{(1-\gamma_1)(1-\rho)} + \frac{\gamma_0}{(1-\gamma_1)(1-\rho\gamma_1)}$.

Table I: *Implied dividend yields with different maturities*

This table reports the summary statistics (Panel A) and the unconditional correlations (Panel B) for the implied dividend yields (IDY) with maturities between 3 months (3m) and 7 months (7m). The period is from January 1994 through December 2009.

	IDY (3m)	IDY (4m)	IDY (5m)	IDY (6m)	IDY (7m)
Panel A: Summary statistics					
Mean	0.0207	0.0206	0.0207	0.0206	0.0203
Std. Dev.	0.0116	0.0086	0.0077	0.0074	0.0073
Panel B: Unconditional correlations					
IDY (3m)	1.0000	0.8702	0.7932	0.7776	0.7541
IDY (4m)	.	1.0000	0.9565	0.9260	0.8769
IDY (5m)	.	.	1.0000	0.9781	0.9252
IDY (6m)	.	.	.	1.0000	0.9789
IDY (7m)	.	.	.	.	1.0000

Table II: *Summary statistics*

This table reports the summary statistics (Panel A) and the unconditional correlations (Panel B) for annualized S&P 500 log monthly returns (r_t^M), annualized log monthly dividend growth rates (Δd_t^M), log dividend-price ratio (dp_t), annualized log implied dividend growth rates (idg_t), and log corrected dividend-price ratio (dp_t^{Corr}). The period is from January 1994 through December 2009.

	r_t^M	Δd_t^M	dp_t	idg_t	dp_t^{Corr}
Panel A: Summary statistics					
Mean	0.0733	0.0361	-4.0198	0.0602	-3.8946
Std. Dev.	0.5443	0.1463	0.2687	0.1868	0.5364
Skewness	-0.9465	-0.1728	0.3758	-1.0235	-0.9809
Kurtosis	4.6857	4.2485	2.6228	5.0496	4.6206
AR (1)	0.1292	0.0871	0.9786	0.5286	0.7359
Panel B: Unconditional correlations					
r_t^M	1.0000	-0.0328	-0.0735	0.0725	0.0157
Δd_t^M	.	1.0000	-0.0854	0.1641	0.0761
dp_t	.	.	1.0000	0.3091	0.7248
idg_t	.	.	.	1.0000	0.8793
$dp_t^{Corr.}$	.	.	.	.	1.0000

Table III: *Dividend growth regressions*

This table reports in-sample regression results for predicting annualized S&P 500 log dividend growth. The predictor variables include log dividend-price ratio (dp_t) and log implied dividend growth (idg_t). All of the regressions are based on monthly observations. For regressions with non-overlapping observations ($h = 1$), t-statistics are calculated according to OLS and are reported in parentheses. For regressions with overlapping observations ($h = 2, 3$ or 6), t-statistics are computed according to Hodrick (1992) and are reported in brackets. The period is from January 1994 through December 2009.

Dividend growth horizon (h)	1	2	3	6
Const.	-0.2728	-0.2644	-0.2423	-0.2013
	(-1.728)	[-1.506]	[-1.344]	[-1.164]
dp_t	-0.0767	-0.0749	-0.0696	-0.0599
	(-1.958)	[-1.735]	[-1.571]	[-1.405]
adj. R^2	0.0147	0.0304	0.0373	0.0328
Const.	0.0248	0.0248	0.0261	0.0283
	(2.284)	[2.206]	[2.264]	[2.517]
idg_t	0.1800	0.1983	0.1906	0.1902
	(3.249)	[4.401]	[4.262]	[4.761]
adj. R^2	0.0479	0.1160	0.1501	0.1842
Const.	-0.4919	-0.5019	-0.4713	-0.4415
	(-3.067)	[-2.463]	[-2.239]	[-2.187]
dp_t	-0.1276	-0.1301	-0.1228	-0.1158
	(-3.229)	[-2.626]	[-2.402]	[-2.355]
idg_t	0.2368	0.2562	0.2454	0.2440
	(4.165)	[4.630]	[4.393]	[4.826]
adj. R^2	0.0931	0.2093	0.2661	0.3081

Table IV: *Return regressions*

This table reports in-sample regression results for predicting annualized log S&P 500 returns. The predictor variables include log dividend-price ratio (dp_t), log implied dividend growth (idg_t), and log corrected dividend-price ratio (dp_t^{Corr}). All of the regressions are based on monthly observations. For regressions with non-overlapping observations ($h = 1$), t-statistics are calculated according to OLS and are reported in parentheses. For regressions with overlapping observations ($h = 2, 3$ or 6), t-statistics are computed according to Hodrick (1992) and are reported in brackets. The period is from January 1994 through December 2009.

Return horizon (h)	1	2	3	6
Const.	0.8234	0.9316	0.9726	1.1540
	(1.393)	[1.339]	[1.418]	[1.784]
dp_t	0.1870	0.2137	0.2238	0.2686
	(1.275)	[1.241]	[1.317]	[1.672]
adj. R^2	0.0033	0.0145	0.0264	0.0702
Const.	0.1689	0.5018	0.6385	0.9402
	(0.276)	[0.634]	[0.832]	[1.317]
dp_t	0.0348	0.1137	0.1461	0.2189
	(0.231)	[0.589]	[0.779]	[1.250]
idg_t	0.7075	0.4637	0.3579	0.2173
	(3.271)	[2.066]	[1.712]	[1.201]
adj. R^2	0.0520	0.0503	0.0571	0.0871
Const.	0.9647	0.7863	0.7061	0.6404
	(3.415)	[3.097]	[2.802]	[2.715]
dp_t^{Corr}	0.2292	0.1833	0.1626	0.1458
	(3.192)	[2.743]	[2.445]	[2.333]
adj. R^2	0.0461	0.0529	0.0619	0.0856

Table V: *In-sample bias correction*

This table reports the effect of small sample bias on the statistical significance of predictor variables in the regressions for predicting annualized log monthly S&P 500 returns (r_{t+1}^M). β and $t - stat.$ are the slope estimate and its corresponding t-statistic according to OLS. β^c and $t^c - stat.$ are the slope estimate and its corresponding t-statistic according to Amihud and Hurvich (2004) bias correction methodology (see regression (33) and equation (34) in the main text). The predictor variables include log dividend-price ratio (dp_t) and log corrected dividend-price ratio (dp_t^{Corr}). The correlation between the innovations to the predictor variable and the errors of the predictive regression is denoted by ρ. The period is from January 1994 through December 2009.

Dependent variable: r_{t+1}^M

	OLS	*Bias correction*
	$\beta_x, (t_x)$	$\beta_x^c, (t_x^c), [\rho\ ]$
dp_t	0.1870	-0.0407
	(1.275)	(-0.273)
		[-0.966]
$dp_t^{Corr.}$	0.2292	0.2234
	(3.192)	(3.098)
		[-0.236]

Table VI: *Out-of-sample predictability*

This table reports out-of-sample R^2_{OS} for predicting annualized log monthly S&P 500 returns (r^M_{t+1}) The predictor variables include log dividend-price ratio (dp_t), log corrected dividend-price ratio (dp^{Corr}_t), and log corrected dividend-price ratio adjusted for the look-ahead bias $(dp^{Corr-NLAB}_t)$. The R^2_{OS} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast).

Dependent variable: r^M_{t+1}

	R^2_{OS}
dp_t	-0.0015
dp^{Corr}_t	0.0606
$dp^{Corr-NLAB}_t$	0.0609

Correlation($dp^{Corr}_t, dp^{Corr-NLB}_t$)=0.9985

Table VII: *Trading strategies*

This table reports certainty equivalent (CE) and Sharpe ratio (SR) of a trading strategy based on timing log monthly S&P 500 returns with different predictor variables. The predictor variables include historical average return estimated through period t $(\overline{r_t})$, log dividend-price ratio (dp_t), log corrected dividend-price ratio $\left(dp_t^{Corr}\right)$, and log corrected dividend-price ratio adjusted for the look-ahead bias $\left(dp_t^{Corr-NLAB}\right)$. The CE and the SR are calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). All the values are annualized.

	CE	SR
$\overline{r_t}$	0.0175	-0.1035
dp_t	0.0278	0.0888
dp_t^{Corr}	0.0507	0.4113
$dp_t^{Corr-NLAB}$	0.0512	0.4170

Table VIII: *Alternative predictors: Summary statistics*

This table presents the summary statistics (Panel A) and the unconditional correlations (Panel B) for the log dividend-price ratio (dp_t), log dividend-price ratio corrected for the implied dividend growth $\left(dp_t^{Corr}\right)$, log dividend-price ratio corrected for the 12 month average monthly dividend growth $\left(dp_t^{HIST}\right)$, variance risk premia (vrp_t), log earnings-price ratio (ep_t), and consumption-to-wealth ratio (cay_t) (monthly observations of cay_t are defined by the most recently available quarterly observation). The period is from January 1994 through December 2009.

	dp_t	dp_t^{Corr}	dp_t^{HIST}	vrp_t	ep_t	cay_t
Panel A: Summary statistics						
Mean	-4.0198	-3.8946	-3.9282	18.2236	-3.2247	-0.0016
Std. Dev.	0.2687	0.5364	0.3007	22.2653	0.4179	0.0204
Skewness	0.3758	-0.9809	-0.2751	-2.8320	-2.2244	0.1480
Kurtosis	2.6228	4.6206	2.1317	36.9296	9.0900	1.8789
AR (1)	0.9786	0.7359	0.9687	0.3014	0.9419	0.9669
Panel B: Unconditional correlations						
dp_t	1.0000	0.7248	0.8858	-0.0979	0.1497	0.4858
$dp_t^{Corr.}$	.	1.0000	0.7416	-0.1146	0.3286	0.3402
dp_t^{HIST}	.	.	1.0000	-0.2295	0.4906	0.2622
vrp_t	.	.	.	1.0000	-0.3095	0.1349
ep_t	.	.	.	.	1.0000	0.0224
cay_t	.	.	.	.	.	1.0000

Table IX: *Alternative predictors: Monthly return regressions*

This table presents in-sample and out-of-sample results for predicting annualized log monthly S&P 500 returns $\left(r_{t+1}^{M}\right)$. In-sample results are based on the period from January 1994 through December 2009. Out-of-sample R_{OS}^{2} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). The predictor variables include log dividend-price ratio (dp_t), log dividend-price ratio corrected for the implied dividend growth $\left(dp_t^{Corr}\right)$, log dividend-price ratio corrected for the 12 month average monthly dividend growth $\left(dp_t^{HIST}\right)$, variance risk premia (vrp_t), log earnings-price ratio (ep_t), and consumption-to-wealth ratio (cay_t).

Dependent variable: r_{t+1}^{M}

Const.	0.8234	0.9647	0.6552	-0.0233	-0.0695	0.0765	0.9045	0.9671	0.9454
	(1.393)	(3.415)	(1.261)	(-0.467)	(-0.221)	(1.943)	(1.565)	(3.525)	(1.842)
dp_t	0.1870						0.2320		
	(1.275)						(1.611)		
dp_t^{Corr}		0.2292						0.2576	
		(3.192)						(3.667)	
dp_t^{HIST}			0.1486						0.2503
			(1.126)						(1.897)
vrp_t				0.0052			0.0054	0.0059	(0.0059)
				(3.007)			(3.165)	(3.506)	(3.381)
ep_t					-0.0438				
					(-0.452)				
cay_t						3.1200			
						(1.621)			
adj. R^2	0.0033	0.0461	0.0014	0.0406	-0.0042	0.0085	0.0487	0.0999	0.0537
R_{OS}^{2}	-0.0015	0.0606	-0.0033	0.0450	-0.0157	0.0083	0.0447	0.1155	0.0496

Table X: *Robustness check: Maturity of implied dividend yield I*

This table presents in-sample and out-of-sample results for predicting annualized log monthly S&P 500 returns $\left(r^M_{t+1}\right)$. In-sample results are based on the period from January 1994 through December 2009. Out-of-sample R^2_{OS} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). The predictor variables include log dividend-price ratio (dp_t), and log corrected dividend-price ratio $\left(dp^{Corr}_t(T)\right)$ defined as $dp^{Corr.}_t(T) = dp_t + 2.08 * (idy^T_t - dp_t)$, where idy^T_t is log implied dividend yield with maturities between 3 months $(3m)$ and 7 months $(7m)$.

Dependent variable: r^M_{t+1}

Const.	0.4054	0.4154	0.8284	0.9647	1.0117	0.8234
	(2.905)	(2.595)	(3.283)	(3.415)	(3.458)	(1.393)
$dp^{Corr}_t(3m)$	0.0849					
	(2.491)					
$dp^{Corr}_t(4m)$		0.0871				
		(2.215)				
$dp^{Corr}_t(5m)$			0.1939			
			(3.035)			
$dp^{Corr}_t(6m)$				0.2292		
				(3.192)		
$dp^{Corr}_t(7m)$					0.2408	
					(3.241)	
dp_t						0.1870
						(1.275)
adj. R^2	0.0267	0.0201	0.0414	0.0461	0.0477	0.0033
R^2_{OS}	0.0365	0.0268	0.0560	0.0606	0.0608	-0.0015

Table XI: *Robustness check: Maturity of implied dividend yield II*

This table presents in-sample and out-of-sample results for predicting annualized log monthly S&P 500 returns $\left(r_{t+1}^{M}\right)$. In-sample results are based on the period from January 1994 through December 2009. Out-of-sample R_{OS}^{2} is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). The predictor variables include annualized log dividend-price ratio based on dividends summed over the past 6 months $\left(dp_t^{6m}\right)$ and log corrected dividend-price ratio $\left(dp_t^{Corr}(T)\right)$ defined as $dp_t^{Corr}(T) = dp_t^{6m} + 2.08*(idy_t^T - dp_t^{6m})$, where idy_t^T is log implied dividend yield with maturities between 3 months $(3m)$ and 7 months $(7m)$.

Dependent variable: r_{t+1}^{M}

Const.	0.4066	0.4166	0.8327	0.9701	1.0154	0.8303
	(2.907)	(2.597)	(3.290)	(3.423)	(3.462)	(1.396)
$dp_t^{Corr}(3m)$	0.0850					
	(2.494)					
$dp_t^{Corr}(4m)$		0.0872				
		(2.217)				
$dp_t^{Corr}(5m)$			0.1946			
			(3.042)			
$dp_t^{Corr}(6m)$				0.2301		
				(3.200)		
$dp_t^{Corr}(7m)$					0.2412	
					(3.246)	
dp_t^{6m}						0.1890
						(1.278)
adj. R^2	0.0268	0.0202	0.0417	0.0464	0.0478	0.0033
R_{OS}^2	0.0366	0.0269	0.0563	0.0609	0.0610	-0.0000

Table XII: *Robustness check: Options moneyness and backward-looking data*

This table presents in-sample and out-of-sample results for predicting annualized log monthly S&P 500 returns $\left(r_{t+1}^{M}\right)$. In-sample results are based on the period from January 1994 through December 2009. Out-of-sample R_{OS}^2 is calculated over the period January 2000 through December 2009 (the period from January 1994 through December 1999 is used to make the first forecast). The predictor variables include log dividend-price ratio (dp_t) and log corrected dividend-price ratio $\left(dp_t^{Corr}\right)$ defined as $dp_t^{Corr} = dp_t + 2.08 * (idy_t - dp_t)$, where idy_t is log implied dividend yield calculated from different number of days of backward-looking data (d), or based on restricted moneyness levels of options (M).

Dependent variable: r_{t+1}^{M}

Const.	0.9647	0.9743	0.8224	0.6675	0.9849	0.8234
	(3.415)	(3.537)	(3.847)	(2.639)	(3.309)	(1.393)
dp_t^{Corr}	0.2292					
	(3.192)					
$dp_t^{Corr}(0.8 \leq M \leq 1.2)$		0.2315				
		(3.309)				
$dp_t^{Corr}(0.9 \leq M \leq 1.1)$			0.1900			
			(3.570)			
$dp_t^{Corr}(-5 \leq d \leq 0)$				0.1531		
				(2.385)		
$dp_t^{Corr}(-15 \leq d \leq 0)$					0.2341	
					(3.095)	
dp_t						0.1870
						(1.275)
adj. R^2	0.0461	0.0498	0.0582	0.0241	0.0432	0.0033
R_{OS}^2	0.0606	0.0632	0.0594	0.0235	0.0599	-0.0015

Figure 1: *Open interest by maturity*

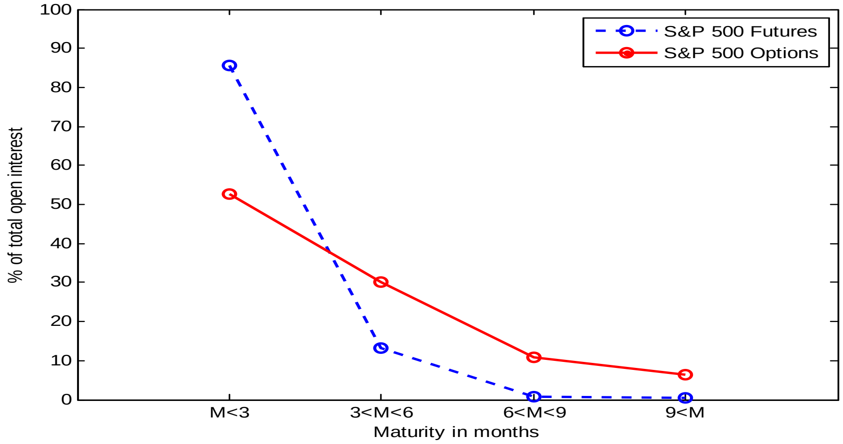

This figure plots the percentage of open interest by maturity for S&P 500 options and S&P 500 futures. The percentage of open interest is calculated as the total open interest for a given maturity over the total open interest for all the maturities. The period is from January 1994 through December 2009.

Figure 2: *Term structure for the implied dividend yields*

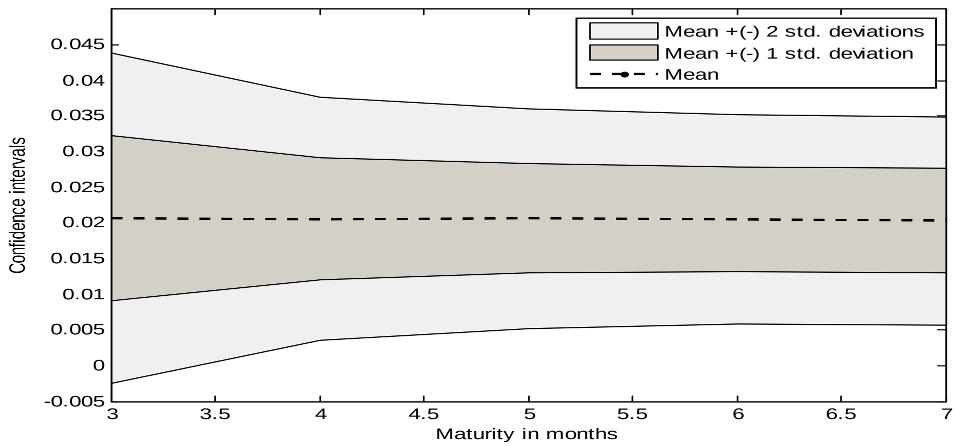

This figure plots the term structure for the implied dividend yields. The figure is based on the summary statistics for the implied dividend yields reported in Table I. Dashed line denotes the mean of the implied dividend yields with different maturities. Dark grey color denotes the area that is one standard deviation away from the mean of the implied dividend yields. Bright grey color denotes the area that is two standard deviations away from the mean of the implied dividend yields. The period is from January 1994 through December 2009.

Figure 3: *Implied dividend growth*

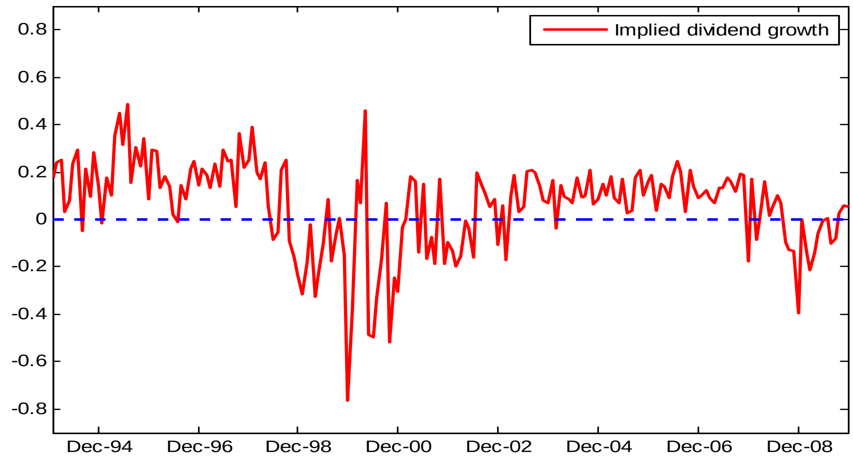

This figure plots log implied dividend growth for the S&P 500. The period is from January 1994 through December 2009.

Figure 4: *Dividend-price ratio and corrected dividend-price ratio*

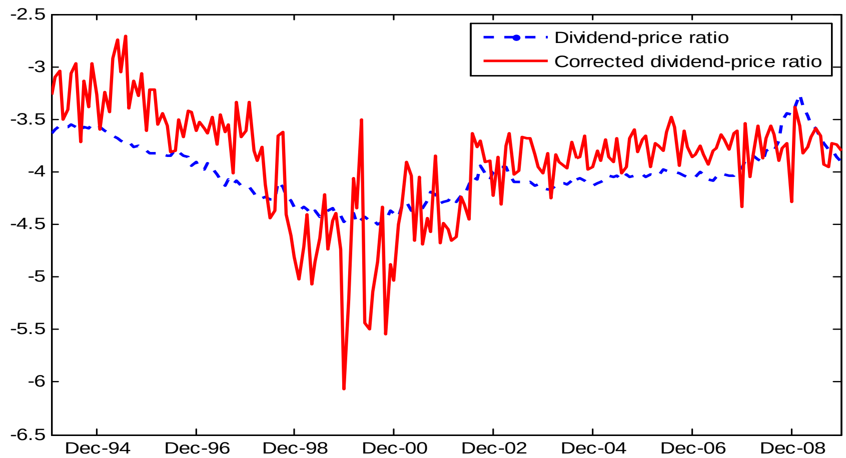

This figure plots log dividend-price ratio and log corrected dividend-price ratio for the S&P 500. The period is from January 1994 through December 2009.

Figure 5: *Implied interest rate, T-bill rate and LIBOR rate*

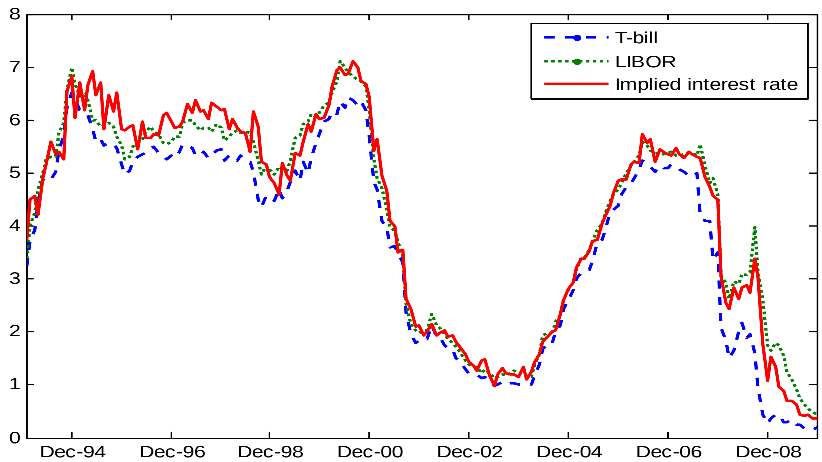

This figure plots the 6-month implied interest rate along with the 6-month T-bill rate and the 6-month LIBOR rate. The period is from January 1994 through December 2009.

Figure 6: *Realized vs. forecasted returns*

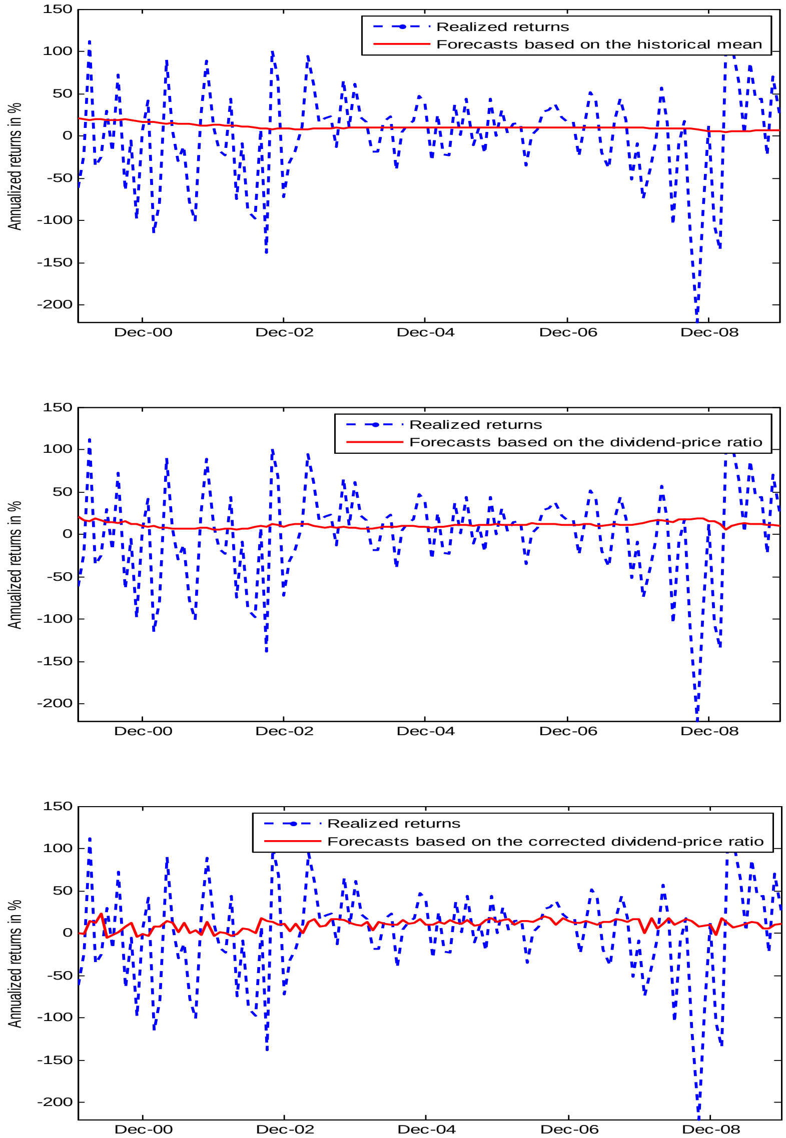

The figures plot annualized monthly S&P 500 log returns along with the return forecasts. The forecasts are based either on historical average return (upper figure), on the dividend-price ratio (middle figure), or on the corrected dividend-price ratio (lower figure). The period is from January 1994 through December 2009.