Stock Market Valuation: Explanations, Non-Explanations, and Some Open Questions *

Stefan Nagel†
University of Chicago,
NBER, CEPR, and CESIfo

July 2024

*This is an expanded version of comments given at the NBER Summer Institute 2024 Asset Pricing group meeting.

†University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637, e-mail: stefan.nagel@chicagobooth.edu

I. Introduction

More than four decades ago, Shiller (1981) made the case that that aggregate stock prices are excessively volatile to be explained by a valuation model with constant expected returns. A large number of research papers since then have fleshed out empirical implications of this excess volatility hypothesis and empirical evidence on this hypothesis.

The leading framework for empirical work in this area is the approximate present value identity of Campbell and Shiller (1988). It relates the current value of the log price-dividend ratio, $p_t - d_t$, to future realizations of dividend growth and returns[1]

$$p_t - d_t = \text{const.} + \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{1}$$

Let $\mathbb{E}_t[.]$ denote objective expectations that reflect the actual law of motion of dividends, returns, and the conditioning variables observed by an econometrician. Taking these expectations on both sides of (1) yields

$$p_t - d_t = \text{const.} + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{2}$$

The Campbell-Shiller framework clarified that a robust empirical prediction of Shiller's excess volatility hypothesis is that expected returns vary over time. Accordingly, a useful definition of excess volatility is that

$$\text{var}\left(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j} \right) > 0. \tag{3}$$

How big $\text{var}\left(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j} \right)$ depends on the conditioning information considered by the econometrician. In what follows, I assume that the set of conditioning variables includes, at a minimum, the log price-dividend ratio.

Excess volatility according to this definition does not necessarily imply that stock prices

1. I use small letters to denote logs, e.g., $d_t = \log(D_t)$ is the natural logarithm of the dividend level D_t.

would be less volatile in a counterfactual world with constant expected returns. Whether this would be the case depends on the correlation of expected return news and cash flow news. If this correlation is negative, innovations to expected returns would dampen the price effects of cash flow news. Instead, what this definition means is that stock prices fluctuations cannot be fully explained by a constant discount rate model under rational expectations. In this case, the constant expected return model may predict volatility that is too high or too low, but it would not predict the correct magnitude of volatility.

The research that I discuss in Sections II and III is concerned with the existence and magnitude of excess volatility.

While the vast majority of work on resolving the excess volatility puzzle has focused on mechanisms that generate time-varying expected returns under the rational expectations assumption that endows investors with objective expectations, much recent work has considered the possibility that investors' subjective expectations differ from objective expectations—for example because investors are not have perfect knowledge of the functional form or the parameters of the data-generating process and must learn about those from realized history, or because investors are subject to behavioral biases when forming expectations. Taking subjective expectations $\widetilde{\mathbb{E}}_t[.]$ of (1) we get

$$p_t - d_t = \text{const.} + \widetilde{\mathbb{E}}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \widetilde{\mathbb{E}}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{4}$$

As an alternative to time-varying discount rate theories of asset prices, researchers in the subjective expectations literature look for mechanisms that can make $\widetilde{\mathbb{E}}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}$ volatile. If subjective expectations of cash-flow growth are time-varying, while subjective expected returns are not, and objective expectations of cash-flow growth are constant, then equating (2) and (4) shows that expected returns will be time-varying under objective expectations:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j} = -\widetilde{\mathbb{E}}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}. \tag{5}$$

Times of high subjective expectations of cash flow growth are then times of low objective expected returns.

The work that I discuss in parts of Section III and Section IV uses equity analyst earnings forecasts as proxies for cash flow expectations.

II. Is there Excess Volatility? Cash-Flow Dark Matter Explanations

Atkeson, Heathcote, and Perri (2024) (AHP) assert that there is no excess volatility puzzle. They propose that time-variation in stock market valuation can actually be explained by time-variation in rational expectations of future cash flows. Their work follows in the footsteps of a few earlier papers that aim to resurrect valuation models with (mostly) constant discount rates. Barsky and De Long (1993) argue that there is a random-walk component in dividend growth rates that can explain the volatility of stock prices, but the data in their empirical tests do not allow to convincingly demonstrate the existence of this random-walk component. The long-run risks model in Bansal and Yaron (2004) (and to a much lesser extent the version of the long-run risks model in Bansal, Kiku, and Yaron (2012)) largely shares this property that changes in stock market valuation levels reflect changes in rational expectations of long-run cash flow growth rates. Emphasizing the role of intangible capital, Hall (2001) argues that expected cash flow growth in conjunction with risk-free rate variation is sufficient to explain time-variation in stock market valuation levels. Likewise, McGrattan and Prescott (2001) argue that the stock market boom in the 1990s is fully explained by the value of intangible capital. McGrattan and Prescott (2004) argue that rational cash flow expectations explain the high stock market valuation levels in 1929 before the subsequent crash.

The earlier papers in this strand of the literature either do not offer any empirical evidence on the time-series relationship between valuation levels and future cash-flow growth (Hall; McGrattan and Prescott), inconclusive evidence (Barsky and DeLong), or the predicted cash-flow predictability is much stronger than it is in the data (Bansal and Yaron; see also Beeler and Campbell (2012)). In light of the lack of convincing evidence for a large role of rational expectations of cash-flow growth in this earlier work, it is surprising that Atkeson, Heathcote, and Perri (2024) seemingly succeed in explaining stock market valuation levels with time-variation in expected cash flow. But, as I will argue, they actually don't.

II.A. Analysis based on AHP's affine model

I first work with the same data as AHP and with AHP's affine valuation model. Later I will reach similar conclusions with a data series constructed in a more conventional way and with the log-linear approximate present value identity.

AHP use annual data from 1929 to 2023. They construct dividends from annual returns on the CRSP value-weighted index with (`vwretd`) and without dividends (`vwretx`). Binsbergen and Koijen (2010) show that extracting dividends using the difference $\texttt{vwretd}_t - \texttt{vwretx}_t$ in annual data leads to contamination of dividends with unexpected returns because this calculation implies that dividends are reinvested in the stock market during the annual return measurement periods. To avoid this problem, when I later turn to the log-linear version of the model, I construct extract dividends from monthly CRSP returns and I sum those over 12-month periods. This avoids the contamination problem. However, my conclusions regarding the analysis in AHP do not depend on the method used to construct the dividend series.

Here is a simplified version of AHP's model. Let C denote aggregate nominal personal consumption expenditure and D the per-share dividends of the CRSP index. AHP's model allows the D/C ratio to have persistent variation around an endpoint X_t,

$$\frac{D_{t+1}}{C_{t+1}} - X_t = \psi\left(\frac{D_t}{C_t} - X_t\right) + \sigma_D \varepsilon_{D,t+1}, \tag{6}$$

where X_t has random-walk dynamics,

$$X_{t+1} = X_t + \sigma_X \varepsilon_{X,t+1}, \tag{7}$$

with IID innovations $\varepsilon_{D,t+1}$ and $\varepsilon_{X,t+1}$. With constant expected returns, the ratio of the price of an index share to C is

$$\frac{P_t}{C_t} = \gamma_D \left(\frac{D_t}{C_t} - X_t \right) + \gamma_X X_t + \phi, \tag{8}$$

for some constant ϕ. In this model, variation in $\frac{P_t}{C_t}$ has two sources. First, $\frac{D_t}{C_t}$ has short-run dynamics around X_t: When $\frac{D_t}{C_t}$ is high relative to X_t, the $\frac{P_t}{C_t}$ ratio is temporarily elevated because in the near future D will be high relative to C. Second, the random-walk component X_t drives long-run dynamics: When X_t is high, future D will be permanently elevated relative to C, which induces high $\frac{P_t}{C_t}$.

In AHP's calibration $\gamma_X = 80$ (ratio of the value of a claim to the aggregate consumption stream to C), they estimate $\psi = 0.9447$, and $\gamma_D \approx 13.9$ (which is a function of ψ and γ_X).[2] AHP extract then X_t from the observed $\frac{P_t}{C_t}$ and $\frac{D_t}{C_t}$ by solving (8) for X_t. I denote this extracted value with $\widehat{X}_t$:

$$\widehat{X}_t = \frac{1}{\gamma^X - \gamma^D} \left(\frac{P_t}{C_t} - \gamma^D \frac{D_t}{C_t} - \phi \right) \tag{9}$$

In AHP's model this yields, of course, $\widehat{X}_t = X_t$.

Figure I presents the components of AHP's decomposition of P/C in (8). As the figure shows, time-variation in $\gamma_X \widehat{X}_t$ almost perfectly matches the time-variation of P/C. But what is captured by $\widehat{X}_t$ in this decomposition? The essence of AHP's story is that when $\widehat{X}_t$ varies, it captures permanent shifts in expected future D/C. But does it really?

The evidence reported in the paper that is meant to support the notion that $\widehat{X}_t$ captures shifting endpoints of D/C is based on the following regression:

$$\frac{D_{t+h}}{C_{t+h}} - \frac{D_t}{C_t} = a + \beta \left(\widehat{X}_t - \frac{D_t}{C_t} \right) + e_{t+h}. \tag{10}$$

2. AHP use the notation ρ for this constant, but I will use a different notation because I use ρ for a different constant in the log-linear present value model later

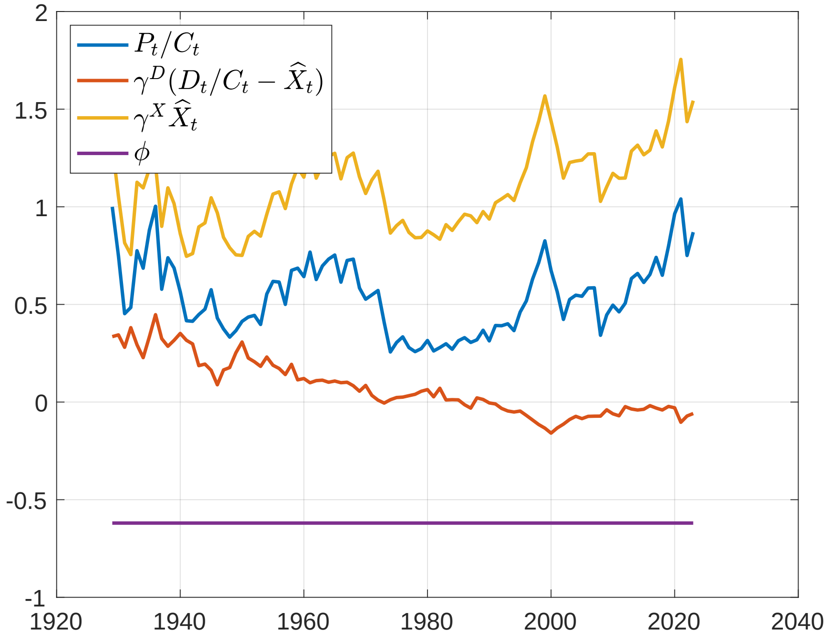

FIGURE I
AHP's decomposition of the price-consumption ratio

AHP suggest that the β coefficient in this regression captures the predictability of D/C associated with the extracted $\widehat{X}_t$. According to (6), the slope coefficient should be $\beta = 1-\psi^h$, depending on the forecast horizon h. Figure IIa replicates a similar figure in AHP, showing the theoretically implied coefficients as a function of h, as well as the empirically estimated ones. They are quite close. But are they close *because* $\widehat{X}_t$ *predicts* D/C at long-horizons?

The regression (10) mixes two effects: the persistence of deviations of D/C around $\widehat{X}_t$, and the shifting endpoints that time-varying $\widehat{X}_t$ is meant to capture. So is it D/C or $\widehat{X}_t$ that forecasts future D/C? To find out, I break up the predictor variable into its two components:

$$\frac{D_{t+h}}{C_{t+h}} - \frac{D_t}{C_t} = a + \beta_x \widehat{X}_t + \beta_{DC}\left(-\frac{D_t}{C_t}\right) + e_{t+h}. \tag{11}$$

If $\widehat{X}_t$ forecasts long-run D/C as AHP claim, I should find that β_x is equal to its theoretically implied value of $1-\psi^h$. Figure IIb shows that this is far from being true. The coefficient on $\widehat{X}_t$ is even slightly negative. Basically, all the predictive power of $\widehat{X}_t - \frac{D_t}{C_t}$ in AHP's regression

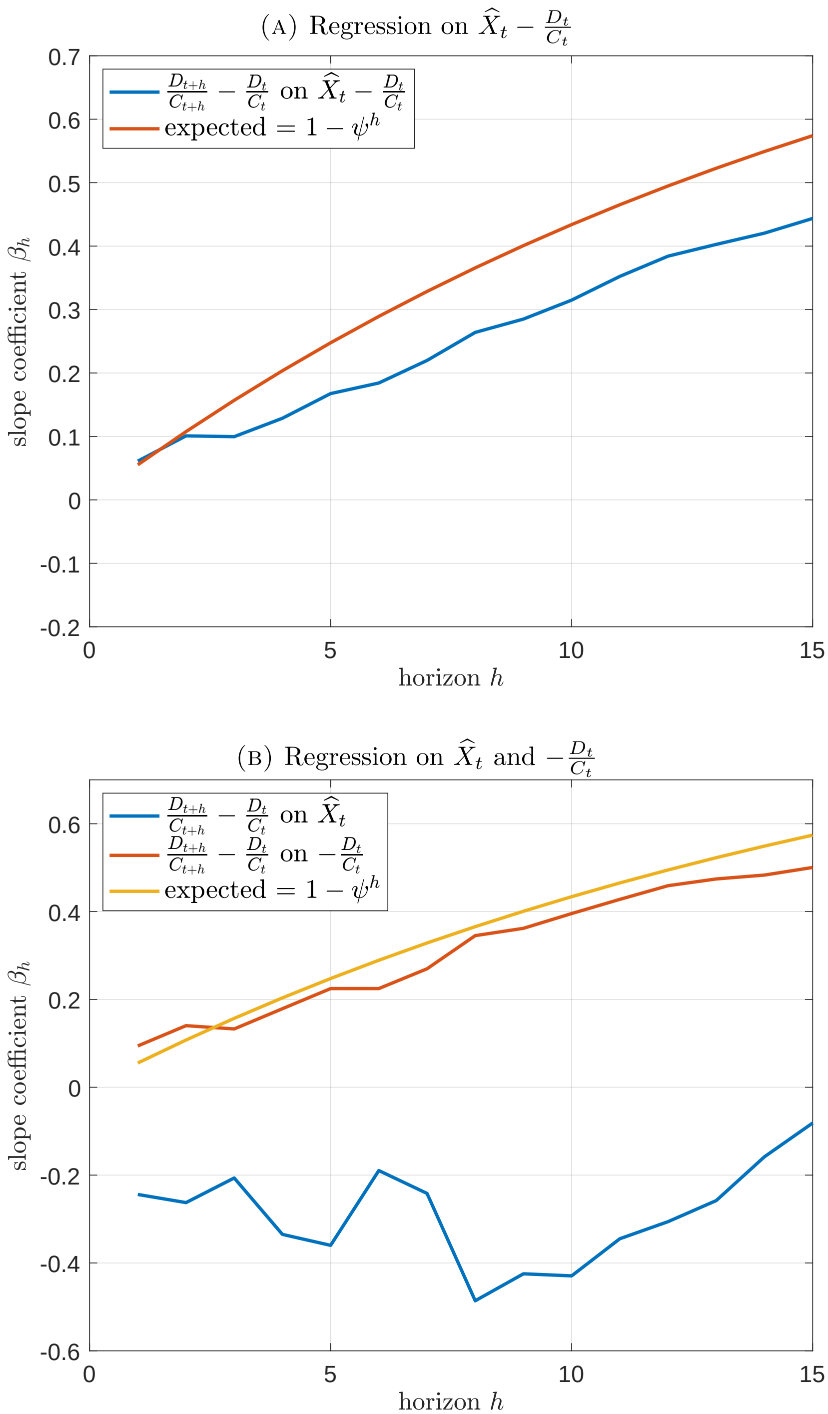

FIGURE II
Slope Coefficient Estimates

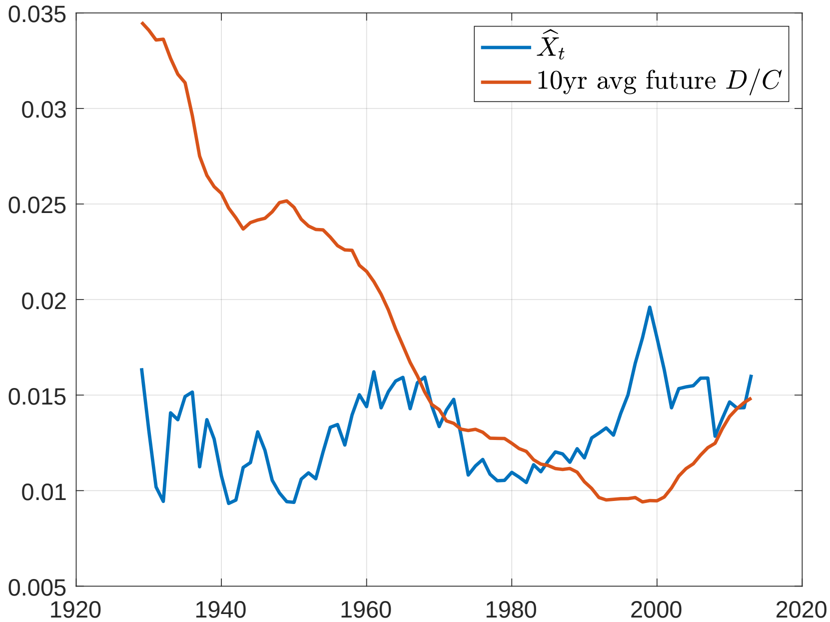

FIGURE III
$\widehat{X}_t$ and future long-run average D/C

(10) comes from the persistence of deviations of D_t/C_t from its mean, and essentially nothing from time-variation in $\widehat{X}_t$. Times of high $\widehat{X}_t$ are not times of high expected future D/C, contrary to AHP's assertion.

Figure III further illustrates the disconnect between $\widehat{X}_t$ and long-run D/C. The figure shows that there is basically no relation between the level of $\widehat{X}_t$ and the average D/C over the next 10 years. In particular, recall from Figure I that the movements in $\widehat{X}_t$ up and down over periods of around a decade are extremely highly correlated with the P/C ratio and thus key to explaining why stock market valuation levels change over time. This figure shows that movements in $\widehat{X}_t$ at those frequencies have nothing to do with long-run D/C.

II.B. An alternative interpretation

If movements in $\widehat{X}_t$ are unrelated to long-run D/C, then what does $\widehat{X}_t$ capture? Here is an alternative model. Consider a possibly highly persistent but stationary D/C with that varies

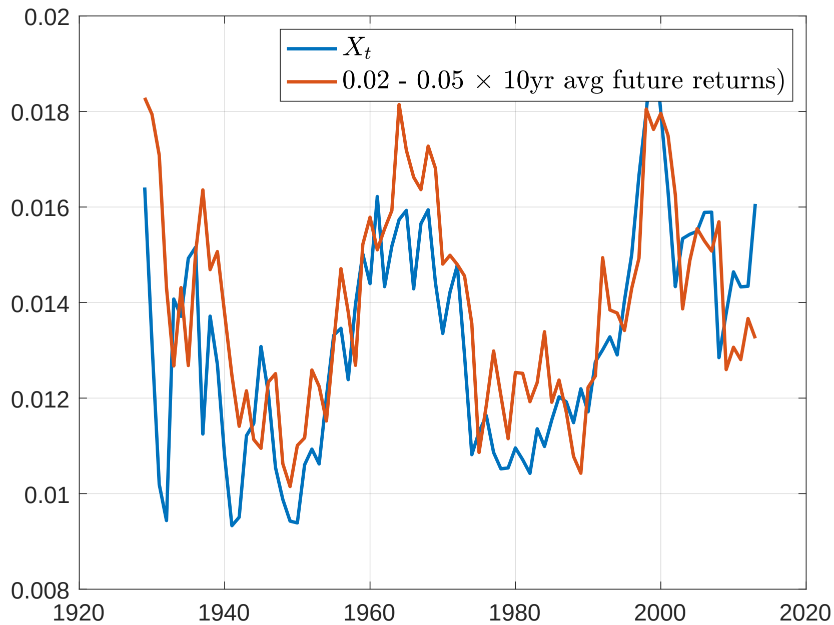

FIGURE IV
$\widehat{X}_t$ and future long-run returns

around a constant X:

$$\frac{D_{t+1}}{C_{t+1}} - X = \psi\left(\frac{D_t}{C_t} - X\right) + \sigma_D \varepsilon_{D,t+1}. \tag{12}$$

In valuation, I now allow for time-varying expected returns, captured by a time-varying ϕ_t:

$$\frac{P_t}{C_t} = \gamma_D\left(\frac{D_t}{C_t} - X\right) + \gamma_X X + \phi_t. \tag{13}$$

Now extract $\widehat{X}_t$ as before with the same calculation as in (9), with $\phi = \mathbb{E}[\phi_t]$. In this alternative model, this calculation yields

$$\widehat{X}_t = X + \phi_t - \phi$$

i.e., in this alternative model $\widehat{X}_t$ perfectly isolates time-variation in expected returns!

Now let's examine empirically whether $\widehat{X}_t$ is related to expected returns, as this alternative model predicts. Figure IV shows the time series of $\widehat{X}_t$ and (an affine function of) the average

TABLE I

Predictive regressions with $\widehat{X}_t$ as predictor

	R_{t+1}	R_{t+2}	R_{t+3}	R_{t+4}	R_{t+5}
Panel A: Return prediction with P_t/D_t					
100× coeff.	-0.27	-0.24	-0.20	-0.17	-0.16
(t-stat.)	(-1.98)	(-1.69)	(-1.25)	(-1.23)	(-1.19)
R^2	0.03	0.02	0.01	0.01	0.00
Panel B: Return prediction with $\widehat{X}_t$					
coeff.	-33.20	-29.00	-19.66	-15.18	-16.04
(t-stat.)	(-3.86)	(-3.22)	(-2.26)	(-1.50)	(-1.68)
R^2	0.13	0.10	0.04	0.02	0.03

stock market return over the subsequent 10 years. Consistent with the alternative model, $\widehat{X}_t$ is very highly correlated with future average returns!

Table I shows predictive regressions.[3] For comparison, Panel A shows OLS regressions of annual stock market index returns R_{t+h} on the price-dividend ratio at the end of period t. Panel B shows regressions on $\widehat{X}_t$. As evident from the t-statistics and R^2, $\widehat{X}_t$ is a much better predictor of future returns than the price-dividend ratio.

All this leads to exactly the opposite of AHP's conclusion. They argue that they have identified a component of P/C that is capturing future cash flows, thereby explaining away the excess volatility puzzle. Instead, this component is actually strongly related to expected returns, which strengthens the excess volatility puzzle.

II.C. A log-linear version of the alternative model

To connect the discussion of AHP's analysis with issues that I will touch on later, it is useful to go through a version of the alternative model in a log-linear framework. Consider dynamics

3. To reduce clutter, I do not report the intercept estimates in this table and the others that follow.

for the log dividend-consumption ratio, $d_t - c_t = \log(D_t/C_t)$, of the following form:

$$d_{t+1} - c_{t+1} - x = \psi(d_t - c_t - x) + \sigma_d \varepsilon_{d,t+1}. \tag{14}$$

As before, dividends are measured as per-share dividends, and consumption is aggregate personal consumption expenditure with IID growth dynamics

$$\Delta c_{t+1} = g + \sigma_c \varepsilon_{c,t+1}, \tag{15}$$

where x and g are constants and $\varepsilon_{d,t+1}$ and $\varepsilon_{c,t+1}$ are IID residuals. Combining (14) and (15) yields

$$\Delta d_{t+1} = -(1-\psi)(d_t - c_t - x) + g + \sigma_d \varepsilon_{d,t+1} + \sigma_c \varepsilon_{c,t+1}, \tag{16}$$

and

$$\mathbb{E}_t \, \Delta d_{t+j} = -\psi^{j-1}(1-\psi)(d_t - c_t - x) + g. \tag{17}$$

In this model, when d_t deviates from c_t by more than x, it will have a tendency to revert back to $c_t + x$. Hence, $d_t - c_t - x$ predicts future dividend growth.

Applying the Campbell-Shiller PV identity and using (17), I obtain

$$p_t - d_t = \text{const.} - \underbrace{\left(\frac{1-\psi}{1-\psi\rho}\right)(d_t - c_t)}_{\text{Expected dividend growth}} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{18}$$

Variation of $p_t - d_t$ therefore partly reflects variation in expected dividend growth associated with variation in $d_t - c_t$. This distorts the signal in $p_t - d_t$ about expected returns. But in analogy to AHP's calculation of $\hat{X}$ in (9), we can perform an adjustment to the log price-dividend ratio based on $d_t - c_t$ that removes the expected dividend growth component and

TABLE II

Predictive regressions with $\hat{x}_t$ as predictor

	r_{t+1}	r_{t+2}	r_{t+3}	r_{t+4}	r_{t+5}
Panel A: Return prediction with $p_t - d_t$					
$p_t - d_t$	-0.083	-0.101	-0.064	-0.065	-0.075
(t-stat.)	(-1.92)	(-2.43)	(-1.48)	(-1.54)	(-1.78)
R^2	0.04	0.06	0.02	0.03	0.04
#Obs.	92	91	90	89	88
Panel B: Return prediction with $\hat{x}_t$					
$\hat{x}$	-0.192	-0.178	-0.095	-0.128	-0.131
(t-stat.)	(-3.27)	(-3.05)	(-1.54)	(-2.13)	(-2.16)
R^2	0.11	0.09	0.03	0.05	0.05
#Obs.	92	91	90	89	88

perfectly isolates time-varying expected returns:

$$\begin{aligned} \hat{x}_t &= p_t - d_t + \left(\frac{1-\psi}{1-\psi\rho}\right)(d_t - c_t) \\ &= \text{const.} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \end{aligned} \tag{19}$$

Figure V presents the log-linear model equivalents to Figures III and IV. The top panel shows that $\hat{x}_t$ is unrelated to future dividend growth (the correlation of the two series is 0.00), which is similar to the lack of positive correlation between $\widehat{X}_t$ and future D/C in Figure III. The bottom panel shows that $\hat{x}_t$ is strongly related to future returns during the subsequent 10 years (the correlation of the two series is 0.86), which is similar to the strong correlation between $\widehat{X}_t$ and future returns in Figure IV.

Table II shows predictive regressions for future returns. Analogous to Table I, $\hat{x}_t$ is a better predictor of future returns than $p_t - d_t$.

(A) $\hat{x}_t$ and future dividend growth $\sum_{j=0}^{9} \rho^j \Delta d_{t+1+j}$

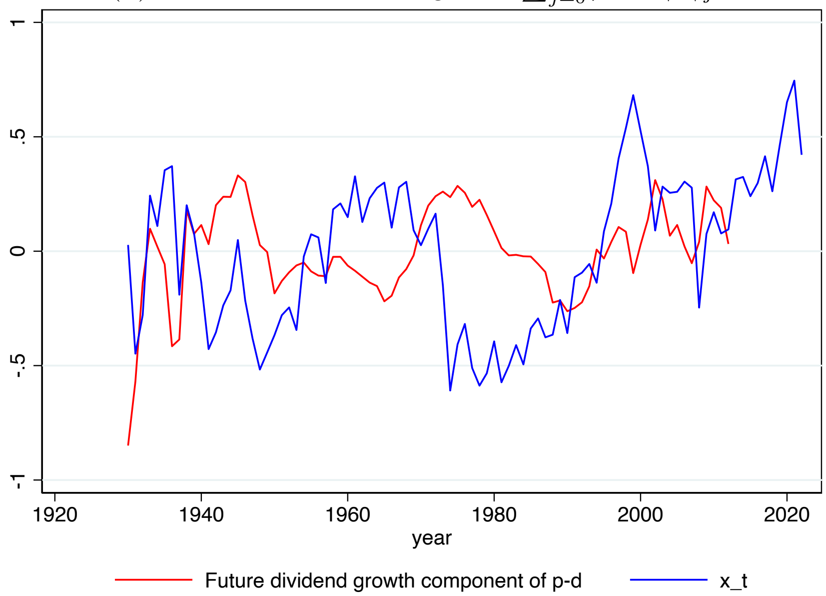

(B) $\hat{x}_t$ and negative future returns $-\sum_{j=0}^{9} \rho^j r_{t+1+j}$

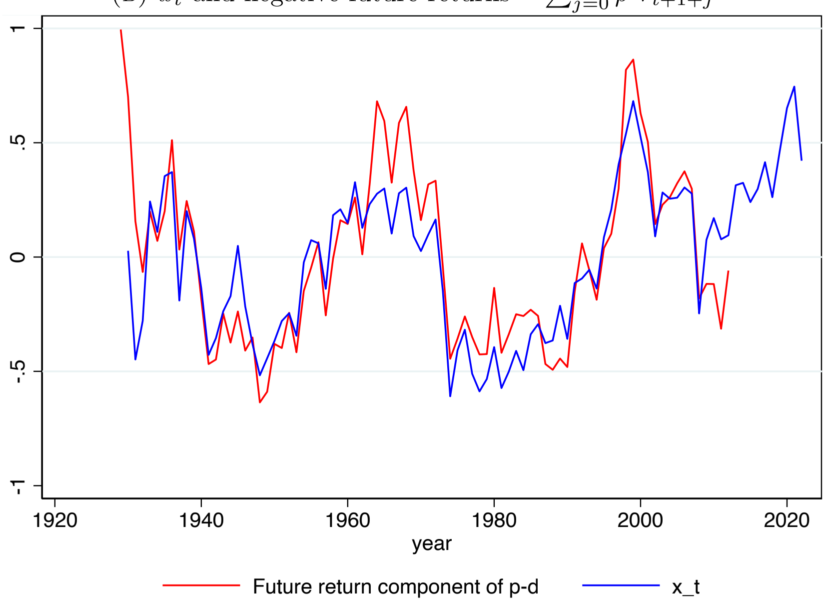

Figure V
Relationship of $\hat{x}_t$ with future cash flows and returns

TABLE III
Dividend growth prediction with $d_t - c_t$ as predictor

	Δd_{t+1}	Δd_{t+2}	Δd_{t+3}	Δd_{t+4}	Δd_{t+5}
$d_t - c_t$	-0.077	-0.069	-0.023	0.001	0.003
(t-stat.)	(-2.86)	(-2.58)	(-0.97)	(0.05)	(0.13)
R^2	0.08	0.07	0.01	0.00	0.00
#Obs.	92	91	90	89	88

II.D. Why is $\hat{x}_t$ is a better predictor of future returns than $p_t - d_t$?

(17) and (18) in the log-linear alternative model show the reason why $\hat{x}_t$ predicts returns better than $p_t - d_t$ does: $p_t - d_t$ is not a clear signal of future returns because it is distorted by time-variation in expected dividend growth that is captured by $d_t - c_t$. Table III presents evidence that this predicted relationship between $d_t - c_t$ and future dividend growth is also present in the data. The calculation of $\hat{x}$ removes this distortion from $p_t - d_t$.

II.E. Why does $d_t - c_t$ predict dividend growth?

This now pushes the question one level further. Why does $d_t - c_t$ predict dividend growth? To understand, we need to look at AHP's definition of the dividend-consumption ratio. AHP calculate

$$\frac{D_t}{C_t} = \frac{\text{Dividends per share}}{\text{Aggregate PCE}}. \tag{20}$$

So with Q_t as index of quantity of outstanding shares, which implies that aggregate dividends are given by $D_{agg,t} = Q_t D_t$, we can write this as

$$\frac{D_t}{C_t} = \frac{\text{Aggregate Dividends}}{\text{Aggregate PCE}} \times \frac{1}{Q_t}, \tag{21}$$

or in logs,

$$d_t - c_t = d_{agg,t} - c_t - q_t. \tag{22}$$

TABLE IV
Dividend growth prediction with $d_{agg,t} - c_t$ and q_t

	Δd_{t+1}	Δd_{t+2}	Δd_{t+3}	Δd_{t+4}	Δd_{t+5}
$d_{agg,t} - c_t$	-0.224	-0.242	-0.098	-0.012	0.020
(t-stat.)	(-3.56)	(-3.88)	(-1.65)	(-0.20)	(0.32)
q_t	0.094	0.088	0.031	0.000	-0.005
(t-stat.)	(3.48)	(3.35)	(1.28)	(0.01)	(-0.19)
R^2	0.15	0.16	0.03	0.00	0.00
#Obs.	92	91	90	89	88

Now, changes in index composition can affect both q_t, and thereby $d_t - c_t$, and, at the same time, future dividend-per-share growth. For example, if new issues come to the market that pay very low current dividend but have high expected future dividend growth, then current dividends per share, $d_{agg,t} - q_t$ fall, $d_t - c_t$ falls as well, and at the same time the expected per share dividend growth for the index rises. The evidence presented in Jank (2015) that high D/P firms exited and low D/P firms entered public markets in the 1970s and the following decades is broadly consistent with this story.

Table IV shows similar predictive regressions for dividend growth as those in Table III, but now with $d_t - c_t$ replaced by its components $d_{agg,t} - c_t$ and q_t. As the table shows, variation in q_t is, indeed, strongly related to future dividend growth during the next years. A note of caution, though: q_t is extremely persistent, so the OLS regressions in this table may not deliver reliable estimates of this relationship.

II.F. Summary

Contrary to AHP's claim, there is no evidence that changes in stock market valuation levels over time are associated with shifts in the expected long-run dividends-per-share to consumption ratio. Hence, the excess volatility puzzle is alive and well. In fact, AHP's method of modifying the price consumption ratio accidentally uncovered a method for removing from the price-consumption ratio (or the price-dividend ratio in my log-linear version of the model) the distorting effects of predictably time-varying future cash flows. Removing this expected cash flow component yields a return predictor that has a much stronger relationship to future returns than the unadjusted price-consumption or price-dividend ratio. While AHP claim to show that there is no excess volatility puzzle, they have unintentionally discovered a return predictor that *strengthens* the excess volatility puzzle.

III. Interpretation of Valuation Ratio Variance Decompositions

Variance decompositions of valuation ratios are a common approach to tackling the excess volatility question. Present value identities and variance decompositions are available for a variety of different types of valuation ratios. Campbell and Shiller (1988) decompose the price-dividend ratio. Vuolteenaho (2000) looks at the market index market-to-book equity ratio, while Cohen, Polk, and Vuolteenaho (2003) and Cohen, Polk, and Vuolteenaho (2009) study variance decompositions of market-to-book equity ratios in the cross-section of firms. Maio and Xu (2020) examine a variance decomposition of the index-level earnings yield, De La O, Han, and Myers (2023) (DHM) look at the price-to-earnings ratio in the cross-section. Motivated by a present value identity written in terms of aggregate rather than per-share quantities, Bansal and Yaron (2006), Larrain and Yogo (2008) and Pruitt (2023) argue that one should look at cash flows to shareholders in aggregate, which include share repurchases and, as a negative cash flow, equity issues. They therefore examine a variance decomposition of an aggregate price-to-payout ratio.

While Campbell and Shiller (1988) found that most of the time-variation in the price-dividend ratio is associated with future returns rather than future cash flows, variance decompositions of other valuation ratios often produce a bigger role for expected future cash flows. These findings on the sources of time-variation in *valuation ratios* lead some authors to conclude that more *stock price* variation is attributable to expected cash flow variation than one would think based on the price-dividend ratio variance decomposition.

The motivation for focusing on variance decomposition of valuation ratios, $p_t - z_t$, rather than, say, a variance decomposition of unexpected returns, $r_{t+1} - \mathbb{E}_t r_{t+1}$, as in Campbell (1991) is to figure out how much expected return and expected cash flow variation matters for price levels. For example, if expected returns are mean-reverting fast, expected return variation could matter a lot for the variance of unexpected returns, but play only a small role in generating price level variation.

III.A. Lack of interpretability of valuation ratio variance decompositions

The problem with interpreting the results from valuation ratio variance decompositions is that the variance share of the expected return component can be almost anything depending on the choice of the denominator in the valuation ratio. Take the Campbell-Shiller PV identity and add $d_t - z_t$ on both sides

$$p_t - z_t = \text{const.} + d_t - z_t + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{23}$$

Even if we keep the conditioning information that pins down the expectations fixed, the variance share

$$\frac{\text{var}\left(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}\right)}{\text{var}(p_t - z_t)} \tag{24}$$

is not invariant to the choice of z because, as one can see by taking the variance of the right-hand side of (23), $\text{var}(p_t - z_t)$ depends on $\text{var}(d_t - z_t)$ and the covariance of $d_t - z_t$ with the other terms on the right-hand side of the present value identity. For instance, by picking a variable $z_t = d_t +$ noise with sufficiently big noise variance, one can make the variance share of the expected return component arbitrarily small.

However, while changing the denominator of a valuation ratio can change the share of *valuation ratio* variance attributed to expected cash flows and expected return movements, changing the denominator of a valuation ratio does not change how much of *stock price* movement is associated with expected cash flows and how much with expected return movements. If one wishes to understand whether stock prices move too much to be explained by movements in expected cash flows without movements in expected returns, a variance decomposition of valuation ratios does not directly answer this question.

In essence, the question whether stock prices move too much to be explained by movements in expected cash flows without movements in expected returns really asks whether there are transitory components in stock prices. Consider the Campbell-Shiller present value identity slightly rearranged, with only p_t on the left-hand side:

$$p_t = \text{const.} + d_t + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{25}$$

If the price moves in period $t+1$ because of news about future dividend growth in periods $t+2$ and beyond, with expected returns remaining fixed, then

$$r_{t+1} - \mathbb{E}_t\, r_{t+1} = \Delta\, \mathbb{E}_{t+1} \sum_{j=1}^{\infty} \rho^j \Delta d_{t+1+j}. \tag{26}$$

Since there is no change in future expected returns, this price change induced by cash flow news is permanent. In contrast, if the price moves in period $t+1$ because of news about future expected returns in periods $t+2$ and beyond, with expected dividend growth fixed, and also $d_{t+1} = \mathbb{E}_t\, d_{t+1}$, so no cash flow news at all, then

$$r_{t+1} - \mathbb{E}_t\, r_{t+1} = -\Delta\, \mathbb{E}_{t+1} \sum_{j=1}^{\infty} \rho^j r_{t+1+j}. \tag{27}$$

The entire price change induced by expected return news is temporary, because it is subsequently reversed by expected returns in opposite directions, as captured by all the future expected return terms on the right-hand side of this equation.[4]

4. Since $\rho < 1$, a one unit increase in r_{t+1} is subsequently reversed by a sum of log returns bigger than one unit. In the Campbell-Shiller log-linearization, $\rho = 1/(1 + \exp(\mathbb{E}[d_t - p_t]))$, so this difference in magnitude between initial log return and subsequent reversal is bigger when $\mathbb{E}[d_t - p_t]$ is higher. Following Campbell

Therefore, $V_r = \text{var}(\mathbb{E}_t \sum_{j=1}^{\infty} \rho^j r_{t+1+j})$ represents the variance of the transitory component of stock prices:

$$\text{var}(\underbrace{p_t - d_t - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}}_{\text{transitory component of } p_t}) = \text{var}(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}) \tag{28}$$

For example, a one standard deviation move in the transitory component (with cash flows fixed) has magnitude $\sqrt{V_r}$. This number is meaningful on its own, without putting it in relation to some other variance. It tells that a 1 S.D. move in the transitory component (w/ c.f. fixed) of the price is roughly $100 \times \sqrt{V_r}$ percent.

Importantly, fixing the set of conditioning variables fully pins down V_r. How big this variance is does not depend on the choice of denominator in a valuation ratio. It only depends on which predictor variables are used, and which statistical model is employed by the econometrician to forecast returns. Focusing on the variance of $\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}$, rather than a variance decomposition, therefore seems a better way to address the question of "what moves stock prices?" or "is the price right?".

III.B. DHM: Price-to-book vs. price-to-earnings ratios

A good example for the lack of interpretability of valuation ratio variance decomposition is provided by DHM. They study the cross-section of log price-to-earnings ratios $p_t - e_t$ and they compare their findings to the study of log price-to-book ratios, $p_t - b_t$, in Cohen, Polk, and Vuolteenaho (2003) and Cohen, Polk, and Vuolteenaho (2009). DHM find that expected return variation accounts for a much larger share of variation in $p_t - e_t$ than in $p_t - b_t$.

The essence of DHM's analysis can be understood by assuming $d_t \approx \text{const.} + e_t$, which

(1991), this can be understood through the following example: Suppose stock returns will be higher 10 periods from now while the path of dividends is fixed. The price falls today, but this fall in price today also means that the dividend yield will be higher in the next 9 periods. For this higher dividend yield not to yield not to imply higher returns during the next 9 periods (by assumption, we are only considering higher returns 10 periods from now), the price has to fall further during the next nine periods to generate a capital loss that offsets the higher dividend yield.

means that the payout ratio is approximately constant. By using this relationship to substitute out d_t on both sides, the Campbell-Shiller present value identity then can be written as

$$p_t - e_t = \text{const.} + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta e_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{29}$$

We can further get to a present value identity for $p_t - b_t$ by adding profitability, $e_t - b_t$, on both sides,

$$p_t - b_t = \text{const.} + \underbrace{e_t - b_t + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta e_{t+1+j}}_{\text{can be expressed as future log ROE}} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{30}$$

DHM correctly point out that $e_t - b_t$ adds profitability variation on the right-hand side that lowers the variance share of expected returns. This is a special case of the generic case I discussed above: Valuation ratio variance decompositions are sensitive to which variable is used in the denominator of the valuation ratio.

Table V illustrates this with DHM's annual panel data for 5 value/growth portfolios at holding horizons (after portfolio formation) between 1 and 15 years. I estimate a VAR on this panel. The VAR includes both $p_t - e_t$ and $p_t - b_t$ as predictor variables, as well as lagged earnings growth and returns. The table shows the variance decompositions implied by the estimated VAR. The results are similar to those in Tables I and II of DHM, but I show the magnitudes of the covariances and variances, not only the variance decomposition shares.

As the table shows, in column (1), the $p_t - b_t$ ratio is far more volatile than the $p_t - e_t$ ratio. And much of the difference comes from the fact that $p_t - b_t$ is exposed to the variance of $e_t - b_t$, as shown in column (2). Much of this variance reflects the fact that profitability levels differ substantially between value and growth stocks. As a consequence, expected return variation commands a much lower share in the variance decomposition of $p_t - b_t$ than in the variance decomposition of $p_t - e_t$, as shown in column (4). This is in line with DHM's argument.

TABLE V
Variance decomposition of $p_t - e_t$ and $p_t - b_t$ for panel of value/growth portfolios

	Covariance of $p-z$ with				
	$p-z$	$e-z$	$\mathbb{E}_t \sum \rho^j \Delta e$	$-\mathbb{E}_t \sum \rho^j r$	
	(1)	(2)	(3)	(4)	(4)/(1)
z = e	0.071	n/a	0.017	0.055	0.775
z = b	0.196	0.093	0.020	0.088	0.449

At the same time, it's not clear why one would want to focus on variance decompositions of different valuation ratios as we can get an unambiguous answer about the amount of transitory variation in stock prices (and hence to what extent the "price is right"): Just look at $\text{var}(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j})$. This quantity is exactly the same in (29) and (30). Specifically, the estimated VAR implies $\text{var}(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j})^{\frac{1}{2}} = 0.209$. Hence, a one standard deviation move in the transitory component of stock prices (with cash flows kept fixed) is roughly 20.9 percent.

III.C. Price-dividend vs. price-payout ratios

Similar comments apply to papers that have discussed the use of price-payout ratios instead of price-dividend ratios. The Campbell-Shiller present value identity considers the perspective of an investor who holds on to a fixed number of shares of a stock or an index. This investor does not participate in repurchases or new stock issues and hence the only cash flows this investor receives are cash dividends. An alternative perspective is to consider a representative investor who holds all outstanding shares. This investor receives cash flows from repurchases and experiences negative cash flows from equity issues.

Let Q_t be the number of shares outstanding after repurchases and equity issues in period t have been completed. Let Y_t denote net payout, i.e., dividends plus repurchases minus equity issues, expressed in dollars per share outstanding at time $t-1$. Aggregate net payout then is Y_tQ_{t-1}. For simplicity, assume $Y_t > 0$.[5] One can write the one period gross return as $R_{t+1} = (P_{t+1}Q_{t+1} + Y_{t+1}Q_t)/(P_tQ_t)$.[6] Following the same steps as in the derivation of the Campbell-Shiller present value identity then yields an expression for the log price to aggregate payout ratio $P_tQ_t/(Y_tQ_{t-1})$ (see Appendix A):

$$p_t - y_t + \Delta q_t = \text{const.} + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j (\Delta y_{t+1+j} + \Delta q_{t+j}) - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j} \tag{31}$$

While returns are the same, cash-flow growth is not the same in the aggregate and per-share version of the present value identity. For example, if firms substitute repurchases for dividends today, this has no effect on aggregate payouts now or in the future, but it lowers today's dividends per share, it raises the current price-dividend ratio, and it raises future per-share growth rates of dividends. As a consequence, the variance decompositions of price-dividend and price-payout ratios are not the same. Specifically, as Larrain and Yogo (2008) and Bansal and Yaron (2006) show

$$\frac{\text{var}\left(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}\right)}{\text{var}(p_t - y_t + \Delta q_t)} < \frac{\text{var}\left(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}\right)}{\text{var}(p_t - d_t)}. \tag{32}$$

Alas, the same comment applies here as in the comparison of $p_t - b_t$ to $p_t - e_t$: Given the same conditioning information, $\text{var}\left(\mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}\right)$ is the same in each variance decomposition. So per-share and aggregate PV identity yield the same answer about magnitude of transitory movement in stock prices. Switching from the per-share to the aggregate present value identity does not change how much transitory movement there is in stock prices.

Another way of seeing this is to consider an unexpected return decomposition as in Camp-

5. Larrain and Yogo (2008) and Bansal and Yaron (2006) show how with minor modifications one can write a present value identity that allows for negative payouts.

6. Note that this return is equal to the per-share return $(P_{t+1} + D_{t+1})/P_t$. Returns are the same on an aggregate or per-share basis.

bell (1991). Based on the price-dividend ratio present value identity we get

$$r_{t+1} - \mathbb{E}_t\, r_{t+1} = \Delta\, \mathbb{E}_{t+1} \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \Delta\, \mathbb{E}_{t+1} \sum_{j=1}^{\infty} \rho^j r_{t+1+j}. \tag{33}$$

If we do it based on the aggregate PV identity, we get

$$r_{t+1} - \mathbb{E}_t\, r_{t+1} = \Delta\, \mathbb{E}_{t+1} \sum_{j=0}^{\infty} \rho^j (\Delta y_{t+1+j} + \Delta q_{t+j}) - \Delta\, \mathbb{E}_{t+1} \sum_{j=1}^{\infty} \rho^j r_{t+1+j}. \tag{34}$$

Since the term of the left-hand side is the same in both equations, and the second term on the right-hand side is, too, the remaining term must be the same in both equations, too. (As before, we assume the same conditioning information in both cases). Hence, while the variance decomposition of the aggregate payout ratio $p_t - y_t + \Delta q_t$ may differ from the variance decomposition of $p_t - y_t$ or $p_t - d_t$, the decomposition of unexpected returns is identical.[7]

III.D. Do short- or long-run cash flow growth expectations explain asset price movements?

Different views have appeared in the literature about the role that short- and long-run earnings expectations play in explaining asset price movements for the aggregate stock market. De La O and Myers (2021) find that short-term growth expectations (STG) explain large share of time-variation in the price-earnings ratio. In contrast, Nagel and Xu (2022) and Bordalo, Gennaioli, La Porta, and Shleifer (2024a) find that long-term growth expectations (LTG) rather than STG explain much of the valuation cycles in the stock market and the return predictability associated with these.

A comparison of the earnings definition in DHM and De La O and Myers (2021) provides a hint at a resolution. This case illustrates once more how the results of variance decompositions

7. Appendix B presents an example to show how a surprise repurchase that does not affect returns, and hence neither the left-hand side, nor the second term on the right-hand side of both equations, also has no effect on the cash-flow growth terms in both equations (even though it results in a change of some of the per-share expected dividend growth terms in (33), but these changes cancel out).

TABLE VI

Share of $\text{var}(p_t - e_t)$ attributed to STG variation with different earnings definitions

CRSP vw. index, 1985 - 2023, annual, end of June prices and forecasts.

	$\frac{\text{cov}(p_t - e_t, \widetilde{\mathbb{E}}_t \Delta e_{t+1})}{\text{var}(p_t - e_t)}$
De La O and Myers (2021) definition	52%
DHM definition	16%

can change a lot depending on which variable is in the denominator of the valuation ratio. De La O and Myers (2021) define earnings as before extraordinary items. In contrast, DHM define earnings as before extraordinary items, special items, and non-recurring taxes.

The exclusion of special items from DHM's earnings definition is the important difference. During recessions, special items become strongly negative. As a consequence, even though stock prices fall, the earnings in the De La O and Myers (2021) definition fall even more, leading to a sharp rise in $p_t - e_t$. At the same time, analysts expectations of next year's earnings are much higher than current earnings depressed by special items, and so $\widetilde{\mathbb{E}}_t e_{t+1} - e_t$ is very high. This generates the very strong positive covariance between $p_t - e_t$ and $\widetilde{\mathbb{E}}_t e_{t+1} - e_t$ in De La O and Myers (2021). As Table VI shows, the share of price-earnings ratio variance attributed to STG variation drops from 52% to 16% when special items are excluded. This shows that time-variation in STG only plays a relatively small role in explaining changes in stock prices. Much of STG's high variance share in De La O and Myers (2021) is due to the quickly reverting noise that special items add to earnings, not because of price movements associated with STG.

III.E. Valuation ratios as return predictors

While variance decompositions of valuation ratios are not suitable for assessing the amount of transitory variation in stock prices, judicious choice of valuation ratios can be useful to improve prediction of returns. Different valuation ratios may perform better or worse in predicting returns. So which valuation ratios an econometrician should condition expectations on is a meaningful question.

Here is an example in which the price-payout ratio (on a per-share basis) predicts returns better than the price-dividend ratio. Suppose repurchases substitute for dividends and leave aggregate payouts, $\Delta y_t + \Delta q_{t-1}$, unchanged. Suppose further that aggregate payout growth is unpredictable. This implies that per-share payouts Δy_{t+1} are negatively correlated with Δq_t: Higher repurchases in period t, which lowers Δq_t, raises per-share payout growth Δy_{t+1} in the next period. For the same reason, it raises per-share dividend growth Δd_{t+1}, too, while lowering current dividends d_t due to the expense for the repurchases. This means that the repurchase at time t generates positive correlation between $p_t - d_t$ on the left-hand side of the present value identity

$$p_t - d_t = \frac{\kappa}{1-\rho} + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j} \tag{35}$$

and $\mathbb{E}_t \Delta d_{t+1}$ on the right-hand side. This variation in $p_t - d_t$ associated with $\mathbb{E}_t \Delta d_{t+1}$ induced by repurchases distorts the ability of $p_t - d_t$ to predict returns. In contrast, in the present value identity for the per-share log price-payout ratio (see Appendix A),

$$p_t - y_t = \frac{\kappa}{1-\rho} + \frac{\rho}{1-\rho}\Delta q_t + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}, \tag{36}$$

the term with Δq_t on the right-hand side is negatively correlated with $\mathbb{E}_t \Delta d_{t+1}$ which cancels the predictable variation in $\mathbb{E}_t \Delta d_{t+1}$. Intuitively, while repurchases depress current dividends per-share, they do not depress current payouts per share. Hence, $p_t - y_t$ is not distorted in its ability to predict returns.

The adjustments to the price-consumption or price-dividend ratio discussed in Section II building on AHP's analysis are another example of how modifications of valuation ratios can improve their ability to predict returns.

III.F. Summary

Understanding the sources of variation in valuation ratios may be of interest for a number of different reasons. But these variance decompositions do not offer a clear view of the magnitude of excess volatility. Variance decompositions of different valuation ratios deliver seemingly different messages about the role of expected return variation in explaining stock price movements because they show what explains valuation ratio movements, not stock price movement. Researchers who are interested in estimating the amount of transitory variation in stock prices should focus on estimating the variance of the transitory component directly.

IV. Are Analyst Forecast a Useful Proxy of Investor Expectations?

I now turn to studies that examine proxies for subjective expectations of investors. As analytical framework, I use the Campbell-Shiller present value identity under subjective expectations (4), restated here:

$$p_t - d_t = \text{const.} + \widetilde{\mathbb{E}}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \widetilde{\mathbb{E}}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j} \tag{37}$$

Given the empirical evidence that stock price movements have a large transitory component associated with time-varying expected returns under objective probabilities, the key question is whether investors' subjective expectations of returns share this time-variation or whether time-varying cash flow growth expectations are instead the driver of valuation levels under the subjective expectations of investors. A variety of different surveys provide return expectations for different types of market participants and forecasters (see, e.g., Nagel and Xu 2023), but for cash flow expectations typically the only source of data are equity analysts' earnings forecasts. There are a number of open questions about the interpretation of these analyst forecasts.

IV.A. Do analysts back out growth rates to fit observed prices?

One concern with analyst earnings expectations as a proxy for investor expectations is that analysts might simply report growth rates that are backed out from observed prices. If analysts use the same discount rate as investors, then this would not really be a problem, as the cash flow expectations backed out from prices would then represent investors' expectations (assuming analysts also use roughly the same term structure of cash flow growth as investors do).

But the main concern is that investors might price stocks with a time-varying risk premium and constant $\mathbb{E}_t\,\Delta d_{t+1+j}$ under rational expectations so that

$$p_t - d_t = \text{const.} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{38}$$

If analysts work with constant discount rate and back out $\widetilde{\mathbb{E}}_t \Delta d_{t+1+j}$ that fit $p_t - d_t$ by solving

$$p_t - d_t = \text{const.} + \widetilde{\mathbb{E}}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} \tag{39}$$

for cash flow growth expectations, then analysts' "cash-flow expectations" actually capture investors' risk premia:

$$\widetilde{\mathbb{E}}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} = \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{40}$$

In this case, analysts' forecast errors would also be predictable and correlated with predictable returns, just as if analysts beliefs reflected mistaken investor beliefs that expected cash flow growth is time-varying.

Bordalo, Gennaioli, La Porta, and Shleifer (2024b) (BGLS) and Bordalo, Gennaioli, La Porta, and Shleifer (2024a) offer a number of empirical observations that are meant to refute the idea that analysts back out their reported LTG from prices. However, upon closer inspection, these empirical facts cannot rule out that by inferring LTG from prices, analysts end up reporting LTG that is contaminated with risk premia.

First, Bordalo, Gennaioli, La Porta, and Shleifer (2024a) show that changes in LTG are correlated with the difference between current earnings and cyclically-adjusted (i.e., multi-year smoothed) earnings five years earlier. Relatedly, BGLS show that earnings news have substantial explanatory power in a regression of their expectations-based return (EBR), which is driven to a large extent by changes in LTG, on returns and earnings news. However, in models in which risk premia move counter-cyclically with the business cycle, risk premia can also covary with multi-year earnings growth. Then, if low risk premia lead analysts to infer high LTG, changes in LTG end up being positively related to this multi-year earnings growth variable. Moreover, as returns also reflect short-term cash flow news in addition to changes in risk premia, returns are not perfectly correlated with LTG changes. Therefore, returns and earnings growth jointly explain changes in LTG under this alternative story, too.

Second, Bordalo, Gennaioli, La Porta, and Shleifer (2024a) show that LTG predicts returns controlling for valuation ratios. BGLS run a similar test in the cross-section. However, as long the valuation ratio is not a perfect signal for the risk premium—for example, because there is some predictable variation in short-term expected cash flow growth that contributes to valuation ratio variance—LTG will help to predict returns even controlling for the valuation ratio under the alternative risk premium story, too.

Third, BGLS show that variables that predict returns also predict analyst forecast errors. This is again observationally equivalent to what would happen under the alternative inferring-from-prices story. For example, if analysts report optimistic expectations when risk premia are low, they will have predictably negative forecast errors.

Thus, while it may be difficult to come up with a specific, empirically measurable and theoretically grounded risk premium variable that produces all of these facts, it is also true that these tests cannot unequivocally reject the concern that analysts' reported LTG inadvertendly captures risk premia. Very likely, this question cannot be settled with data on valuation ratios, returns, analyst consensus forecasts, and earnings alone.

One avenue for getting more clarity is to look for exogenous shocks to prices that are plausibly free of cash flow news and then check whether analysts revise earnings expectations in response to these shocks. Chaudhry (2023) follows this approach and exploits flow- and benchmarking-induced price changes. He finds a substantial effect of non-cash-flow-news price shocks on short- and long-term growth expectations of analysts. This evidence is not necessarily inconsistent with an interpretation of analyst forecasts as proxies for investor cash flow expectations. After all, models exist in which cash flow expectations are driven by extrapolation from past returns (see, e.g., Jin and Sui 2022). But the evidence also does not reject the possibility that analysts could end up reporting earnings growth forecasts that embody risk premia.

Looking beyond just consensus earnings forecasts, the outputs produced by analysts have a number of properties that seem difficult to reconcile with the idea that analyst forecasts are largely inferred from market prices.

First, there is substantial disagreement about LTG between analysts. And this disagreement has shown to be useful in a subjective beliefs asset pricing approach (see, e.g., Diether, Malloy, and Scherbina 2002; Hong and Sraer 2016). If analysts mainly looked at market prices to come up with their reported expectations, it would be difficult to explain why they disagree so much. After all, they are all seeing the same market prices.

Relatedly, an examination of the different moving parts of an analysts valuation model could help shed light on how they come up with their reported expectations. If analysts disagree, or use different assumptions on some dimensions of their analysis, but they are trying to get close to matching the market price with their valuation, then they have to disagree on the remaining dimensions in a negatively correlated way. For example, if they disagree on WACC because they use different betas, or different risk premium assumptions, then analysts that use higher WACC should have lower growth expectations. Focusing on analysts who use discounted cash flow analysis, Décaire, Sosyura, and Wittry (2024) do not find evidence for such a negative correlation between WACC and growth expectations.

Finally, it may be useful to consider earnings forecasts jointly with target prices and stock recommendations issued by analysts. Existing work already offers some results that cast doubt on the expectations-inferred-from-prices story. For example, Bradshaw (2004) finds that LTG is an important explanatory factor for analysts' stock recommendations. It wouldn't really make sense that analysts back out LTG from current prices, but then, at the same time, say, optimism about LTG leads them to suggest that the current price undervalues the firm. Jung, Shane, and Yang (2012) find that announcements of analysts' stock recommendations generate stronger price reactions when they are accompanied by a long-term growth forecast. If LTG was simply backed out from market prices, there would be no reason why investors should regard a stock recommendation as more informative if it's published with an LTG estimate.

None of the above is bullet-proof refutation of the view that analysts report expectations inferred from prices. But the examples illustrate that more research that goes beyond the study of consensus earnings forecasts may help to get a better handle on how analysts come up with their estimates.

IV.B. Do analysts calculate target prices based on trailing multiples without forward-looking reasoning?

Ben-David and Chinco (2024) (BC) raise a point that is somewhat related to the idea that analysts look for guidance from market prices to come up with the outputs in their reports. BC argue that analysts mechanically apply trailing price-earnings ratios to near-term earnings forecasts to come up with a target price.

Looking at the Gordon growth formula shows that, in principle, there is nothing wrong with a multiples approach to valuation,

$$\hat{P}_t = \widetilde{\mathbb{E}}_t[E_{t+1}] \left(\frac{1}{r-g} \right), \tag{41}$$

if the multiple $1/(r-g)$ that is applied to the earnings forecast $\widetilde{\mathbb{E}}_t[E_{t+1}]$ reflects forward-looking assessment of $r - g$. But BC assert that this multiple is not based on forward-looking reasoning but mechanically based on trailing price-earnings ratios:

$$\hat{P}_t = \widetilde{\mathbb{E}}_t[E_{t+1}] \times \text{trailing } \frac{P}{E}. \tag{42}$$

Proving that this is a *mechanical* reliance on trailing price-earnings ratios, and that forward-looking reasoning is absent, is difficult, though. A strong positive correlation between the multiple implicitly or explicitly used by the analyst and trailing price-to-earnings is not necessarily evidence of absence of forward-looking reasoning. If investors and analysts to some extent have similar views about future cash flows and risks, then the multiple that a forward-looking analyst would apply, based on their expectations of future growth and risks, would be highly correlated with recent multiples seen in market prices.

In this regard, it seems questionable whether BC's finding that a regression of analysts' target-price implied price-earnings multiple on trailing price-earnings ratios produces a slope coefficient of 0.58 and R^2 of 54% is really evidence of mechanical application of trailing price-earnings ratios. This is a far from perfect correlation. It seems plausible that this level of correlation could well be explained simply by the fact that investors and analysts to some extent have similar expectations.

Rather than trying to prove mechanical reliance on past price-earnings ratios, it may be more interesting to explore the variation that is not explained by trailing price-earnings ratios. Existing literature gives some hints about factors that may play a role. Bradshaw (2002) finds that the PEG ratio based valuation,

$$\hat{P}_{PEG,t} = \frac{\widetilde{\mathbb{E}}_t[E_{t+1}]}{LTG_t} \tag{43}$$

helps explain the wedge between target prices and current market prices. (In light of the earlier discussion above whether analysts infer LTG from market prices, the fact that LTG helps explain this wedge is also another piece of evidence that analysts use LTG in ways that does not fit with the notion that LTG is just backed out from market prices.) Yin, Peasnell, and Hunt III (2018) find that industry-adjusted forward price-earnings ratios implied by target prices correlate with industry-adjusted LTG.

V. Conclusion

Excess volatility alive and well. AHP's cash flow dark matter explanation does not find empirical support. In fact, the adjustment to the price-consumption ratio based on the dividend-consumption ratio that they devised actually strongly predicts stock market index returns rather than future cash flows—which strengthens the case for excess volatility.

This then brings up an interesting open question: Why does this adjustment produce a stronger return predictor? I have sketched an explanation based on entry of low dividend-paying, high dividend-growth firms into the stock market. But perhaps alternative mechanisms are at work. This seems like an interesting area for further research.

Much disagreement in the literature about the economic magnitude of excess volatility seems to stem from the use of variance decompositions of different valuation ratios. Given that predictable variation in future returns reveal transitory components of stock prices, it is better to focus on the magnitude of the variance of this transitory component directly, rather than looking at the variance share in a variance decomposition of a valuation ratio. The magnitude of the variance of the transitory component only depends on the conditioning information, while results from variance decompositions of valuation ratios can be very sensitive to the choice of denominator variable.

That variance decompositions of valuation ratios can be sensitive to the specification of the valuation ratio also applies to variance decompositions that employ subjective expectations data. The surprisingly dominant role that short-term earnings growth expectations seem to play in a variance decomposition of the price-earnings ratio is strongly diminished when earnings are measured excluding special items that add a lot of noise to earnings.

One weak spot of the literature on subjective cash flow expectations is that, due to a lack of alternative data sources, so much research in this area is based on equity analyst earnings forecasts. Whether they are a good proxy for investor cash flow expectations is not yet well understood. There are still a number of unresolved questions about how analysts come up with their forecasts. Among the more important question is whether analysts infer expectations from market prices. If analysts are doing this, their reported cash flow expectations could end up picking up risk premium variation. I have argued that existing attempts to answer this question do not yet provide sufficient clarify. More can be learned by studying the factors that generate disagreement between analysts and from a more detailed analysis of how analysts come up with their forecasts. Data on cash flow expectations of investors rather than analysts would also help, but such data may not be available at this point, especially not over longer time frames.

References

Atkeson, Andrew, Jonathan Heathcote, and Fabrizio Perri, 2024, "There is No Excess Volatility Puzzle," Working paper, National Bureau of Economic Research.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review* 1, 183–221.

Bansal, Ravi, and Amir Yaron, 2004, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance* 59, 1481–1509.

Bansal, Ravi, and Amir Yaron, 2006, "The Asset Pricing-Macro Nexus and Return-Cash Flow Predictability," Working paper, Duke University.

Barsky, Robert B, and J Bradford De Long, 1993, "Why Does the Stock Market Fluctuate?," *Quarterly Journal of Economics* 108, 291–311.

Beeler, Jason, and John Y. Campbell, 2012, "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review* 1, 141–182.

Ben-David, Itzhak, and Alex Chinco, 2024, "Expected EPS × Trailing P/E," Working paper, The Ohio State University.

Binsbergen, Jules H van, and Ralph SJ Koijen, 2010, "Predictive Regressions: A Present-Value Approach," *Journal of Finance* 65, 1439–1471.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2024a, "Belief Overreaction and Stock Market Puzzles," *Journal of Political Economy* 132, 1450–1484.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2024b, "Finance Without (Exotic) Risk," Working paper, Harvard University.

Bradshaw, Mark T, 2002, "The Use of Target Prices to Justify Sell-Side Analysts' Stock Recommendations," *Accounting Horizons* 16, 27–41.

Bradshaw, Mark T, 2004, "How do Analysts Use Their Earnings Forecasts in Generating Stock Recommendations?," *The Accounting Review* 79, 25–50.

Campbell, John Y, 1991, "A Variance Decomposition for Stock Returns," *Economic Journal* 101, 157–179.

Campbell, John Y., and Robert J. Shiller, 1988, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies* 1, 195–228.

Chaudhry, Aditya, 2023, "Prices Impact Analyst Cash Flow Expectations," Working paper, Ohio State University.

Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2003, "The Value Spread," *Journal of Finance* 58, 609–641.

Cohen, Randolph B, Christopher Polk, and Tuomo Vuolteenaho, 2009, "The Price is (Almost) Right," *Journal of Finance* 64, 2739–2782.

De La O, Ricardo, Xiao Han, and Sean Myers, 2023, "The Return of Return Dominance: Decomposing the Cross-section of Prices," Working paper, Wharton School.

De La O, Ricardo, and Sean Myers, 2021, "Subjective Cash Flows and Discount Rates," *Journal of Finance* 76, 1339–1387.

Décaire, Paul H., Denis Sosyura, and Michael D. Wittry, 2024, "Resolving Estimation Ambiguity," Working paper, Arizona State University.

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina, 2002, "Differences in Opinion and the Cross Section of Stock Returns," *Journal of Finance* 57, 2113–2140.

Hall, Robert E, 2001, "Struggling to Understand the Stock Market," *American Economic Review* 91, 1–11.

Hong, Harrison, and David A Sraer, 2016, "Speculative Betas," *Journal of Finance* 71, 2095–2144.

Jank, Stephan, 2015, "Changes in the Composition of Publicly Traded Firms: Implications for the Dividend-Price Ratio and Return Predictability," *Management Science* 61, 1362–1377.

Jin, Lawrence J, and Pengfei Sui, 2022, "Asset Pricing with Return Extrapolation," *Journal of Financial Economics* 145, 273–295.

Jung, Boochun, Philip B Shane, and Yanhua Sunny Yang, 2012, "Do Financial Analysts' Long-Term Growth Forecasts Matter? Evidence from Stock Recommendations and Career Outcomes," *Journal of Accounting and Economics* 53, 55–76.

Larrain, Borja, and Motohiro Yogo, 2008, "Does Firm Value Move Too Much to be Justified by Subsequent Changes in Cash Flow?," *Journal of Financial Economics* 87, 200–226.

Maio, Paulo, and Danielle Xu, 2020, "Cash-Flow or Return Predictability at Long Horizons? The Case of Earnings Yield," *Journal of Empirical Finance* 59, 172–192.

McGrattan, Ellen R, and Edward C Prescott, 2001, "Is the Stock Market Overvalued?," Working paper, National Bureau of Economic Research.

McGrattan, Ellen R, and Edward C Prescott, 2004, "The 1929 Stock Market: Irving Fisher Was Right," *International Economic Review* 45, 991–1009.

Nagel, Stefan, and Zhengyang Xu, 2022, "Asset Pricing with Fading Memory," *Review of Financial Studies* 35, 2190–2245.

Nagel, Stefan, and Zhengyang Xu, 2023, "Dynamics of Subjective Risk Premia," *Journal of Financial Economics* 150, 103713.

Pruitt, Seth, 2023, "Dogs and Cats Living Together: A Defense of Cash-Flow Predictability," Working paper, Arizona State University.

Shiller, Robert J., 1981, "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review* 71, 321 – 436.

Vuolteenaho, Tuomo, 2000, "Understanding the Aggregate Book-to-Market Ratio and Its Implications to Current Equity-Premium Expectations," Working paper, Harvard University.

Yin, Yuan, Ken Peasnell, and Herbert G Hunt III, 2018, “How Do Sell-Side Analysts Obtain Price-Earnings Multiples to Value Firms?,” *Accounting and Business Research* 48, 108–135.

Appendix

A. Aggregate present value identity

The gross return for investor who holds aggregate market and gets aggregate total payout $Y_{t+1}Q_t$ that includes net repurchases and cash dividends

$$R_{t+1} = \frac{P_{t+1}Q_{t+1} + Y_{t+1}Q_t}{P_tQ_t}. \tag{A.1}$$

This return is the same as the return of an investor who keeps holding shares held at time t, so

$$\frac{P_{t+1} + D_{t+1}}{P_t} = \frac{P_{t+1}Q_{t+1} + Y_{t+1}Q_t}{P_tQ_t}, \tag{A.2}$$

and hence

$$P_{t+1} + D_{t+1} = P_{t+1}(Q_{t+1}/Q_t) + Y_{t+1}. \tag{A.3}$$

Rearranged,

$$Y_{t+1} = P_{t+1}(1 + D_{t+1}/P_{t+1} - Q_{t+1}/Q_t), \tag{A.4}$$

and in logs

$$y_{t+1} = p_{t+1} + \log\left[1 + \exp(d_{t+1} - p_{t+1}) - \exp(\Delta q_{t+1})\right]. \tag{A.5}$$

A first-order Taylor approximation around $d_{t+1} - p_{t+1} = \bar{dp}$ and $\Delta q_{t+1} = 0$ yields

$$\begin{aligned} y_{t+1} &= p_{t+1} + \exp(\bar{dp}) + [d_{t+1} - p_{t+1} - \exp(\bar{dp})] - \frac{1}{\exp(\bar{dp})}\Delta q_{t+1} \\ &= d_{t+1} - \frac{\rho}{1-\rho}\Delta q_{t+1}, \end{aligned} \tag{A.6}$$

so

$$\Delta y_{t+1} = \Delta d_{t+1} - \frac{\rho}{1-\rho}\Delta^2 q_{t+1}, \tag{A.7}$$

where Δ^2 denotes a twice differenced quantity. Then, using (A.6) we can get from

$$p_t - d_t = \frac{\kappa}{1-\rho} + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j} \tag{A.8}$$

to a present value identity for the price to per-share payout ratio,

$$p_t - y_t = \frac{\kappa}{1-\rho} + \frac{\rho}{1-\rho}\Delta q_t + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{A.9}$$

Aggregate payouts are defined as $Y_{t+1}Q_t$ so aggregate log payout growth is $\Delta y_{t+1} + \Delta q_t$ and the aggregate price-to-payout ratio $P_tQ_t/(Y_tQ_{t-1})$ in logs is $p_t - y_t + \Delta q_t$. The aggregate PV identity, log-linearized around same point for $y_t - p_t$ that we log-linearized the per-share return around, so that we get the same ρ, is

$$p_t - y_t + \Delta q_t = \frac{\kappa}{1-\rho} + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j (\Delta y_{t+1+j} + \Delta q_{t+j}) - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{A.10}$$

To relate this to the per-share identity (A.9), use (A.7),

$$\begin{aligned}\sum_{j=0}^{\infty} \rho^j (\Delta y_{t+1+j} + \Delta q_{t+j}) &= \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} + \Delta q_{t+1} - \frac{\rho}{1-\rho}\Delta^2 q_{t+1}) \\ &= \Delta d_{t+1} + \Delta q_t - \frac{\rho}{1-\rho}\Delta q_{t+1} + \frac{\rho}{1-\rho}\Delta q_t \\ &\quad + \rho \Delta d_{t+2} + \rho \Delta q_{t+1} - \frac{\rho^2}{1-\rho}\Delta q_{t+2} + \frac{\rho^2}{1-\rho}\Delta q_{t+1} \\ &\quad + \ldots \end{aligned} \tag{A.11}$$

Summing the Δq from the same time period, and noting that $-\rho + \rho(1-\rho) + \rho^2 = 0$, all the Δq terms after time t cancel, and we get

$$\sum_{j=0}^{\infty} \rho^j (\Delta y_{t+1+j} + \Delta q_{t+j}) = \frac{1}{1-\rho}\Delta q_t + \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}. \tag{A.12}$$

Therefore,

$$p_t - y_t + \Delta q_t = \frac{\kappa}{1-\rho} + \frac{1}{1-\rho}\Delta q_t + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}, \tag{A.13}$$

or

$$p_t - y_t = \frac{\kappa}{1-\rho} + \frac{\rho}{1-\rho}\Delta q_t + \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - \mathbb{E}_t \sum_{j=0}^{\infty} \rho^j r_{t+1+j}. \tag{A.14}$$

which of course gets us back the per-share PV identity (A.9).

B. Surprise repurchases example

Suppose at t all future Δq_{t+j} for $j \geq 1$ were expected to be zero, but then a surprise repurchase $\Delta q_{t+1} < 0$ takes place, while all future Δq_{t+j} for $j \geq 2$ are still expected to be zero. Assume the repurchase is just substitution for cash dividend payouts, so it has no effect on aggregate payouts. So nothing changes in (34) due to the surprise repurchase. Per-share payouts however change, as they will be calculated in $t+1$ based on the now lower number of shares Q_{t+1}:

$$\Delta E_{t+1} \Delta y_{t+2} = -\Delta\, \mathbb{E}_{t+1}\, \Delta q_{t+1}. \tag{A.15}$$

Expectations of payouts in other periods remain unchanged. Cash dividends change in two periods: in period $t+1$ due to the expense for the repurchase and in period $t+2$ due to the now changed shares outstanding. Using (A.7), we get

$$\Delta\,\mathbb{E}_{t+1}\,\Delta d_{t+1} = \frac{\rho}{1-\rho}\Delta\,\mathbb{E}_{t+1}\,\Delta q_{t+1} \tag{A.16}$$

and

$$\begin{aligned}\Delta\,\mathbb{E}_{t+1}\,\Delta d_{t+2} &= \Delta\,\mathbb{E}_{t+1}\,\Delta y_{t+2} + \frac{\rho}{1-\rho}\Delta\,\mathbb{E}_{t+1}\,\Delta q_{t+1} \\ &= -\Delta\,\mathbb{E}_{t+1}\,\Delta q_{t+1} + \frac{\rho}{1-\rho}\Delta\,\mathbb{E}_{t+1}\,\Delta q_{t+1} \\ &= -\frac{1}{1-\rho}\Delta\,\mathbb{E}_{t+1}\,\Delta q_{t+1}. \end{aligned} \tag{A.17}$$

However, the total effect of the two expected dividend growth changes on the PV identity is zero because

$$\frac{\rho}{1-\rho}\Delta\,\mathbb{E}_{t+1}\,\Delta q_{t+1} - \rho\frac{1}{1-\rho}\Delta\,\mathbb{E}_{t+1}\,\Delta q_{t+1} = 0. \tag{A.18}$$