Subjective Cash Flows and Discount Rates*

Ricardo De la O
Stanford University

Sean Myers
Stanford University

December 4, 2017

Abstract

What drives stock prices? Using survey forecasts for dividend growth and returns for the S&P 500 index, we find that changes in subjective dividend growth expectations are the main driver of movements in the price-dividend ratio. Subjective dividend growth expectations vary substantially over time and match future dividend growth remarkably well, while subjective return expectations are relatively flat and weakly negatively correlated with future returns. One-year subjective dividend growth expectations explain a large amount of the movements in the price-dividend ratio, accounting for 36% of the total variation. Using longer horizon subjective expectations, we estimate that subjective dividend growth expectations account for at least 73% of the variation in the price-dividend ratio, while subjective return expectations account for at most 27%. These findings highlight the importance of time-varying dividend growth expectations in determining aggregate stock prices.

1 Introduction

A central question in finance is what drives stock price movements. Specifically, we want to know what drives the large movements in the aggregate price-dividend ratio, a measure of how cheap or expensive stocks are at a given time. A stock's price should equal the expected discounted value of future dividends. An increase in the price-dividend ratio, i.e. a market boom, must therefore be due to higher expected dividend growth or lower expected returns. Similarly, a sudden drop in the price-dividend ratio, such as the market bust in late 2008, must be due to a drop in expected dividend growth or investors demanding higher expected returns.

*We thank John Campbell, John Cochrane, Robin Greenwood, Samuel Hanson, Ben Hébert, Lawrence Jin, Rafael La Porta, Hanno Lustig, Monika Piazzesi, Martin Schneider, Andrei Shleifer, Ken Singleton, and seminar participants at the Harvard Business School Behavioral Finance and Financial Stability Group, 2017 Chicago Initiative in Theory and Empirics, the Stanford GSB Finance lunch, and the Stanford Macroeconomics lunch for valuable comments. Email addresses: delao@stanford.edu, seanmyers@stanford.edu.

Regressions using historical price and dividend data for the S&P 500 index have shown that a high price-dividend ratio is typically not followed by high future dividend growth but by low future returns (Campbell and Shiller [1988a], Cochrane [2008, 2010]). This result has motivated many models where price movements are driven by agents' time-varying expected returns, e.g. habit formation, stochastic volatility, and time variation in disaster probabilities. Many of these models assign little or no importance to fluctuations in expected dividend growth for explaining price movements.

This paper measures investors' subjective expectations using survey data on the S&P 500 index and finds that movements in the price-dividend ratio are predominantly explained by changes in subjective dividend growth expectations, not changes in subjective return expectations. During a market boom both subjective dividend growth expectations and subjective return expectations are high. The high subjective dividend growth expectations give direct evidence that the market boom can be explained by investors expecting high future dividends. The high subjective return expectations mean that the market boom cannot be due to lower expected discount rates, which provides indirect evidence that the market boom must be driven by increased expected dividends. Likewise, market busts are associated with low subjective dividend growth expectations and low subjective return expectations. The fall in the price-dividend ratio is not driven by investors demanding higher returns (possibly to compensate for increased risk), but instead by investors expecting lower future dividends. Based on the direct evidence from subjective dividend growth expectations and the indirect evidence from subjective return expectations, we estimate that subjective expectations of future dividend growth explain at least 73% of the variation in the price-dividend ratio.

For 2003Q1-2015Q3, we construct one-year and two-year subjective expectations of S&P 500 dividends from the Thomson Reuters I/B/E/S Estimates Database by aggregating analyst forecasts for individual firms in the S&P 500. One-year subjective dividend growth expectations are highly correlated with future one-year dividend growth and vary substantially over time. For the same sample, we collect one-year and ten-year subjective expectations of S&P 500 returns from the quarterly Graham-Harvey Global Business Outlook Survey, which surveys CFO's of major U.S. corporations. We also examine larger sample periods going back to 1952 by constructing subjective return expectations from additional surveys. In contrast to subjective dividend growth expectations, we find that subjective return expectations are quite flat over time and one-year subjective return expectations are weakly negatively correlated with future one-year returns. Longer horizon subjective expectations show that investors do not believe that changes in short-term dividend growth or returns will be persistent. We refer to this as low expected persistence, meaning that changes in short-term subjective expectations are only associated with small changes in long-term subjective expectations.

Using the survey data and a variance decomposition introduced by Campbell and Shiller [1988b], we calculate how much of the variation in the price-dividend ratio comes from changes in subjective dividend growth and return expectations. We find three key results: a large positive contribution from subjective dividend growth expectations, a small negative contribution from subjective return expectations, and a dominance of short-term subjective expectations. Subjective dividend growth expectations vary significantly over time and are high when the price-dividend ratio is high, explaining most of the movement in the price-dividend ratio. Changes in subjective return expectations cannot explain much movement in the price-dividend ratio because these expectations are relatively flat over time. Further, since subjective return expectations are slightly higher when the price-dividend ratio is high, they actually make a negative contribution to explaining the variation in the price-dividend ratio. Rather than being the driving force behind price volatility, we find that changes in subjective return expectations if anything dampen the movements in the price-dividend ratio. Lastly, we find that the price-dividend ratio is largely driven by changes in short-term subjective dividend growth expectations, with one-year subjective dividend growth expectations accounting for 36% of the variation in the price-dividend ratio. Because expected persistence is low, long-term subjective expectations do not play a large role in explaining price movements.

Based on the results of the decomposition, we test a simple asset pricing model in which expected returns are constant, one-year dividend growth expectations are taken from the survey data, and expected persistence is low. The first two tests focus on the prices of short-term and long-term claims to the S&P 500. Volatile one-year dividend growth expectations and constant return expectations imply that the price of a one-year dividend strip (i.e. a claim to all dividends from the S&P 500 for the next year) should be quite volatile. To test this, we compare the strip price implied by the model with the one-year dividend strip price calculated by Van Binsbergen, Brandt, and Koijen [2012] using option price data. Second, low expected persistence implies that shocks to short-term dividend growth expectations do not have large effects on long-term dividend growth expectations and the price of a long-term claim to the S&P 500 should reflect this. We compare the model implied price-dividend ratio for the S&P 500 to the observed price-dividend ratio to test if the observed value is consistent with volatile short-term dividend growth expectations and low expected persistence. The third test focuses on the accuracy of the model for predicting future price movements. If shocks to dividend growth expectations are not persistent then movements in the price-dividend ratio should only be temporary. We test this implication by comparing the model forecast for the one-year change in the price-dividend ratio to the future one-year observed change in the price-dividend ratio. For the one-year strip price, the price-dividend ratio, and the change in the price-dividend ratio, the model values match the observed values quite well. Regressing the observed values on the model values gives significant coefficients of 1.01, 0.96, and 0.91 and high R^2's of 0.75, 0.70, and 0.36 respectively.

In the final section we perform robustness checks on our main results. First, we show that the failure of subjective return expectations to explain price movements is not limited to the Graham-Harvey survey respondents nor the sample period. This result holds over the 5 return surveys, with one going back to 1952, validating the results for different sample periods, forecast horizons, and respondent groups (CFO's, professional forecasters, consumers, etc.). This implies that subjective dividend growth expectations, not subjective return expectations, must be driving price movements over these samples. Then, we explore the possibility that survey responses do not represent a direct measure of subjective expectations. For example, a high response on the dividend survey may simply mean that investors are optimistic about dividends, but the reported amount may not literally represent the expected value. We construct an alternative model in which the survey responses are simply treated as data about investors' expectations and construct fitted expectations by regressing future one-year returns and dividend growth on the survey responses. We then use these fitted expectations in the decomposition and the result are almost identical with our main decomposition. Finally, we relax the no-bubble condition for our price decomposition and measure the potential for bubbles to explain price movements. We find that the possibility of bubbles in subjective expectations does not affect the contribution of dividend growth expectations in explaining the volatility of the price-dividend ratio. Subjective dividend growth expectations still explain at least 73% of the volatility of the price-dividend ratio regardless of the existence of bubbles.

There is a growing literature challenging the irrelevance of expected dividend growth and the dominance of expected returns in driving prices. Chen and Zhao [2009] criticize the method of inferring the role of dividend growth expectations indirectly from return expectations. The reason is that the small predictive power of the price-dividend ratio on future returns may create a large misspecification error that gets inherited by the expected dividend growth. One of the advantages of our study is that we can estimate the role of dividend growth expectations and return expectations separately and obtain similar quantitative results. Others like Ang [2012] and Koijen and Van Nieuwerburgh [2011] argue that sample selection and reinvesting dividends in a particular way (i.e. changing the calculation of dividend growth) could lead to results were dividend growth expectations have nontrivial significance. Our paper sticks to the standard definitions of dividend growth and returns and shows that even though price-dividend ratio movements in our sample are followed by changes in observed returns, subjective return expectations have almost no contribution to price-dividend ratio movements.

Our paper also relates to literature identifying the contribution of short-term assets to price volatility. Using S&P 500 option prices, Van Binsbergen, Brandt, and Koijen [2012] measure the price of dividend strips at different horizons. They conclude that excess volatility in the aggregate stock market must be explained by excess volatility in short-term dividend strip prices. This implies that short-term expectations play an important role in price volatility, which matches our results. We further refine this finding by showing that short-term subjective dividend growth expectations, rather than subjective return expectations, drive short-term dividend strip price volatility.

There is also a recent interest in using surveys of expectations, rather than statistical expectations based on regressions, to understand price movements. Greenwood and Shleifer [2014] use a wide variety of surveys to show that subjective expectations of stock market returns are negatively correlated with statistical expectations of future returns. Piazzesi, Salomao, and Schneider [2015] and Koijen, Schmeling, and Vrugt [2015] have found subjective return expectations to be negatively correlated with future returns and statistical expectations of returns in bond, stock, and currency markets both in the US and abroad. We confirm their findings about subjective stock return expectations and complement them with an independent measure of subjective dividend growth expectations which shows investors do have accurate expectations about dividends. Chen, Da, and Zhao [2013] cast doubt on the importance of return expectations in moving prices using firm-level earnings expectations data. Their objects of study are different from the traditional dividend growth and return decomposition in the existing literature, as they decompose the price into the implied cost of equity capital, which they treat as a measure of discount rates, and a residual measure of cash flows. They show that the implied cost of equity capital cannot explain the majority of price movements and infer that their measure of cash flows must then be driving the price movements. We measure subjective expectations for the standard cash flow and discount rate definitions, namely aggregate stock dividend growth and returns, and are able to directly observe both the cash flow component and discount rate component.

Recent theoretical models (e.g. Adam, Marcet, and Beutel [2015], Barberis, Greenwood, Jin, and Shleifer [2015], Jin and Sui [2017]) have been proposed to reconcile the negative correlations of one-year subjective return expectations and future returns via extrapolative or highly persistent expectations. Our paper sheds light beyond this one-year correlation and shows three salient patterns in subjective expectations that should be used to compare the predictions of asset pricing models. First, investors report significant time-varying subjective dividend growth expectations that are positively correlated with the price-dividend ratio, or in other words, changes in subjective dividend growth expectations must be important for price determination. Second, subjective return expectations have low volatility and no (or positive) comovement with the price-dividend ratio. Thus, they do not play a large role in driving price movements. Third, both subjective dividend growth and return expectations show low expected persistence, implying that investors believe changes in dividend growth and returns will not be persistent. The first and second patterns will pose a challenge for most rational expectations models since the price-dividend ratio is not strongly correlated with observed future dividend growth and is negatively correlated with observed future returns.[1] The third

[1] Further, under rational expectations, dividend growth expectations cannot be more volatile than return expectations. This is because the price-dividend ratio does not significantly comove with observed future dividend growth. Any information that raises dividend growth expectations must also raise return expectations by an equal amount, otherwise the price-dividend ratio would rise and there would be comovement between the price-dividend ratio and observed future dividend growth. We thank John Cochrane for making this point.

pattern, on the other hand, would be difficult to match by long-run risk or learning models where shocks cause agents to update both their short-term and long-term expectations.

The sections are organized as follows. Section 2 introduces the Campbell-Shiller decomposition and discusses our approach in light of its current treatment in the literature. Sections 3 and 4 describe the data sources and compare the subjective expectations with observed future one-year dividend growth and returns. Section 5 calculates the role of dividend growth and returns in explaining movements in the price-dividend ratio using subjective expectations as well as a statistical expectations benchmark. Section 6 tests a model of time-varying dividend growth expectations taken from the survey data and constant discount rates by comparing the model implied asset prices to observed asset prices. Section 7 covers robustness checks on our main findings. Section 8 concludes.

2 Decomposing Price Movements

A stock's price is the discounted value of future dividends, which means the value of the S&P 500 is the discounted value of future dividends paid by the constituent firms. The vast majority of S&P 500 firms pay dividends. By market value, the dividend paying firms represent 80-90% of the entire index over our sample, and dividends are the main method by which S&P 500 firms make cash distributions to their share holders[2]. Campbell and Shiller [1988b]'s log-linearization of the return identity states the price-dividend ratio in terms of future returns, r_{t+1}, future dividend growth, Δd_{t+1}, and the future price-dividend ratio, pd_{t+1}, all in logs:

$$pd_t = \kappa + \Delta d_{t+1} - r_{t+1} + \rho pd_{t+1}, \tag{1}$$

where κ is a constant, $\rho = \frac{e^{\bar{pd}}}{1+e^{\bar{pd}}} < 1$ and $\bar{pd}$ is the mean value of the log price-dividend ratio. By imposing a no-bubble condition[3], $\lim_{T\to\infty} \rho^T pd_{t+T} = 0$, we can iterate the equation and apply expectations to write the price-dividend ratio as the sum of a constant plus two main factors,

$$pd_t = \frac{1}{1-\rho}\kappa + \sum_{j=1}^{\infty} \rho^{j-1} E_t \left[\Delta d_{t+j}\right] - \sum_{j=1}^{\infty} \rho^{j-1} E_t \left[r_{t+j}\right]. \tag{2}$$

This equation directly shows that an increase in the price-dividend ratio must be due to higher expected dividend growth or lower expected returns. While equation (2) also holds without expectations, applying expectations makes all of the components known at time t. This allows us to see what drives the change in the price-dividend ratio, i.e. did prices rise because investors are optimistic about future dividends or because they are expecting lower returns. To measure the relative importance of dividend growth expectations and

[2]Dividends represent 80% of total payouts made by S&P 500 firms over our sample, where total payouts are measured as dividends plus share repurchases minus share issuance.

[3]In section 7.3 we analyze the implications of relaxing the no-bubble condition.

return expectations, we separate the variance of the price-dividend ratio into its covariance with expected dividend growth and its covariance with expected returns to get the following decomposition:

$$1 = \frac{Cov\left(\sum_{j=1}^{\infty}\rho^{j-1}E_t\left[\Delta d_{t+j}\right], pd_t\right)}{Var(pd_t)} + \frac{Cov\left(-\sum_{j=1}^{\infty}\rho^{j-1}E_t\left[r_{t+j}\right], pd_t\right)}{Var(pd_t)}. \tag{3}$$

There are two possible approaches to empirically address the relative importance of these two components. A common approach in the literature is to statistically infer these expectations from historical dividend and price data. These *statistical expectations* will satisfy the property $Cov\left(\sum_{j=1}^{T}E_t\left[\Delta d_{t+j}\right], pd_t\right) = Cov\left(\sum_{j=1}^{T}\Delta d_{t+j}, pd_t\right)$ for every horizon T as long as the price-dividend ratio is used in the inference. Consequently the component of expected dividend growth can be approximated by $\frac{Cov\left(\sum_{j=1}^{T}\rho^{j-1}\Delta d_{t+j}, pd_t\right)}{Var(pd_t)}$, which under stationarity is just the OLS coefficient of a simple regression of future dividend growth on pd_t. Similarly, we can obtain the contribution of expected returns by regressing future returns on pd_t. Findings by Campbell and Shiller [1988b], Fama and French [1988], Cochrane [2008] and others suggest that the contribution of expected dividend growth is virtually zero and all price-dividend ratio movements must come from expected returns. Since the price-dividend ratio does not covary with observed future dividend growth, many economic models assume expected dividend growth is constant or unimportant for stock market volatility and that time-varying risk premia are the primary factor driving prices in the economy.

There is a second approach to measure the importance of dividend growth expectations versus return expectations. Rather than inferring the statistical expectations, we can directly measure the expectations held by investors at each point in time. We use forecast surveys to construct robust measures of dividend growth and return expectations at different horizons. With these *subjective expectations* denoted by $E^*[\cdot]$, we revisit the relative importance of the two components of price-dividend ratio volatility. This way, we can re-evaluate if the current models of time-varying risk premia and constant dividend growth expectations align with actual investor expectations or if more focus should be placed on modeling agents with large time-varying dividend growth expectations.

3 Data and Variable Construction

In this section, we explain the data sources used for our main calculations and the construction process to build the aggregate dividend growth and return expectations.

3.1 S&P 500 Index

From Compustat, we create a list of all companies in the S&P 500 at the end of each quarter and record their price per share, dividends per share and number of outstanding shares. We calculate a quarterly dividend measure for the index by aggregating the total ordinary dividends paid by each company and adjusting them by the S&P 500 index divisor. We build the S&P 500 index divisor by taking the total market capitalization of the S&P 500 companies and dividing by the S&P 500 index at the end of each quarter.

3.2 Subjective Dividend Growth Expectations

We construct subjective dividend growth expectations for the S&P 500 using the Summary Statistics of the Thomson Reuters I/B/E/S Estimates Database. I/B/E/S is a comprehensive forecast database containing analyst estimates for more than 20 forecast measures - including DPS (dividends per share). The Summary Statistics contains the average forecasts on different horizons for U.S. publicly traded companies. We build a measure of the aggregate dividend expectations using the constituents of the S&P 500 at each point in time. Constituents are weighted by the market value of outstanding common shares. This procedure is analogous to the process in which the S&P 500 index is calculated and is explicitly derived in Appendix A.1.

Two features of the data must be dealt with in order to calculate aggregate dividend expectations. First, the I/B/E/S database contains DPS estimates for up to five Annual Fiscal Periods (FY1-FY5), four Quarter Fiscal Periods (Q1-Q4) and a Long Term Growth measure. Because not all companies have the same fiscal year end, we interpolate across the different horizons to obtain a precise expectation of dividends over the next twelve months following the response of the analyst. For example, if the fiscal year of Firm A ends in 9 months after a given point in time, we may only have available a 9-month dividend expectation and a 21-month dividend expectation for that firm. We interpolate these two measures in order to ensure that every expectation is exactly twelve months ahead. We repeat the analogous procedure to construct two-year expectations. The second feature of the data is that I/B/E/S expectations were not available for all S&P 500 companies. We take the aggregate dividend expectation of those companies in the S&P 500 with expectations available and multiply it by the ratio of total S&P 500 market value to the market value of the forecasted companies. The assumption behind this normalization is that the forecasted companies are a representative sample of the S&P 500. We test this assumption in the tables below and find that it holds quite well. Appendix A.1-A.2 gives more detail on our methodology.

Table 1 shows tests of our dividend construction methodology. Since we cannot know the expected dividends for companies that do not report forecasts, we test our methodology using realized dividends. We construct an aggregate realized dividend using the same method applied to our subjected expected dividend.

Table 1: Correlations of S&P 500 dividend measures

	Levels			Growth		
	Forecasted	Shiller	SPY	Forecasted	Shiller	SPY
All Companies	0.997	0.995	0.994	0.935	0.930	0.916
Forecasted		0.995	0.992		0.935	0.919
Shiller			0.997			0.953

Note. The table features four quarterly time series spanning 2003Q1-2015Q3. *All Companies* contains the aggregate quarterly dividends paid out by all S&P 500 companies. *Forecasted* contains the aggregate quarterly dividends paid out only by S&P 500 companies for which a one-year subjective dividend expectation exists. *Shiller* contains the quarterly S&P 500 dividends obtained from Shiller [2015]. *SPY* contains the quarterly dividends paid out by the SPDR S&P 500 ETF. Under *Levels* columns, we calculate the correlation of the four series. Under *Growth* columns, we calculate the annual percent change of each of the four series and then take its correlation.

The first three columns of Table 1 give the correlation of our aggregate dividend measure with Robert Shiller's S&P 500 dividend and the dividend for SPY, a popular S&P 500 replicating ETF. The first dividend measure is our aggregate dividend using all companies in the S&P 500. The high correlation of this measure with Shiller and SPY dividends shows that our aggregation technique is accurate. The second measure is identical to the first, except it only uses companies for which we have a one-year subjective dividend expectation and is scaled by the ratio of total S&P 500 market value to total forecasted companies market value. The high correlation between the first two measures shows that the forecasted companies are representative of the entire set of constituents.

The second set of columns in Table 1 show the correlation of dividend growth for each of the four measures. As before, the high correlation between All Companies dividend growth and Forecasted Companies dividend growth shows that the reporting companies are a representative subset. The high correlation of these two measures with Shiller and SPY dividend growth shows that our dividend aggregation procedure is accurate.

3.3 Subjective Return Expectations

Our main measure of subjective return expectations is taken from a survey conducted by John Graham and Campbell Harvey of Duke University's Fuqua School of Business. The survey is completed by 200 to 500 chief financial officers (CFO's) of major U.S. corporations quarterly. Among other things, the survey solicits CFO views about the U.S. economy. In particular, they report their expectations of returns on the S&P 500 index over the next twelve months, which we will label $E_t^*[r_{t+1}]$, and their expectations of the average annual returns over the next ten years. The sample includes CFO's from both public and private companies representing a broad range of industries, geographic areas and sizes. The data is available from the third quarter of 2000 until the second quarter of 2016. For robustness, we show in the Section 7.1 that our results hold for 5 different surveys of return expectations (2 one year return surveys, 1 two to three year return survey and 2 ten year return surveys) that span different populations and samples.[4]

4 Short-Term Subjective Expectations

4.1 Dividend growth

In this section, we take a first look at the data and how the short-term subjective expectations perform in the face of future and current dividend growth and returns. Figure 1 shows that subjective dividend growth expectations, denoted as $E_t^*[\Delta d_{t+1}]$, have similar volatility to that of observed future one-year dividend growth and track it quite well. Table 2 shows subjective dividend growth expectations are strongly correlated with observed future dividend growth, with a correlation of 0.75. The accuracy of subjective dividend growth expectations makes it unlikely that investors are not reporting their actual dividend growth expectations or are altering their responses to the surveys to rationalize the current prices.

At first sight, one may think that the high correlation between $E_t^*[\Delta d_{t+1}]$ and Δd_{t+1} is just due to high persistence in the dividend growth process. For instance, if Δd_{t+1} follows an AR1 process, $\Delta d_{t+1} = \phi \Delta d_t + \varepsilon_{t+1}$, with a high persistence then current dividend growth would carry useful information about future dividend growth, and we would expect a high correlation of subjective dividend growth expectations and observed future one-year dividend growth. However, the correlation between $E_t^*[\Delta d_{t+1}]$ and Δd_t, though positive, is noticeably lower (0.22) than $\text{corr}(E_t^*[\Delta d_{t+1}], \Delta d_{t+1})$, suggesting there is a component of the prediction that is unexplained by current dividend growth.

4.1.1 Earnings growth

The price decomposition in equation (2) shows that dividend growth expectations are our cash flow measure of interest. However, we show that the main features of subjective dividend growth expectations also hold for subjective earnings growth expectations. Subjective dividend growth expectations and subjective earnings growth expectations both vary substantially over time and they are highly correlated with future dividend

[4] The different sampling periods, methodology and population targets of these extra surveys makes them excellent measures of external validation for our main results. In addition to the Graham-Harvey one-year and ten-year return expectations, we use the Federal Reserve Bank of Philadelphia's Livingston Survey (1952-2016), the University of Michigan Survey of U.S. consumers (2000-2005), and the Survey of Professional Forecasters (1992-2016). We choose the Graham-Harvey survey as our main source for return expectations because it provides both short-term and long-term return forecasts and aligns with our dividend forecast sample. A more detailed description of the additional surveys is available in Appendix A.4.

Figure 1: Expected and realized one-year dividend growth

The figure compares the one-year subjective dividend growth expectation and the observed future one-year dividend growth for the S&P 500. The solid line is the one-year subjective dividend growth expectation based on survey data. The dotted line is the observed future one-year dividend growth.

Table 2: Correlation of expected and observed dividend growth and returns

	Δd_{t+1}	Δd_t	pd_t		r_{t+1}	r_t	pd_t
$E_t^*[\Delta d_{t+1}]$	0.75	0.22	0.84	$E_t^*[r_{t+1}]$	-0.03	0.41	0.65
	(0.24)	(0.22)	(0.09)		(0.14)	(0.17)	(0.11)
Δd_{t+1}		0.29	0.70	r_{t+1}		-0.09	-0.20
		(0.17)	(0.17)			(0.16)	(0.17)
Δd_t			0.26	r_t			0.63
			(0.20)				(0.17)

Note. The table shows pairwise correlations using quarterly data from 2003Q1-2015Q3. $E_t^*[\Delta d_{t+1}]$ and $E_t^*[r_{t+1}]$ are the one-year subjective dividend growth and return expectations calculated from I/B/E/S and the G-H return survey respectively. Δd_t, Δd_{t+1}, r_t , r_{t+1} and pd_t are the observed current and future one-year dividend growth and returns on the S&P 500 and the price-dividend ratio for the S&P 500. Small-sample adjusted Newey-West standard errors in parenthesis.

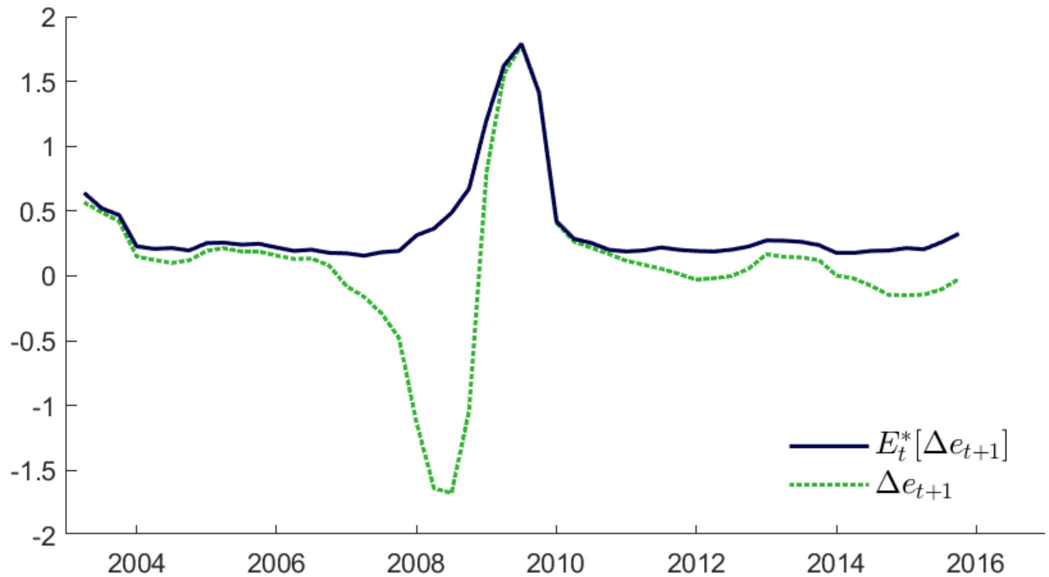

Figure 2: Expected and realized one-year earnings growth

The figure compares the one-year subjective earnings growth expectation and the observed future one-year earnings growth for the S&P 500. The solid line is the one-year subjective earnings growth expectation based on survey data. The dotted line is the observed future one-year earnings growth.

growth and future earnings growth, respectively.

Using the I/B/E/S earning expectations for individual firms, we calculate an aggregate measure of subjective earnings growth expectations in an analogous procedure to the construction of the subjective dividend growth expectations. Figure 2 shows the subjective earnings growth expectations and the observed future one-year earnings growth. While earnings growth expectations fail to predict the -1.67 log points (-81%) change in earnings during the crisis, they do predict the quick recovery and track future earnings growth reasonably well for the rest of the sample. The correlation of subjective earnings growth expectations with observed future earnings growth is 0.59 and is significant. Subjective earnings growth expectations are highly volatile with a standard deviation of 36%. Although the large recovery after the recession is the main episode of volatility, earnings growth expectations are still volatile outside the recession. Subjective earnings growth expectations have a standard deviation of 9.5% outside the recession, a similar value to the 7.6% standard deviation of dividend growth expectations.

4.2 Returns

The behavior of subjective return expectations is shown in Figure 3. Compared to subjective dividend growth expectations, subjective return expectations $E_t^*[r_{t+1}]$ look quite flat. Moreover, Table 2 shows that subjective return expectations and future returns are weakly negatively correlated. There is, however, a strong positive correlation (0.41) between r_t and $E_t^*[r_{t+1}]$, meaning that subjective return expectations are more related to

Figure 3: Expected and realized one-year returns

The figure compares the one-year subjective return expectation and the observed future one-year return on the S&P 500. The solid line is the one-year subjective return expectation based on survey data. The dotted line is observed future one-year return.

current returns than future returns. These results are consistent with Greenwood and Shleifer [2014], whose study rejects a positive correlation between subjective return expectations and future returns and highlights instead a strong role of recent returns in shaping subjective return expectations. Further, subjective return expectations are significantly positively correlated with pd_t, which reverses the negative relationship between r_{t+1} and pd_t.

We can conclude thus far that subjective expectations have strong predictions for future one-year dividend growth, but not for one-year returns. Further, we find that subjective dividend growth expectations are significantly more volatile than the subjective return expectations. Even though realized one-year dividend growth is 40% less volatile than the realized one-year returns, the volatility of subjective dividend growth expectations is 6 times larger than the volatility of subjective return expectations.

5 Decomposition of Price-Dividend Ratio Volatility

In this section, we decompose the variance of the price-dividend ratio into movements in subjective dividend growth and return expectations. First, we use one-year subjective expectations to estimate the relative importance of short-term subjective dividend growth and return expectations relative to the long-term component. Then, using longer horizon subjective expectations, we estimate the importance of the full horizon subjective dividend growth expectations relative to the full horizon subjective return expectations. The portion of price-dividend ratio movements that is due to changes in subjective dividend growth expectations is defined as cash flow news, and the portion that is due to changes in subjective return expectations is defined as discount rate news.

5.1 One-year Decomposition

From equation (1), we have that

$$
\begin{aligned}
Var\left(pd_t\right) &= Cov\left(E_t^*\left[\Delta d_{t+1}\right], pd_t\right) + Cov\left(-E_t^*\left[r_{t+1}\right], pd_t\right) + \rho Cov\left(E_t^*\left[pd_{t+1}\right], pd_t\right) \\
1 &= \underbrace{\frac{Cov\left(E_t^*\left[\Delta d_{t+1}\right], pd_t\right)}{Var\left(pd_t\right)}}_{CF_1} + \underbrace{\frac{Cov\left(-E_t^*\left[r_{t+1}\right], pd_t\right)}{Var\left(pd_t\right)}}_{DR_1} + \underbrace{\rho\frac{Cov\left(E_t^*\left[pd_{t+1}\right], pd_t\right)}{Var\left(pd_t\right)}}_{LT}. \qquad (4)
\end{aligned}
$$

Our measures CF_1 and DR_1 capture the influence of one-year subjective dividend growth expectations (cash flow news) and one-year subjective return expectations (discount rate news). The influence of subjective dividend growth and return expectations looking more than one year ahead is all captured in our measure of long-term influence LT. We can directly measure short-term subjective dividend growth and return expectations, while the one-year subjective price-dividend ratio expectation ($E_t^*\left[pd_{t+1}\right]$) is inferred from the current price-dividend ratio, one-year subjective return expectations, and one-year subjective dividend growth expectations.

A useful feature of this decomposition is that the one-year cash flow news and one-year discount rate news are estimated completely separately. There is no concern that subjective return expectations are affecting the estimate of cash flow news or that the subjective dividend growth expectations are affecting the estimate of discount rate news. This separation of the two types of subjective expectations means that these estimates will still be accurate even if the investors answering the return surveys and the investors answering the dividend surveys disagree on their beliefs. If the two groups of investors have different subjective expectations, then LT can simply be interpreted as the portion of the price-dividend ratio variation that is not explained by movements in the first group's subjective return expectations or the second group's subjective dividend growth expectations.

To provide a benchmark for our estimates, we also calculate the decomposition using statistical expectations. Let $y_t = \left[\Delta d_t \, r_t\right]$. As long as pd_t is used in the statistical expectations, the relation $E_t\left[y_{t+1}|pd_t\right] = y_{t+1} + \varepsilon_{t+1}$ will hold with $Cov\left(\varepsilon_{t+1}, pd_t\right) = 0$. Because the forecast error will be uncorrelated with the observable pd_t, the statistical expectations will satisfy $Cov\left(E_t\left[y_{t+1}\right], pd_t\right) = Cov\left(y_{t+1}, pd_t\right)$. Using this fact, we can use observed values of y_{t+1} to calculate the decomposition of the price-dividend ratio under statistical expectations. Table 3 shows the results of the decomposition when we use subjective expectations and when we use statistical expectations. As a robustness check, we run the decompositions for the full sample

2003Q1-2015Q3 and with the NBER recession (2007Q4-2009Q2) removed.

There are two relevant properties of the return expectations. First, we see that under subjective expectations DR_1 is negative in both the full sample and when the recession is removed. In other words, investors tend to report higher expected one-year returns when the price-dividend ratio is high. For comparison, under statistical expectations DR_1 is positive because a high price-dividend ratio predicts low one-year returns. The second interesting outcome is that under subjective expectations DR_1 is small in magnitude. Contrary to the standard view that stock market volatility is driven mainly by changes in expected returns, we find that subjective return expectations play a negligible role in moving the price-dividend ratio. This stems from the fact that subjective return expectations are flat in comparison to both one-year returns and subjective dividend growth expectations. Even if the sign of DR_1 was reversed, meaning that subjective return expectations helped to explain the volatility of the price-dividend ratio, it would not make a substantial contribution because there is not enough volatility in subjective return expectations to explain large price-dividend ratio movements. This positive subjective relationship extends beyond the one-year horizon. The CFO survey, as well as the Survey of Professional Forecasters, gives expected average returns over the next ten years. These subjective expectations are positively correlated with the current price-dividend ratio.[5] Therefore, the contribution of subjective return expectations must be negative for at least the first 10 years.

In comparison, the effect of one-year subjective dividend growth expectations is large and positive. In market booms, investors tend to report higher subjective dividend growth expectations. Because subjective dividend growth expectations vary significantly over time, they account for a large portion of the volatility of the price-dividend ratio. In Appendix A.3, we show that this also holds for two-year subjective dividend growth expectations. One-year dividend growth can have a large effect on prices because it affects the levels of both short-term and long-term dividends. Holding dividend growth fixed for all following years, a 10 percentage point drop in one-year dividend growth means that all future dividends fall by 10%, which causes the price to fall by 10%.

Under statistical expectations, the full sample decomposition also puts a large weight on future one-year dividend growth, however, this is only due to comovement of the price-dividend ratio and future one-year dividend growth during the recession. Between 2008-2010 there is a large drop in the current price-dividend ratio and observed future one-year dividend growth, but the variables have little comovement beyond that. When the NBER recession is removed, the coefficient on statistical dividend growth expectations falls to 0.13 and is insignificant, which matches the typical result found in the literature.

Removing the recession has virtually no effect on the importance of subjective dividend growth or return expectations. While dividend growth expectations are on average quite accurate, the forecast errors tend to

[5] See A.3-A.4 for more details.

Table 3: Variance decomposition of price-dividend ratio using one-year estimates

		CF_1	DR_1	LT
Subjective Expectations	Full Sample	0.36	−0.05	0.69
		(0.04)	(0.01)	(0.04)
	Without recession	0.35	−0.06	0.71
		(0.07)	(0.02)	(0.08)
Statistical Expectations	Full Sample	0.38	0.19	0.43
		(0.12)	(0.12)	(0.20)
	Without recession	0.13	0.29	0.57
		(0.11)	(0.15)	(0.20)

Note. This table shows the importance of one-year cash flow news (CF_1) and one-year discount rate news (DR_1) in the price-dividend ratio variance decomposition. Under *Subjective Expectations*, survey data from I/B/E/S and Graham-Harvey is used and CF_1 and DR_1 are the coefficients obtained by regressing $E_t^*[\Delta d_{t+1}]$ on pd_t and $E_t^*[r_{t+1}]$ on pd_t respectively. Under *Statistical Expectations*, observed dividend and return data from Compustat are used and CF_1 and DR_1 are the coefficients obtained by regressing Δd_{t+1} on pd_t and r_{t+1} on pd_t respectively. LT is inferred from the decomposition identity $1 = CF_1 + DR_1 + LT$. For each specification, we use quarterly data running from 2003Q1 to 2015Q3 and then remove the NBER recession period spanning 2007Q4 to 2009Q2. Small-sample adjusted Newey-West standard errors in parenthesis.

be correlated with the price-dividend ratio. Excluding the recession, investors tend to overestimate dividend growth during market booms and underestimate during market busts. This is why the one-year cash flow news component under subjective expectations remains high when the recession is removed, even though cash flow news under statistical expectations falls significantly. Given that subjective return expectations are relatively flat throughout the entire sample and are positively correlated with the price-dividend ratio, it is not surprising that their contribution to price-dividend ratio volatility remains small and negative when the recession is removed.

5.2 Full Horizon Decomposition

This section estimates the full horizon cash flow news $CF = \frac{Cov\left(\sum_{j=1}^{\infty}\rho^{j-1}E_t^*[\Delta d_{t+j}], pd_t\right)}{Var(pd_t)}$ and discount rate news $DR = \frac{Cov\left(-\sum_{j=1}^{\infty}\rho^{j-1}E_t^*[r_{t+j}], pd_t\right)}{Var(pd_t)}$ that comprise equation (3). Using the two-year subjective dividend growth expectations and the average ten-year subjective return expectations, we estimate a simple decay model of investor expectations given by

$$\begin{aligned} E_t^*[\Delta d_{t+1+j}] - \mu_d &= \phi_d^j \left(E_t^*\left[\Delta d_{t+1}\right] - \mu_d\right) + \varepsilon_{t,j}^d \\ E_t^*\left[r_{t+1+j}\right] - \mu_r &= \phi_r^j \left(E_t^*\left[r_{t+1+j}\right] - \mu_r\right) + \varepsilon_{t,j}^r. \end{aligned}$$

This functional form is consistent with an agent who believes in an underlying AR1 model. After the agent forms her one-year expectations, she simply believes that dividend growth and returns will gradually decay back to their mean values μ_d, μ_r. We refer to ϕ_d, ϕ_r as the expected persistence of dividend growth and returns, meaning how persistent the investors believe shocks to dividend growth and returns will be. We choose this form because of its simplicity and the fact that it holds for most standard asset pricing models, due to the fact that stock fundamentals are typically written as AR1 processes. We estimate the expected persistence of dividend growth using the two-year subjective dividend growth expectations obtained from I/B/E/S. For returns, we use the subjective expected returns for the next 10 years, $E_t^*\left[r_{t+1,t+10}\right]$, and the one-year subjective return expectations to calculate subjective return expectations for years 2 through 10, $E_t^*\left[r_{t+2,t+10}\right]$. We then use this value to estimate the expected persistence of returns.

With this simple specification, we have a straightforward definition for full horizon cash flow news and discount rate news, $CF = \frac{1}{1-\rho\phi_d}CF_1$ and $DR = \frac{1}{1-\rho\phi_r}DR_1$. Even if subjective expectations do not follow a simple decay process, this definition of cash flow news and discount rate news will still be correct as long as $\sum_{j=1}^{\infty}\rho^{j-1}\varepsilon_{t,j}^d$ and $\sum_{j=1}^{\infty}\rho^{j-1}\varepsilon_{t,j}^r$ are not correlated with the current price-dividend ratio. Using the two-year subjective dividend growth expectations and ten-year subjective return expectations, we do not find any evidence that the error terms are correlated with pd_t. In Section 6, we explicitly test the accuracy of the simple decay functional form with errors that are uncorrelated with pd_t and find that it is a good fit for matching the movements in the price-dividend ratio. Combining our definitions for CF and DR with (3) gives three equations that determine ϕ_d, ϕ_r:

$$\begin{aligned} E_t^*\left[\Delta d_{t+2}\right]-\mu_d &= \phi_d\left(E_t^*\left[\Delta d_{t+1}\right]-\mu_d\right)+\nu_t^d && (5)\\ E_t^*\left[r_{t+2,t+10}^*\right]-9\mu_r &= \phi_r\frac{1-\phi_r^9}{1-\phi_r}\left(E_t^*\left[r_{t+1}\right]-\mu_r\right)+\nu_t^r && (6)\\ 1 &= \frac{1}{1-\rho\phi_d}CF_1+\frac{1}{1-\rho\phi_r}DR_1. && (7) \end{aligned}$$

The benefit of having both subjective dividend growth expectations and subjective return expectations is that we have independent methods for measuring the size of cash flow news and discount rate news. There are three possible ways to estimate the decomposition. First, we can estimate ϕ_d from the subjective dividend growth expectations, calculate CF and then infer $DR = 1 - CF$. Second, we can estimate ϕ_r from the subjective return expectations, calculate DR and infer CF. Third, we can jointly estimate ϕ_d, ϕ_r such that $CF + DR = 1$ using maximum likelihood. Table 4 shows the results of these three estimations and for comparison also shows two recent estimates of cash flow news and discount rate news under the statistical expectations used in the literature.

All three methods show that the price-dividend ratio is predominantly driven by subjective dividend growth expectations. In the first row of Table 4, where the effect of subjective dividend growth expectations is estimated directly, we clearly see that subjective dividend growth expectations explain the majority of price-dividend ratio movements, 73%. This estimation relies only on subjective dividend growth expectations data and is completely separate from the subjective return expectations data. The second row shows the results when the contribution of subjective return expectations is estimated directly. The contribution of full horizon subjective return expectations is small and negative at -9%. The negative contribution of subjective return expectations means subjective dividend growth expectations must explain over 100% of the price-dividend ratio volatility as it must drive the price-dividend ratio movements and make up for the positive comovement of subjective return expectations and pd_t. Since the surveys are taken from different groups of investors, it is not surprising that the inferred CF differs from the directly observed CF. What is surprising is that both surveys provide strong evidence that prices are predominantly driven by investors' subjective dividend growth expectations.

When the expected persistences of dividend growth and returns are estimated jointly, the result is similar to the case where discount rate news is estimated directly. Because subjective return expectations have low volatility and $DR_1 < 0$, it is difficult to get a large, positive DR. In fact, changing ϕ_r has little impact on DR which is why the standard error for ϕ_r is so large in the third estimation. Thus, it is far more likely that

Table 4: Variance decomposition of price-dividend ratio into full horizon CF and DR

Subjective Expectations Data	ϕ_d	ϕ_r	CF	DR
Dividend growth	0.52		0.73	0.27
	(0.16)		(0.24)	(0.24)
Returns		0.48	1.09	−0.09
		(0.39)	(0.07)	(0.07)
Dividend growth and returns	0.69	0.46	1.09	−0.09
	(0.10)	(2.18)	(0.35)	(0.35)
Cochrane [2010]			−0.35	1.35
Maio and Santa-Clara [2015]			−0.05	1.05

Note. This table calculates the full horizon variance decomposition using different subsets of data sources. In the first row, we use exclusively the subjective dividend expectations from I/B/E/S ($E_t^*[\Delta d_{t+1}]$ and $E_t^*[\Delta d_{t+2}]$) and estimate the expected persistence ϕ_d. Then CF is estimated as $CF_1/(1-\rho\phi_d)$ and DR is inferred as $1-CF$. In the second row, we use exclusively the subjective return expectations from Graham-Harvey ($E_t^*[r_{t+1}]$ and $E_t^*[r_{t+1,t+10}]$) and estimate the expected persistence ϕ_r. Then, DR is estimated as $DR_1/(1-\rho\phi_r)$ and CF is inferred as $1-DR$. The third row uses both sources of data to perform a ML estimation constrained by the identity $1=CF+DR$ and jointly estimate the expected persistences ϕ_r, ϕ_d which determine CF and DR. CF_1 and DR_1 are estimated by regressing $E_t^*[\Delta d_{t+1}], E_t^*[r_{t+1}]$ on pd_t. We use quarterly data running from 2003Q1 to 2015Q3. Small-sample adjusted Newey-West standard errors in parenthesis. Fourth and fifth row show Cochrane [2010] and Maio and Santa-Clara [2015] estimates of CF and DR under statistical expectations.

the direct estimation of CF understates the role of subjective dividend growth expectations than that the indirect estimation of CF from the subjective returns data overstates the role of subjective dividend growth expectations.

In Appendix A.3, we show that similar results can be obtained by simply adding the two-year subjective dividend growth expectations ($E^* [\Delta d_{t+2}]$) and ten-year subjective return expectations ($E^* [r_{t+1,t+10}]$) to the decomposition and calculating the two-year cash flow news and ten-year discount rate news without the need to estimate a decay structure for expectations. Ten-year subjective return expectations slightly rise with the price-dividend ratio meaning that ten-year discount rate news is small and negative, consistent with the one-year discount rate news of Section 5.1 and the full horizon discount rate news of Table 4. Two-year subjective dividend growth expectations rise with the price-dividend ratio, creating large positive cash flow news. Combined, changes in one-year and two-year subjective dividend growth expectations explain 59% of the movements in the price-dividend ratio. This fits well with the full horizon cash flows news of Table 4 which estimates that changes in subjective dividend growth expectations at all horizons explain at least 73% of the variation in the price-dividend ratio.

6 Asset Pricing Using Dividend Surveys

Summarizing the results of the previous sections, we find three key facts for subjective expectations. First, subjective return expectations are relatively flat and do not play a large role in explaining price movements. Second, subjective dividend growth expectations are time-varying and explain the majority of price movements. Third, subjective dividend growth expectations have low expected persistence, which means that prices are primarily driven by changes in short-term dividend growth expectations.

Based on these results, we build a simple asset pricing model in which discount rates are constant and dividend growth expectations are taken from the one-year survey forecasts and the simple decay functional form. We then compare the model implied asset prices to observed prices. First, we check if the observed price for a short-term asset is consistent with volatile short-term dividend expectations and a constant discount rate. Second, we test whether the simple decay functional form and low expected persistence generate enough volatility to match price movements of a long-term asset. Finally, we compare the forecasted price changes implied by the model to the observed future price changes to test if this model has the potential to predict future prices.

The convenient feature of this model is that the agent's entire horizon of dividend growth expectations and discount rates is captured by one time-series, $E_t^* [\Delta d_{t+1}]$. We make the agent's discount rates constant and equal to μ_r, which is the mean one-year subjective return expectation from the survey data. We set the agent's one-year dividend growth expectations equal to the subjective expectations $E_t^* [\Delta d_{t+1}]$ and set her $j+1$ year dividend growth expectations equal to $\mu_d + \phi_d^j (E_t^* [\Delta d_{t+1}] - \mu_d)$, based on the simple decay functional form of Section 5.2. The mean expected dividend growth, μ_d, is set to the mean one-year subjective dividend growth expectation from the survey data and the expected persistence ϕ_d is estimated from (5) using the one-year and two-year subjective dividend growth expectations. Using just this one time-series $E_t^* [\Delta d_{t+1}]$, we will be able to match the prices of a short-term asset and a long-term asset and reasonably predict the future change in the price of the long-term asset. The ability of $E_t^* [\Delta d_{t+1}]$ to explain asset prices without needing time-varying discount rates may cause one to wonder if the subjective divided growth expectations are risk-adjusted. We show in Appendix A.5 that there is no evidence that $E_t^* [\Delta d_{t+1}]$ is risk-adjusted.

For the first test, we look at the observed price of a one-year dividend strip, i.e. a claim to all dividends from the S&P 500 for one year. The log price for a one-year dividend strip is simply the difference between the expected log dividend and the log discount rate. The model log price for a one-year dividend strip is then

$$p_t^{strip,model} = E_t^* [d_{t+1}] - \mu_r.$$

Volatile one-year dividend expectations and a constant discount rate imply that the one-year dividend strip price should be quite volatile. The model one-year strip price rises 0.4 log points (50%) from 2004 to 2008 and then quickly drops 0.3 log points (35%) from 2008 to 2009. Van Binsbergen, Brandt, and Koijen [2012] show that this is exactly the case. Using options data, they measure the one-year dividend strip price and find excessive volatility. The first row of Table 5 shows the regression of the observed dividend strip price on the model strip price. The coefficient on the model price is significant and almost exactly 1, at 1.01, and we cannot reject the null hypothesis that the observed strip price equals the model strip price plus noise, i.e. $a = 0$ and $b = 1$. Since the observed strip price is only provided up to 2009Q3, our sample is 2003Q1-2009Q3. The left-hand side of Figure 4 shows the observed and model strip prices. The first noteworthy result is that the model strip price matches the overall level of the observed strip price. Given that the observed strip price, subjective dividend expectations, and μ_r are all taken from separate sources, it is remarkable that the mean strip price is 3.15 and the mean model strip price is 3.13. The second impressive result is that both measures have similar volatilities, with standard deviations of 0.17 and 0.15. While the model strip price does not fully match the drop in the observed strip price during the financial crisis, movements in the model price explain the majority of the movements in the observed price. As Table 5 shows, the model strip price can explain 75% of the variation in the observed strip price. Thus, the observed strip price matches the prediction from volatile one-year dividend expectations and a constant discount rate.

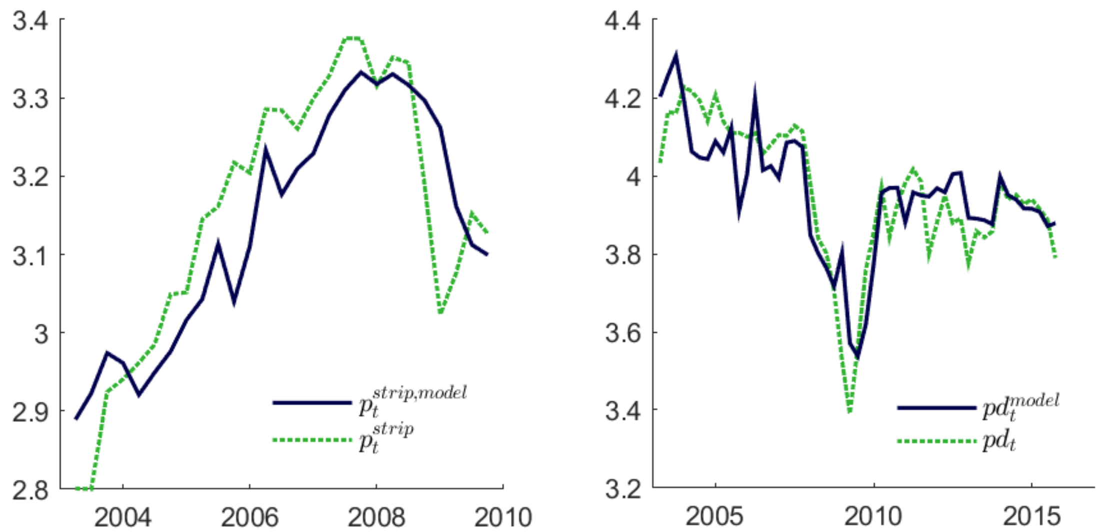

Figure 4: Price of dividend strip and price-dividend ratio

The left figure compares the S&P 500 one-year dividend strip price implied by the model and the S&P 500 observed one-year dividend strip price. The solid line is the one-year dividend strip price implied by the model. The dotted line is the strip price calculated by Van Binsbergen, Brandt, and Koijen [2012] based on options data. The right figure compares the S&P 500 model price-dividend ratio and the S&P 500 observed price-dividend ratio. The solid line is the price-dividend ratio implied by the model. The dotted line is the observed price-dividend ratio for the S&P 500. Since the strip price is not available for our entire sample, the left figure comprises the period 2003Q1-2009Q3. The right figure covers 2003Q1-2015Q3.

The second test checks whether volatile short-term dividend growth expectations with low expected persistence can match the movements in the price-dividend ratio of the S&P 500, i.e. can the model match the price of a long-term asset. Longer horizon dividend growth expectations in the model are calculated based on the one-year dividend growth expectations and the simple decay functional form. The model price-dividend ratio is

$$\begin{aligned} pd_t^{model} &= \bar{pd} + \sum_{j=1}^{\infty} \rho^{j-1} \phi_d^{j-1} \left(E_t^* \left[\Delta d_{t+1} \right] - \mu_d \right) \\ &= \bar{pd} + \frac{1}{1-\rho\phi_d} \left(E_t^* \left[\Delta d_{t+1} \right] - \mu_d \right). \end{aligned}$$

Because discount rates are constant, the price-dividend ratio only deviates from its mean when dividend growth expectations are above or below their mean.

Table 5 shows the results from regressing the observed price-dividend ratio on the model value. Remarkably, movements in the model price-dividend ratio match almost one to one with movements in the observed price-dividend ratio, with a highly significant regression coefficient of 0.96. Further, we cannot reject the null hypothesis that the observed price-dividend ratio exactly equals the model price-dividend ratio plus

Table 5: Model and observed asset prices

Regression	a	b	R^2
$p_t^{strip} = a + b\left(p_t^{strip,model}\right) + \varepsilon_t$	-0.01	1.01	0.75
	(0.51)	(0.16)	
$pd_t = a + b\left(pd_t^{model}\right) + \varepsilon_t$	0.16	0.96	0.70
	(0.43)	(0.11)	
$pd_{t+1} - pd_t = a + b\left(pd_{t+1}^{forecast} - pd_t\right) + \varepsilon_t$	-0.02	0.91	0.36
	(0.04)	(0.26)	

Note. This table reports the results of OLS regressions of observed values on the values implied by the model. The one-year dividend strip price series p_t^{strip} is taken from Van Binsbergen, Brandt, and Koijen [2012] and the regression is run over the sample 2003Q1-2009Q3 because the strip price is not available for our entire sample. The regressions for the observed price-dividend ratio pd_t and the future one-year observed change in the price-dividend ratio $pd_{t+1} - pd_t$ cover 2003Q1-2015Q3. Small-sample adjusted Newey-West standard errors are reported in parenthesis.

noise. The right-hand side of Figure 4 shows the observed and model price-dividend ratio over our sample. While there may be discrepancies in certain quarters, the R^2 is high, at 0.70, meaning that the model value explains the majority of observed price-dividend ratio movements. Overall, the model value matches the general trends well with a high but declining price-dividend ratio before the crisis, a sharp drop and recovery during 2008-2010, and a gradual decline from 2010-2016. Therefore, the observed price-dividend ratio is consistent with a constant discount rate, volatile short-term dividend growth expectations, and low expected persistence.

The third test concerns the accuracy of this model for predicting future movements in the price-dividend ratio. To test this, we can look at two possible series. The first is the price of a futures contract for the S&P 500. Given that we match the observed one-year dividend strip price and the price-dividend ratio quite well, we know that the price of a S&P 500 futures contract implied by the model must match well with the observed futures contract price. This comes from the fact that the one-year dividend strip price can be derived from the futures contract price.

The second series we can use is the observed future change in the price-dividend ratio. If there is no new information over the next year, then the future price-dividend ratio should be the mean price-dividend ratio plus the expected future dividend growth,

$$\begin{aligned} pd_{t+1}^{forecast} &= \bar{pd} + \sum_{j=1}^{\infty} \rho^{j-1} \phi^j \left(E_t^*\left[\Delta d_{t+1}\right] - \mu_d\right) \\ &= \bar{pd} + \frac{\phi_d}{1-\rho\phi_d}\left(E_t^*\left[\Delta d_{t+1}\right] - \mu_d\right). \end{aligned}$$

Figure 5: Change in price-dividend ratio

The figure compares the one-year change in the price-dividend ratio forecasted by the model and the observed future one-year change in the price-dividend ratio for the S&P 500. The solid line is the one-year change in the price-dividend ratio forecasted by the model if no new information is revealed over the next year. The dotted line is the observed one-year future change in the price-dividend ratio.

This is the same as the model value for the current price-dividend ratio, except that a coefficient of ϕ_d is applied to the expected dividend growth, since only ϕ_d of the shock to dividend growth is expected to persist into the next year. This will be the model prediction at time t of the price-dividend ratio at $t+1$. Since $\phi_d = 0.52$, the price-dividend ratio is forecasted to be half-way back its mean value after one year.

Table 5 shows the results of regressing the observed future change in the price-dividend ratio on the model forecasted change in the price-dividend ratio. Since this test deals with future prices, it is not surprising that the coefficient on the forecasted value is slightly lower than the coefficients for the tests of the dividend strip price or the price-dividend ratio. With a coefficient of 0.91, we find that the change in the price-dividend ratio forecasted by the model matches well with the observed change. Figure 5 shows that the forecasted and observed future change in the price-dividend ratio. It is true that the forecasted value does not foresee the crisis. This makes sense since the forecasted value is a measure of what should happen if there is no new information over the next year and we know from the survey data that there was a large amount of negative information about future dividends revealed during the crisis. The forecasted change does predict the quick recovery and generally fits well with the observed change in the price-dividend ratio both before and after the crisis. Overall, we find that the forecasted change explains 36% of the variation in the observed future price-dividend ratio change, which is large given that future price changes are notoriously hard to predict.

Table 6: Comovement of return surveys and the price-dividend ratio

	G-H 1 Year	Livingston 1 Year	Michigan 2-3 Year	G-H 10 Year	SPF 10 Year
Sample	2003-2015	1952-2016	2000-2005	2003-2015	1992-2016
Std. dev.	0.013	0.052	0.012	0.008	0.012
$\frac{-Cov(\cdot,pd_t)}{Var(pd_t)}$	-0.049 (0.013)	-0.033 (0.024)	-0.074 (0.015)	-0.007 (0.012)	-0.003 (0.013)

Note. This table shows the volatility of survey return forecasts and their comovement with the price-dividend ratio. The G-H one-year and ten-year return surveys cover 2003Q1-2015Q3 quarterly. The Livingston survey covers 1952Q2-2016Q2 and is conducted every 6 months. The Michigan survey covers June 2000 - October 2005 and is conducted in 22 months over this period. The Survey of Professional Forecasters covers 1992-2016 annually. Small-sample adjusted Newey-West standard errors in parenthesis.

7 Robustness Checks

7.1 Longer Sample Surveys

The main results of this paper focus on the sample 2003Q1-2015Q3. In this section, we will use additional return surveys to show that subjective dividend growth expectations must be the main driver of the price-dividend ratio over longer samples, even going back to 1952. Specifically, we will show that subjective return expectations are relatively flat and have a small, positive comovement with the price-dividend ratio, consistent with our results for 2003Q1-2015Q3. Because their subjective return expectations cannot explain the movements in the price-dividend ratio, these respondents must have time-varying subjective dividend growth expectations that drive the price-dividend ratio. In other words, we find small, negative subjective discount rate news over these longer samples, implying that there must be large, positive subjective cash flow news.

In addition to the one-year and ten-year return expectations obtained from the Graham-Harvey survey of Duke University, we use the Federal Reserve Bank of Philadelphia's Livingston Survey, the University of Michigan Survey of U.S. consumers and the Survey of Professional Forecasters as additional measures of S&P 500 return expectations over different samples. The Livingston Survey covers 1952-2016 and measures one-year return expectations. The Michigan Survey covers 2000-2005 and measures return expectations over the next 2-3 years. The Survey of Professional Forecasters covers 1992-2016 and measures ten-year return expectations. All expected returns are annualized. Appendix A.4 gives more detail on these additional surveys.

Table 6 shows that subjective return expectations do not explain price-dividend ratio movements in any of the different sample periods. We consistently find that subjective return expectations are relatively flat and slightly rise with the price-dividend ratio. In the second row, we see that all return surveys, except the Livingston survey, have similar standard deviations. These low standard deviations make it unlikely that subjective return expectations are driving the large movements in the price-dividend ratio. In the third row, we see the contribution of each survey to explaining the volatility of the price-dividend ratio. If subjective return expectations fall when the price-dividend ratio is high, then the subjective return expectations would be able to explain some of the price-dividend ratio movements. Instead, we find that all 5 subjective return expectations are high when the price-dividend ratio is high. Because subjective return expectations rise with the price-dividend ratio and have low volatility, they make a small negative contribution to explaining the variance of the price-dividend ratio. Even the greater volatility of the Livingston survey does not translate into a larger covariance with the price-dividend ratio. Overall, we find that subjective return expectations cannot explain the movements in the price-dividend ratio over these samples, which means that subjective dividend growth expectations must be driving price movements.

7.2 Fitting Expectations to Survey Responses

For robustness, we construct an alternative model where the survey responses are not treated as a direct measurement of investors' expectations, but instead are treated as data correlated with investors' expectations. We do this to account for the possibility that analysts do not report their precise expectations, but instead report some biased measure of their market sentiment, optimism, etc. We calculate one-year fitted expectations by regressing future one-year observed dividend growth and future one-year observed returns on the survey responses, providing a greater chance for large time-varying return expectations. We then calculate long-term expectations using the autocorrelation of the survey data.

Since the contributions of dividend growth and return expectations only depend on covariances, we can demean all variables without loss of generality. Let $X_t = (E_t^*[\Delta d_{t+1}] \quad E_t^*[r_{t+1}])$ be the demeaned one-year survey responses and let Δd_{t+1}, r_{t+1} be the demeaned future one-year observed dividend growth and future one-year observed returns. We estimate $\hat{\beta}^d, \hat{\beta}^r, \hat{\varphi}$ from three equations:

$$\begin{aligned}
\Delta d_{t+1} &= \beta^d X_t + \varepsilon_{t+1}^d \\
r_{t+1} &= \beta^r X_t + \varepsilon_{t+1}^r \\
X_{t+1} &= \varphi X_t + \varepsilon_{t+1}.
\end{aligned}$$

The one-year fitted expectations from this specification are $\hat{E}_t\left[\Delta d_{t+1}\right] = \hat{\beta}^d X_t$ and $\hat{E}_t\left[r_{t+1}\right] = \hat{\beta}^r X_t$, and

Table 7: Cash Flow News (CF) and Discount Rate News (DR) under fitted expectations

$\widehat{CF}$	$\widehat{DR}$	$\hat{\beta}^d$		$\hat{\beta}^r$	
0.61	−0.08	[1.00	−0.54]	[0.24	−1.08]
(0.23)	(0.27)	(0.27	0.89)	(0.37	1.80)

Note. This table shows the estimated variance decomposition under fitted expectations. Parameter vectors $\hat{\beta}^d$ and $\hat{\beta}^r$ are obtained by linearly predicting Δd_{t+1} and r_{t+1} from survey responses $X_t \equiv (E_t^*[\Delta d_{t+1}] \quad E_t^*[r_{t+1}])$. This way $\hat{E}_t[\Delta d_{t+1}] = \hat{\beta}^d X_t$ and $\hat{E}_t[r_{t+1}] = \hat{\beta}^r X_t$. To obtain $\widehat{CF}$ and $\widehat{DR}$, we first estimate an autoregressive structure for $X_{t+1} = \varphi X_t + \varepsilon_{t+1}$. This way, we can define $\hat{E}_t[\Delta d_{t+j}] = \hat{\beta}^d \, (\hat{\varphi})^{j-1} \, X_t$ for all j and $\widehat{CF}$ is calculated by regressing $\hat{\beta}^d \, (I - \rho\hat{\varphi})^{-1} \, X_t$ on pd_t. Similarly, $\widehat{DR}$ is calculated by regressing $\hat{\beta}^r \, (I - \rho\hat{\varphi})^{-1} \, X_t$ on pd_t. We use quarterly data running from 2003Q1 to 2015Q3. Small-sample adjusted Newey-West standard errors in parenthesis

the corresponding full horizon fitted expectations are

$$\begin{aligned} \sum_{j=1}^{\infty} \rho^{j-1} \hat{E}_t[\Delta d_{t+j}] &= \hat{\beta}^d \, (I - \rho\hat{\varphi})^{-1} \, X_t \\ \sum_{j=1}^{\infty} \rho^{j-1} \hat{E}_t[r_{t+j}] &= \hat{\beta}^r \, (I - \rho\hat{\varphi})^{-1} \, X_t. \end{aligned}$$

This allow us to calculate the corresponding cash flow news and discount rate news for these fitted expectations as

$$\begin{aligned} \widehat{CF} &= \frac{Cov\left(pd_t, \sum_{j=1}^{\infty} \rho^{j-1} \hat{E}_t[\Delta d_{t+j}]\right)}{Var(pd_t)} \\ \widehat{DR} &= \frac{Cov\left(pd_t, -\sum_{j=1}^{\infty} \rho^{j-1} \hat{E}_t[\Delta d_{t+j}]\right)}{Var(pd_t)}. \end{aligned}$$

Note that under this specification, the contribution of the two components may not sum to one since the fitted expectations are not capturing all of the movements in future dividend growth and returns.

Looking at $\hat{\beta}^d$ in Table 7, we see that future dividend growth moves one-for-one with dividend growth survey responses and is not significantly related to return survey responses. The estimate for $\hat{\beta}^r$ shows that dividend growth surveys do not contain any significant information about future returns, unlike what one would expect if dividend growth survey responses were risk neutral. In Appendix A.5, we look further into the possibility of risk neutral expectations and find no evidence that dividend growth survey responses are risk-adjusted.

The cash flow news and discount rate news estimated in Table 7 show that fitted dividend growth expectations are still the dominant factor in the volatility of the price-dividend ratio. Overall, this alternative model matches our earlier results quite well. Fitted dividend growth expectations are volatile and rise with the price-dividend ratio, creating large, positive cash flow news. Fitted return expectations are relatively flat and rise with the price-dividend ratio due to their positive weight on dividend growth surveys, creating small, negative discount rate news.

7.3 Price Bubbles

The decomposition (3) assumes that the no-bubble condition $E_t^*\left[\lim_{T\to\infty}\rho^T pd_{t+T}\right] = 0$ is satisfied. In this section, we remove this assumption and allow for the possibility that investors believe the price-dividend ratio will be non-stationary and will grow faster than $1/\rho$. In this case, a third time-varying element in the decomposition could be affecting price movements. An increase in the price-dividend ratio could be driven by higher expected dividend growth, lower expected returns or a higher value of the "bubble" term. The new decomposition would be the following:

$$1 = \underbrace{\frac{Cov\left(\sum_{j=1}^{\infty}\rho^{j-1}E_t^*\left[\Delta d_{t+j}\right], pd_t\right)}{Var(pd_t)}}_{CF} + \underbrace{\frac{Cov\left(-\sum_{j=1}^{\infty}\rho^{j-1}E_t^*\left[r_{t+j}\right], pd_t\right)}{Var(pd_t)}}_{DR} + \underbrace{\frac{Cov\left(E_t^*\left[\lim_{T\to\infty}\rho^T pd_{t+T}\right], pd_t\right)}{Var(pd_t)}}_{Bubble}. \tag{8}$$

To estimate the three terms in the decomposition, we can use the value of CF derived from the dividend survey data and the value of DR derived from the return survey data in Section 5.2. We estimated $CF = 0.73$ and $DR = -0.09$, which means that under this specification, 36% of the volatility of the price-dividend ratio could be attributed to movements in the bubble term. This is a non-trivial contribution, but it does not change our main result that cash flow news is the dominant driver of price movements, explaining 73% of the volatility of the price-dividend ratio. The dominance of cash flow news holds even if we do not make any functional form assumptions about the behavior of long horizon subjective expectations. We show in Appendix A.3 that one-year and two-year subjective dividend growth expectations account for 59% of the volatility of the price-dividend ratio. This means that subjective dividend growth expectations can explain the majority of price movements, even when we remove the assumption of the no-bubble condition and remove our assumptions about long horizon expectations.

8 Conclusion

Movements in the price-dividend ratio must be explained by changes in expected dividend growth or changes in expected returns. Using subjective expectations based on survey data, we find that changes in subjective dividend growth expectations are the primary driver of the price-dividend ratio. Subjective dividend growth expectations vary significantly over time and are highly correlated with observed future dividend growth, while subjective return expectations are relatively flat and are weakly negatively correlated with observed future returns. Both subjective dividend growth and return expectations show low expected persistence, meaning that changes in short-term subjective expectations do not have large effects on long-term subjective expectations.

Short-term and long-term subjective dividend growth and return expectations all rise with the price-dividend ratio. This provides direct evidence of large cash flow news because a high price-dividend ratio can be explained by high subjective dividend growth expectations. It also provides indirect evidence of large cash flow news through the lack of discount rate news since a high price-dividend ratio cannot be explained by lower subjective expected returns. For both the direct and indirect estimates of cash flow news, we find that at least 73% of the variation in the price-dividend ratio is explained by changes in subjective dividend growth expectations. This is concentrated in short-term expectations, with changes in one-year subjective dividend growth expectations alone accounting for 36% of price-dividend ratio movements. These findings highlight the importance of time-varying subjective dividend growth expectations in determining aggregate stock prices.

References

Klaus Adam, Albert Marcet, and Johannes Beutel. Stock Price Booms and Expected Capital Gains. *Working Paper*, 2015. doi: 10.2139/ssrn.2634053.

Andrew Ang. Predicting dividends in log-linear present value models. *Pacific Basin Finance Journal*, 20(1): 151–171, 2012. ISSN 0927538X. doi: 10.1016/j.pacfin.2011.08.005.

Nicholas Barberis, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. X-CAPM: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1):1–24, 2015. URL `http://dx.doi.org/10.1016/j.jfineco.2014.08.007`.

John Y. Campbell and Robert J. Shiller. Stock Prices, Earnings, and Expected Dividends. *The Journal of Finance*, 43(3):661–676, 1988a. doi: 10.1111/j.1540-6261.1988.tb04598.x.

John Y. Campbell and Robert J. Shiller. The Dividend-Price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies*, 1(3):195–228, 1988b.

Long Chen and Xinlei Zhao. Return Decomposition. *The Review of Financial Studies*, 22(12):5213–5249, 2009. doi: 10.1093/rfs/hhp017.

Long Chen, Zhi Da, and Xinlei Zhao. What Drives Stock Price Movements? *The Review of Financial Studies*, 0(0):1–36, 2013. doi: 10.1093/rfs/hht005.

John H Cochrane. The dog that did not bark: A defense of return predictability. *Review of Financial Studies*, 21(4):1533–1575, 2008. ISSN 08939454. doi: 10.1093/rfs/hhm046.

John H Cochrane. Presidential Address : Discount Rates. 2010.

Eugene F. Fama and Kenneth R. French. Dividend yields and expected stock returns. *Journal of Financial Economics*, 22(1):3–25, 1988. ISSN 0304405X. doi: 10.1016/0304-405X(88)90020-7.

Robin Greenwood and Andrei Shleifer. Expectations of Returns and Expected Returns. *The Review of Financial Studies*, 0(0), 2014. doi: 10.1093/rfs/hht082.

Lawrence J Jin and Pengfei Sui. Asset Pricing with Return Extrapolation. *Working Paper*, 2017.

Ralph S J Koijen and Stijn Van Nieuwerburgh. Predictability of Returns and Cash Flows. *The Annual Review of Financial Economics*, 3:467–491, 2011. doi: 10.1146/annurev-financial-102710-144905.

Ralph S J Koijen, Maik Schmeling, and Evert B. Vrugt. Survey Expectations of Returns and Asset Pricing Puzzles. *Working Paper*, 2015. ISSN 1556-5068. doi: 10.2139/ssrn.2513701.

Paulo Maio and Pedro Santa-Clara. Dividend Yields, Dividend Growth, and Return Predictability in the Cross Section of Stocks. *Journal of Financial and Quantitative Analysis*, 50(1/2):33–60, 2015. doi: 10.1017/S0022 109015000058.

Monika Piazzesi, Juliana Salomao, and Martin Schneider. Trend and Cycle in Bond Premia. *Working Paper*, pages 1–27, 2015.

Robert J. Shiller. *Irrational exuberance*. Princeton University Press, 2015.

Jules Van Binsbergen, Michael Brandt, and Ralph Koijen. On the Timing and Pricing of Dividends. *American Economic Review*, 102(4):1596–1618, 2012.

A Appendix

A.1 Data Aggregation for Expected Dividend Calculation

Let the following variables be defined for each period t:

$D_{i,t}$ =ordinary dividend per share paid by company i at time t

$P_{i,t}$ =price per share of company i at time t

$S_{i,t}$ =shares of company i at time t

x_t =set of companies in S&P 500 at time t.

The total market value M_t and dividends D_t paid by the S&P 500 constituents are defined as:

$$\begin{aligned} M_t &= \sum_{i \in x_t} P_{i,t} S_{i,t} \\ D_t &= \sum_{i \in x_t} D_{i,t} S_{i.t}. \end{aligned}$$

Standard & Poor's define the S&P 500 index ($SP500_t$) as the total market capitalization of the constituents M_t, adjusted by a divisor. The divisor is defined by Standard & Poor's at every period to satisfy the following identity

$$\frac{\sum_{i \in x_t} P_{i,t} S_{i,t}}{Divisor_t} = SP500_t = \frac{\sum_{i \in x_{t-1}} P_{i,t} S_{i,t-1}}{Divisor_{t-1}} \tag{9}$$

or,

$$Divisor_t / Divisor_{t-1} = \sum_{i \in x_t} P_{i,t} S_{i,t} / \sum_{i \in x_{t-1}} P_{i,t} S_{i,t-1}. \tag{10}$$

In other words, the divisor moves so that the value of the S&P 500 index is not affected by changes in the S&P 500 constituents or the number of outstanding shares issued. One result of this is that the index is not affected by share repurchases. In addition, the divisor is also adjusted whenever a special dividend is issued. Standard and Poor's assume that the share price drops by the amount of the special dividend and adjust the divisor to offset this change in share price. Since the S&P 500 index is not affected by share repurchases or special dividends, we can think of the index as the value of a portfolio that automatically reinvests any special dividends or payments from share repurchases back into the portfolio. Therefore, the only cash flow from this portfolio is the ordinary dividends paid by S&P 500 constituents. This is why we do not include special dividends or share repurchases in our measure of dividends.

Figure 6: Comparison of aggregate dividend measures

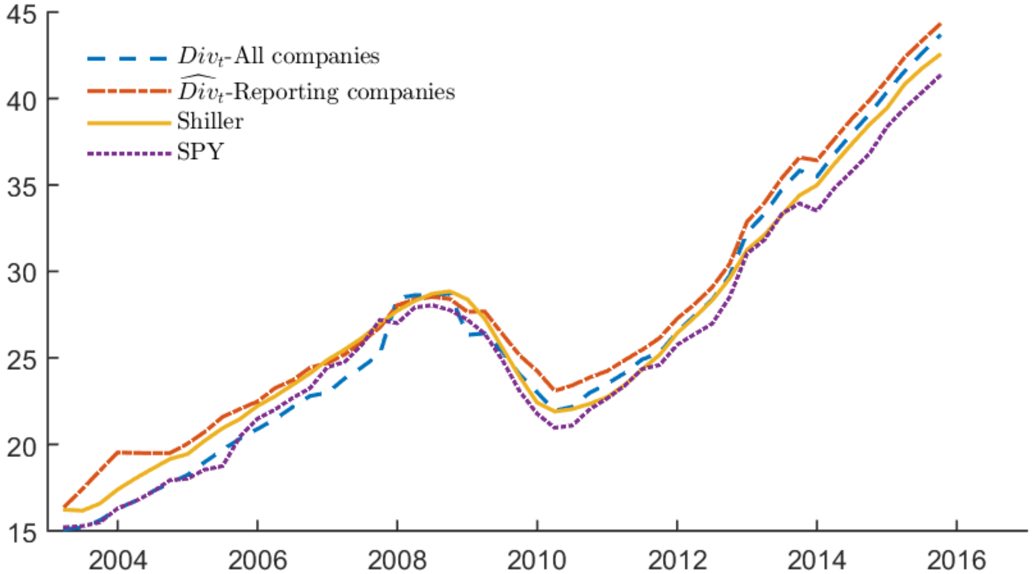

The figure compares four measures of the aggregate dividend for the S&P 500. The solid line *Shiller* contains the quarterly S&P 500 dividends obtained from Shiller [2015]. The dashed line *All Companies* contains the aggregate quarterly dividends paid out by all S&P 500 companies. The dash-dotted line *Forecasted* contains the aggregate quarterly dividends paid out only by S&P 500 companies for which a one-year subjective dividend expectation exists. This value is then scaled by the ratio of the market value of all S&P 500 companies to the market value of S&P 500 companies for which a one-year subjective dividend expectation exists. The dotted line *SPY* contains the quarterly dividends paid out by the SPDR S&P 500 ETF.

The divisor is not publicly available, but we can back out the value of the divisor by using (9) to obtain the simple ratio:

$$\widehat{Divisor}_t \quad = \quad M_t/SP500_t.$$

Once an estimate of $Divisor_t$ is obtained for every quarter, we construct an aggregate dividend index for the S&P 500, expressed as:

$$Div_t \quad = \quad \frac{D_t}{\widehat{Divisor}_t}.$$

This dividend measure has been constructed 'bottom-up' from the individual ordinary dividend payments of each company. We can see in Figure 6 the performance of our dividend measure compared to the S&P 500 dividend reported by Robert Shiller and the dividends paid by the Exchange Traded Fund "SPY", the largest replicating ETF of the S&P 500. The correlation of Div_t and the other dividend estimates in levels is also very high (> 0.99).

If, instead of an aggregate dividend measure, we want to build an aggregate expected dividend measure, we can use similar logic. The one-year subjective expected dividend for the S&P 500 can be described as:

$$E_t^*[Div_{t+1}] = E_t^* \left[\frac{\sum\limits_{i \in x_{t+1}} d_{i,t+1} s_{i,t+1}}{\widehat{Divisor}_{t+1}} \right]. \tag{11}$$

Because dividend forecasts are made in levels, rather than in logs, we approximate subjective expected dividend growth as $E_t^* [\Delta d_{t+1}] \approx \log(E_t^* [Div_{t+1}]) - \log(Div_t)$. As long as volatility is countercyclical, accounting for the Jensen terms from this approximation would only increase the procyclicality of $E_t^* [\Delta d_{t+1}]$ and strengthen our result that subjective cash flow news is large.

In order to build our aggregate estimator $E_t^*[Div_{t+1}]$, we need to make an assumption about how people form expectations about the future constituents and shares outstanding of the S&P 500. We assume that people expect that any changes in constituents or shares outstanding that may affect total dividends will be offset by changes in the divisor. Since the divisor adjusts to offset changes in total market value due to changes in constituents or shares outstanding, this simply means that people expect that changes in constituents or shares outstanding will have the same proportional effect on total dividends as total market value. In other words, we assume that people do not expect changes in constituents or shares outstanding to affect the price-dividend ratio of the S&P 500. A stronger assumption that would also be consistent with our methodology is to simply assume that people do not expect the constituents or shares outstanding to change over the next year. Assumption 1 implies that $E_t^* \left[\sum_{x_{t+1}} D_{i,t+1} S_{i,t+1} / Divisor_{t+1}\right] = E_t^* \left[\sum_{x_t} D_{i,t+1} S_{i,t} / Divisor_t\right] = \sum_{x_t} E_t^* \left[D_{i,t+1}\right] s_{i,t} / Divisor_t$.

Assumption 1. $E_t^* \left[\frac{\sum\limits_{x_{t+1}} D_{i,t+1} S_{i,t+1}}{\sum\limits_{x_t} D_{i,t+1} S_{i,t}} \right] = E_t^* \left[\frac{\sum\limits_{x_{t+1}} P_{i,t+1} S_{i,t+1}}{\sum\limits_{x_t} P_{i,t+1} S_{i,t}} \right].$

Given that we sometimes do not have expectations data for all firms in the S&P 500, we make a second assumption to construct $E_t^*[Div_{t+1}]$. Denote as $x_t^j \subset x_t$ the set of companies that have an expected value for horizon j. We normalize by the ratio of total market value, M_t, to the market value of the firms that have an expected dividend for horizon j, M_t^j. To do this, we assume that the firms that have an expected dividend are a representative sample of the S&P 500. Then $E_t^*[Div_{t+1}] = \left(M_t / M_t^1\right) \sum_{x_t^1} E_t^* \left[D_{i,t+1}\right] S_{i,t} / Divisor_t$.

Assumption 2. $\frac{\sum\limits_{i \in x_t} E_t^* \left[D_{i,t+1}\right] S_{i,t}}{\sum\limits_{i \in x_t^j} E_t^* \left[D_{i,t+1}\right] S_{i,t}} = \frac{M_t}{M_t^j}$ *where M_t^j is the market value of firms in x_t^j.*

Assumption 2 becomes easier to satisfy the higher the coverage of firms with valid forecasts we have. To see that the firms included in the forecast are representative of the S&P 500 index, a fourth measure is shown in Fig 9. The behavior of the measure $\widehat{Div}_t = \left(M_t / M_t^1\right) \sum_{i \in x_t^1} D_{i,t} S_{i,t} / Divisor_t$ is very similar to Div_t. This means that aggregate dividend constructions using only those companies with forecast for a certain horizon look very similar to the main aggregate dividend Div_t. Furthermore, the correlation of all the measures both in growth and levels is very close to one.

A.2 Horizon Interpolation

All firms on the I/B/E/S estimates database do not share the same fiscal year end date. In order to match a proper one-year forecast, different forecasts were used depending on the fiscal period of each firm's estimation.

For instance, if the estimates are taken on April 15, 2004 for a firm with fiscal year ending in January, the FY_1 variable will show a forecast for January 31, 2005, FY_2 will present a forecast for January 31, 2006, Q_1 for May 30th, 2004 and so on. Given that we want an estimate for March 31th, 2005, we will interpolate FY_1 and FY_2 accordingly to obtain that firm's 12 month estimate. Sometimes we make use of $Q_1 - Q_4$ estimates to improve the interpolation procedure. The exclusion of the quarterly estimates does not affect the results.

A.3 Campbell-Shiller Extended Decomposition

We perform a Campbell-Shiller decomposition similar to Section 5.1, but extend it with two extra pieces of data, the average ten-year expected return obtained in the Graham-Harvey CFO survey and the two-year expected dividend growth we obtained from I/B/E/S. We can decompose (3) into:

$$
\begin{aligned}
pd_t &= \sum_{j=1}^{2} \rho^{j-1} E_t^* \left[\Delta d_{t+1}\right] - \sum_{j=1}^{10} \rho^{j-1} E_t^* \left[r_{t+1}\right] + \left(\sum_{j=3}^{\infty} \rho^{j-1} E_t^* \left[\Delta d_{t+j}\right] - \sum_{j=11}^{\infty} \rho^{j-1} E_t^* \left[r_{t+j}\right] \right) \\
1 &= \frac{Cov\left(\sum_{j=1}^{2} \rho^{j-1} E_t^* \left[\Delta d_{t+j}\right], pd_t\right)}{Var\left(pd_t\right)} + \frac{Cov\left(-\sum_{j=1}^{10} \rho^{j-1} E_t^* \left[r_{t+j}\right], pd_t\right)}{Var\left(pd_t\right)} \\
&\quad + \frac{Cov\left(\sum_{j=3}^{\infty} \rho^{j-1} E_t^* \left[\Delta d_{t+j}\right] - \sum_{j=11}^{\infty} \rho^{j-1} E_t^* \left[r_{t+j}\right], pd_t\right)}{Var\left(pd_t\right)} \\
1 &= CF_2 + DR_{10} + \widetilde{LT}.
\end{aligned}
$$

The final variables $\widetilde{LT}$ captures all cash flow news beyond the two-year horizon and all discount rate news beyond the ten-year horizon.

As Table 8 shows, two-year dividend growth expectations are strongly correlated with the price-dividend ratio, resulting in large cash flow news. The two-year cash flow news of 0.59 for the full sample fits well with the one-year cash flow news of 0.36 and full horizon cash flow news of 0.73 from Section 5. Extending the scope of the discount rate news from one to ten years also produces similar results to Section 5. Ten-year

Table 8: Subjective expectations extended decomposition

	CF_2	DR_{10}	DR_{10}^{spf}
Full Sample	0.59	-0.07	-0.15
	(0.10)	(0.11)	(0.03)
Without recession	0.69	-0.34	-0.44
	(0.12)	(0.06)	(0.12)

Note. This table shows the importance of two-year cash flow news (CF_2) and ten-year discount rate news (DR_{10}) on the price-dividend variance decomposition. Using quarterly survey data from I/B/E/S and Graham-Harvey, CF_2 and DR_{10} are the coefficients obtained by regressing $\sum_{j=1}^{2} \rho^{j-1} E_t^* [\Delta d_{t+j}]$ on pd_t and $\sum_{j=1}^{10} \rho^{j-1} E_t^* [r_{t+j}]$ on pd_t respectively. We also show the ten-year discount rate news measured from the annual ten-year return forecast from the Survey of Professional Forecasters DR_{10}^{spf}. For each specification we use data running from 2003Q1 to 2015Q3 and then remove the NBER recession period spanning 2007Q4 to 2009Q2. Small-sample adjusted Newey-West standard errors in parenthesis.

subjective return expectations rise slightly with the price-dividend ratio, producing small, negative ten-year discount rate news. We perform the same test using subjective expected ten-year returns from a completely different survey, the Survey of Professional Forecasters by the Philadelphia Federal Reserve, which is denoted as DR_{10}^{spf}, and find similar results.

A.4 Additional Surveys

In addition to the one-year and ten-year return expectations obtained from the Graham-Harvey survey of Duke University (*G-H*), we use the Federal Reserve Bank of Philadelphia's Livingston Survey (*Livingston*), the University of Michigan Survey of U.S. consumers (*Michigan*) and the Survey of Professional Forecasters (*SPF*) as additional measures of S&P 500 return expectations at different horizons. All surveys are annualized returns.

The Livingston Survey is conducted twice a year by The Federal Reserve Bank of Philadelphia and spans from 1952-2016. The survey elicits forecasts of 18 different variables describing national output, prices, unemployment, and other macroeconomic data from 50-60 experts. Our variable of interest is the one-year expectations of stock market prices. Because it is prices and not returns that are forecasted, we can only build capital gains expectations $E_t^* \left[\frac{P_{t+1}}{P_t}\right]$ and not return expectations $E_t^* \left[\frac{P_{t+1}+D_{t+1}}{P_t}\right]$. Since dividends are very small compared to prices, we expect capital gain movements to be a reasonable proxy for the qualitative behavior of the expected returns. During the first years of the Livingston Survey, the S&P 400 industrial

Table 9: Explaining subjective dividend growth expectations

Regression	a	b^d	b^r
$E_t^*\left[\Delta d_{t+1}\right] = a + b^d \Delta d_{t+1} + b^r r_{t+1} + \varepsilon_t$	0.07	0.60	-0.02
	(0.01)	(0.08)	(0.04)

Notes. This table reports the results of the OLS regression of subjective dividend growth expectations on future realized dividend growth and returns for 2003Q1-2015Q3. Small-sample adjusted Newey-West standard errors are reported in parenthesis.

index was used as the forecasted index. Starting in 1990, the S&P 500 was forecasted instead. We use this survey due to the generous sampling period, obtaining similar results to the rest of the surveys.

The Michigan Survey of Consumers is conducted every month by the Survey Research Center, under the direction of the University of Michigan. The focus of the survey is on three areas: how consumers view prospects for their own financial situation, how they view prospects for the general economy over the near term, and their view of prospects for the economy over the long term. In 22 of the survey months between 2000 and 2005, the expected average return on the S&P 500 over the next 2-3 years was included in the questionnaire.

The Survey of Professional Forecasters is conducted quarterly by the Federal Reserve Bank of Philadelphia. One of the variables in the questionnaire is the forecast for the annualized average rate of return on the S&P 500 over the next 10 years. This variable is available annually from 1992-2015 and the respondents are *professional* forecasters, which is defined as those who produce regular forecasts of economic variables as part of their jobs in the business world or on Wall Street.

A.5 Testing for Risk-adjustment

In this section, we use a simple regression to test if subjective dividend growth expectations are risk-adjusted. If investors are risk-adjusting their dividend responses in the surveys, then we should find that subjective dividend growth expectations are a combination of future dividend growth and future returns, i.e. subjective dividend growth expectation should be soaking up information about future returns. Table 9 shows the results of this regression. We find no evidence that subjective dividend growth expectations are risk-adjusted. The coefficient on future returns is an insignificant -0.02, implying that subjective dividend growth expectations do not appear to be significantly influenced by future returns. In contrast, future dividend growth has a clear effect on subjective dividend growth expectations. This confirms the results of Table 7 in Section 7, which shows that subjective dividend growth expectations are not useful for predicting future returns.