Crimes Against Campbell-Shiller

Itzhak Ben-David* and Alex Chinco†

February 12, 2026

Abstract

The Campbell and Shiller (1988a) log-linear approximation is widely treated as an accounting identity that always holds in realized data, in expectation, and under arbitrary subjective beliefs. None of these claims is true. The formula breaks down when firms repurchase shares. It does not hold in expectation unless investors enforce present-value logic with extreme precision. It does not hold under arbitrary subjective beliefs. For the formula to hold under biased beliefs, the expectation errors must satisfy a precise adding-up constraint. If investors understood their errors well enough to impose such a constraint, they would simply correct them. Popular belief specifications, like diagnostic expectations and extrapolation, can be made compatible with the formula, but only by embedding them in much more elaborate models. Before using the formula, researchers should ask whether the strong assumptions are reasonable in their setting.

Keywords: Campbell-Shiller, Present-Value Relationships, Subjective Beliefs, Behavioral Finance, Return Decomposition

*The Ohio State University and NBER. ben-david.1@osu.edu
†Michigan State University. alexchinco@gmail.com

Introduction

Campbell and Shiller (1988a) is one of the most influential papers in empirical asset pricing. The paper derives a log-linear approximation to the definition of stock returns and uses this formula to decompose the price-to-dividend (PD) ratio into expectations of future dividend growth and future returns

$$\log \text{PD} \;\approx\; \sum_{t=1}^{\infty} \rho^{t-1} \cdot \left\{ \mathbb{E}[\Delta \log \text{Div}_t] \;-\; \mathbb{E}[\text{R}_t] \right\} \tag{1}$$

where $\rho = 1/(1 + e^{-\overline{\log \text{PD}}}) \approx 0.96$ and we have suppressed a constant term.

The original analysis is sharp, self-aware, and careful about its own assumptions. This paper is not a critique of Campbell-Shiller. It is an accounting of how many qualifications are needed to apply the formula beyond the setting it was designed for. Each section identifies a separate assumption that is often taken for granted, and each one narrows the applicable domain.

There are many qualifications. First, Campbell-Shiller would only be an accounting identity in a world where companies never issue or repurchase shares. We do not live on that planet. The point about repurchases is especially problematic because this is how the majority of large US companies return cash (Boudoukh, Michaely, Richardson, and Roberts, 2007; Ben-David and Chinco, 2026). A company that switched from paying dividends to returning all cash exclusively through buybacks would have an undefined dividend growth rate

$$\Delta \log \text{Div}_{t+1} \;=\; \log(\$0/\text{sh}) - \log(\text{Div}_t) \;=\; -\infty \tag{2}$$

Second, suppose that we could safely ignore new issuance and buybacks. In that case, Campbell-Shiller would hold in realized data, but it would not necessarily hold in expectation. Many equity investors do not apply present-value logic. For example, sell-side equity analysts, the professional forecasters whose price targets move billions of dollars of trading volume, typically value stocks using backward-looking trailing multiples rather than forward-looking discount rates (Ben-David and Chinco, 2025).

Even investors who do apply present-value logic must know the relevant discount rates and growth rates with an unrealistic level of precision. "Typical estimates put the equity premium between 4% and 8%." (Cochrane, 2017) By contrast, a back-of-the-envelope calculation shows that an investor would need to know the cap rate of the typical S&P 500 company to within $DY^2 \approx 0.04\%$. If not, parameter uncertainty would make it impossible to keep track of the very next dividend. A present-value relationship cannot be considered an accounting "identity" in a setting where noise makes it difficult to count payments.

Third, suppose new issuance and repurchases can be ignored. Further suppose that investors fastidiously enforce present-value relationships with sufficient precision. If their subjective beliefs are biased, it would still not be safe to simply plug those beliefs into the Campbell-Shiller formula. The formula holds if you replace the rational expectation operator, $\mathbb{E}[\cdot]$, with any other expectation operator, $\widetilde{\mathbb{E}}[\cdot]$, defined using a coherent probability measure. The reason is that expectations under a well-defined probability measure are integrals, and integration is linear. Linearity is what preserves equality.

But most of the belief specifications studied in behavioral finance are not defined using a coherent probability measure. For the decomposition to hold under subjective beliefs, the distortions to return expectations, dividend-growth expectations, and price-dividend-ratio expectations must satisfy an adding-up condition: $\delta_R = \delta_{Div} - \rho \cdot \delta_{PD}$. This is a knife-edge condition: a single linear equation relating three free parameters. It holds automatically under a well-posed probability measure, but it is generically violated when the distortions are specified independently on a variable-by-variable basis.

Diagnostic expectations (Bordalo, Gennaioli, and Shleifer, 2018), extrapolative expectations (Barberis, Greenwood, Jin, and Shleifer, 2015), and natural expectations (Fuster, Laibson, and Mendel, 2010) all fail this test when applied variable by variable. That does not mean these models are wrong. It means that using them with Campbell-Shiller requires a lot of extra baggage. You cannot just say investors have biased return expectations and call it a day. You must provide enough additional structure so that the agent's subjective beliefs are consistent with a coherent distorted probability distribution.

What's more, this adding-up condition is a strange thing to require of biased agents. It says that if investors are too optimistic about dividend growth, they must be correspondingly pessimistic about future valuations or optimistic about returns, in exactly the right proportions. If investors understood the structure of their own errors well enough to enforce such a constraint, they would simply correct those errors and stop making them.

When it snows in Michigan, the elementary-school car pick-up line tends to run slow. Some parents consistently fail to appreciate this fact, and their unfortunate children always arrive late to basketball practice on snowy days. At this point, no one is surprised. It would be surprising if these thick-skulled parents somehow overestimated the driving time to basketball practice on exactly the days when the car pick-up line was slowest. That is the sort of thing it takes for Campbell-Shiller to hold under biased subjective beliefs.

Paper Outline. This paper proceeds as follows. Section 1 reviews the original Campbell-Shiller derivation, walking through the log-linear approximation step by step and clarifying what it is an approximation to. Section 2 shows how repurchases and new issuance break or distort the formula when it is applied to realized data. Section 3 identifies two additional assumptions required for the formula to hold in expectation: that investors price stocks using present-value logic rather than backward-looking multiples, and that they know the relevant parameters with extreme precision. Section 4 derives an error-aggregation condition that subjective beliefs must satisfy for the decomposition to hold, shows that only beliefs derived from a coherent probability measure automatically satisfy it, and examines how three prominent papers in the literature have navigated this constraint.

Related Work. This paper touches on several literatures. The starting point is the Campbell-Shiller log-linear present-value framework itself. Campbell and Shiller (1988a) and Campbell and Shiller (1988b) derive the approximation and apply it to aggregate US stock-market data. Campbell (1991) uses the framework to decompose unexpected stock returns into "cash-flow news" and "discount-rate news," finding that most of the variation in aggregate returns comes from

revisions to expected future discount rates rather than revisions to expected future dividends. Vuolteenaho (2002) applies the framework at the firm level, documenting that firm-level returns are primarily driven by cash-flow news, the opposite of the aggregate result. The textbook treatment in Campbell (2017) presents the decomposition as a workhorse identity and is the source of many of the claims we scrutinize.

A large empirical literature uses Campbell-Shiller to understand the sources of variation in expected returns. Cochrane (2011) argues that the dividend-price ratio forecasts returns rather than dividend growth, implying that virtually all variation in valuations reflects changing discount rates. Lettau, Ludvigson, and Manoel (2018) revisit this finding using improved econometric methods. Our paper does not take a position on the discount-rate-news versus cash-flow-news debate. Instead, we argue that the framework within which the debate takes place rests on assumptions that are routinely violated.

Other papers have studied the accuracy of the approximation itself. Engsted, Pedersen, and Tanggaard (2012) derive an upper bound on the mean approximation error and show that the log-linear approximation remains accurate even in the presence of explosive rational bubbles. Cho, Kremens, Lee, and Polk (2024) propose an alternative present-value formula that expresses cash-flow fundamentals in terms of ROE and investment value added rather than dividends, generating smaller approximation errors in practice.

The issue of share repurchases has been recognized before, most notably by Boudoukh et al. (2007), who propose the "payout yield" (dividends plus repurchases, scaled by price) as a more comprehensive measure of cash returned to shareholders. Our contribution is to show that the problem is not merely one of mismeasurement. Even the payout yield does not save the Campbell-Shiller decomposition, because the log-linear approximation is derived from a per-share return definition that implicitly assumes every share survives. Repurchases violate this assumption by terminating shares.

Our discussion of present-value logic and backward-looking pricing rules connects to the growing literature on demand-system asset pricing (Koijen and Yogo, 2019; Gabaix and Koijen, 2024). In these models, investors do not

necessarily price stocks by discounting future cash flows; instead, prices emerge from the interaction of demand curves and market clearing. The finding in Ben-David and Chinco (2025) that sell-side analysts use trailing multiples rather than discounted cash flow models is consistent with this view. Our paper argues that if the marginal investor does not apply present-value logic, then the Campbell-Shiller decomposition is asking the wrong question.

The most closely related theoretical contribution concerns the behavior of Campbell-Shiller under subjective beliefs. Several influential papers in the belief-based asset-pricing literature assert that the decomposition holds under arbitrary subjective expectations. De la O and Myers (2021) use the framework to measure the contribution of subjective cash-flow and discount-rate expectations to the variation in the PD ratio. Bordalo, Gennaioli, La Porta, and Shleifer (2019, 2024) use it to study the role of diagnostic expectations in driving credit cycles and stock-market fluctuations. In each case, the identity is imposed rather than tested. We derive the conditions under which the identity actually holds and show that meeting those conditions requires either a full equilibrium model or a circularity in which one input is backed out as a residual.

Jin and Sui (2022) pursue a different approach: they embed extrapolative return beliefs inside an equilibrium model with Epstein-Zin preferences, and the Campbell-Shiller decomposition holds because all beliefs derive from a single (if incorrect) probability distribution.

There is a broader literature on expectations in asset pricing. Nagel and Xu (2022) provides a comprehensive survey. Bordalo et al. (2018) introduce diagnostic expectations, a model in which agents overweight states that have become relatively more likely. Greenwood and Shleifer (2014) and Barberis et al. (2015) study models in which agents extrapolate from recent returns. Fuster et al. (2010) propose natural expectations, where agents use simpler forecasting models than the true data-generating process. Adam, Marcet, and Nicolini (2016) and Adam, Marcet, and Beutel (2017) study asset pricing under subjective beliefs with internal rationality. Our paper does not argue against any of these models. Rather, we show that using them with the Campbell-Shiller decomposition requires more infrastructure than is acknowledged.

1 Original Derivation

This section has two goals. Subsection 1.1 walks through the Campbell-Shiller derivation step-by-step. The derivation starts from the definition of a stock return, applies a first-order Taylor approximation, iterates forward, and then takes expectations under a transversality condition. Subsection 1.2 clarifies what all this extra work buys you. The Gordon growth model assumes constant discount rates and constant dividend growth, producing a clean but rigid closed-form expression. Campbell and Shiller's contribution is to linearize the time-varying version of this relationship.

1.1 Log-Linear Approximation

We now walk through the derivation step by step. Start from the definition of a one-period gross stock return. If you buy a share at time t for Price_t, collect a dividend of Div_{t+1}, and sell at Price_{t+1}, then your gross return is

$$1 + \text{R}_{t+1} = \frac{\text{Price}_{t+1} + \text{Div}_{t+1}}{\text{Price}_t} \tag{3}$$

This is a definition, not an approximation. It holds for any stock in any time period, regardless of what investors believe or how they form expectations. Note that shares outstanding does not appear in this formula. The return is defined for a single share that is held from t to $t + 1$. We will return to the consequences of this omission in Section 2.

Multiply the right-hand side by $\frac{\text{Div}_{t+1}/\text{Div}_{t+1}}{\text{Div}_t/\text{Div}_t}$ to express the gross return in terms of the price-to-dividend ratio, $\text{PD}_t = \frac{\text{Price}_t}{\text{Div}_t}$,

$$1 + \text{R}_{t+1} = \left(\frac{\text{Price}_{t+1} + \text{Div}_{t+1}}{\text{Price}_t}\right) \times \left(\frac{\text{Div}_{t+1}/\text{Div}_{t+1}}{\text{Div}_t/\text{Div}_t}\right) \tag{4a}$$

$$= \left(\frac{\text{Price}_{t+1}/\text{Div}_{t+1} + 1}{\text{Price}_t/\text{Div}_t}\right) \times \left(\frac{\text{Div}_{t+1}}{\text{Div}_t}\right) \tag{4b}$$

$$= \left(\frac{\text{PD}_{t+1} + 1}{\text{PD}_t}\right) \times \left(\frac{\text{Div}_{t+1}}{\text{Div}_t}\right) \tag{4c}$$

Taking the natural logarithm of both sides and writing R_{t+1} for the log gross return, $\log PD_t$ for the log price-dividend ratio, and $\Delta \log Div_{t+1}$ for log dividend growth, we get

$$R_{t+1} = \log(PD_{t+1} + 1) - \log PD_t + \Delta \log Div_{t+1} \tag{5}$$

This is still exact. The only nonlinearity is the $\log(PD_{t+1} + 1)$ term, which is a nonlinear function of the future PD ratio. Everything else is already log-linear.

The key step in the Campbell-Shiller derivation is to replace $\log(PD_{t+1} + 1)$ with a first-order Taylor approximation around $\overline{\log PD} = \mathbb{E}[\log PD_t]$, the mean log PD ratio. If we define $\overline{PD} = e^{\overline{\log PD}}$, then a first-order expansion gives

$$\log(PD_{t+1} + 1) \approx \ln(\overline{PD} + 1) + \left(\frac{\overline{PD}}{\overline{PD} + 1}\right) \times (\log PD_{t+1} - \overline{\log PD}) \tag{6}$$

Now define the two linearization constants

$$\rho = \frac{\overline{PD}}{\overline{PD} + 1} \qquad \kappa = \log(\overline{PD} + 1) - \rho \cdot \overline{\log PD} \tag{7}$$

For the aggregate US stock market, typical calibrations put $\rho \approx 0.96$ at an annual frequency (Campbell, 2017).

Substituting Equations (6) and (7) into Equation (5) gives the Campbell-Shiller one-period approximation

$$R_{t+1} \approx \kappa + \Delta \log Div_{t+1} + \rho \cdot \log PD_{t+1} - \log PD_t \tag{8}$$

Or, equivalently, if we solve for the current log PD ratio, we get

$$\log PD_t \approx \kappa + \Delta \log Div_{t+1} - R_{t+1} + \rho \cdot \log PD_{t+1} \tag{9}$$

Today's log PD ratio is approximately equal to next year's dividend growth, minus next year's return, plus a discounted version of next year's log PD ratio. The power of Equation (9) is that it can be applied recursively.

Replace $\log \mathrm{PD}_{t+1}$ on the right-hand side with its own one-period expansion: $\log \mathrm{PD}_{t+1} \approx \kappa + \Delta \log \mathrm{Div}_{t+2} - \mathrm{R}_{t+2} + \rho \cdot \log \mathrm{PD}_{t+2}$. Substituting yields

$$\begin{aligned} \log \mathrm{PD}_t \approx \kappa \cdot (1+\rho) + \{\Delta \log \mathrm{Div}_{t+1} - \mathrm{R}_{t+1}\} \\ + \ \rho \cdot \{\Delta \log \mathrm{Div}_{t+2} - \mathrm{R}_{t+2}\} + \rho^2 \cdot \log \mathrm{PD}_{t+2} \end{aligned} \tag{10}$$

Repeating this substitution H times produces something that looks like Equation (1) from the introduction with an extra term on the end

$$\log \mathrm{PD}_t \approx \sum_{h=1}^{H} \rho^{h-1} \cdot \{\Delta \log \mathrm{Div}_{t+h} - \mathrm{R}_{t+h}\} + \rho^H \cdot \log \mathrm{PD}_{t+H} \tag{11}$$

Note that we are not writing the leading $\kappa \cdot \left(\frac{1-\rho^H}{1-\rho}\right)$ term for simplicity's sake. This is an approximation in realized values. It holds for any stock in any sample, up to the quality of the first-order Taylor approximation. No expectations have been taken. No assumptions about beliefs have been made.

To see what happened to the $\rho^H \cdot \log \mathrm{PD}_{t+H}$ term, let's take conditional expectations of both sides of Equation (11). Since $\log \mathrm{PD}_t$ is known at time t, the left-hand side does not change. On the right-hand side, the objectively correct conditional expectation passes through the sum

$$\log \mathrm{PD}_t \approx \sum_{h=1}^{H} \rho^{h-1} \cdot \{\mathbb{E}_t[\Delta \log \mathrm{Div}_{t+h}] - \mathbb{E}_t[\mathrm{R}_{t+h}]\} + \rho^H \cdot \mathbb{E}_t[\log \mathrm{PD}_{t+H}] \tag{12}$$

We need the terminal term, $\rho^H \cdot \mathbb{E}_t[\log \mathrm{PD}_{t+H}]$, to vanish as $H \to \infty$. This is known as a transversality condition. It says that the expected log PD ratio does not grow so fast that its discounted value explodes. Under this assumption, we arrive at an expected-value version of Campbell-Shiller

$$\log \mathrm{PD}_t \approx \sum_{h=1}^{\infty} \rho^{h-1} \cdot \{\mathbb{E}_t[\Delta \log \mathrm{Div}_{t+h}] - \mathbb{E}_t[\mathrm{R}_{t+h}]\} \tag{13}$$

It says that a stock's log PD ratio is approximately equal to a constant plus the discounted sum of expected future dividend growth net of returns.

1.2 What Is The Upside?

The classic Gordon model says to price stocks using the following formula

$$\text{Price}_t = \frac{\mathbb{E}_t[\text{Div}_{t+1}]}{1+\text{r}} + \frac{\mathbb{E}_t[\text{Price}_{t+1}]}{1+\text{r}} \tag{14a}$$

$$= \frac{\mathbb{E}_t[\text{Div}_{t+1}]}{1+\text{r}} + \sum_{h=2}^{\infty} \frac{\mathbb{E}_t[\text{Div}_{t+h}]}{(1+\text{r})^h} \tag{14b}$$

$$= \mathbb{E}_t[\text{Div}_{t+1}] \times \left(\frac{1}{\text{r}-\text{g}}\right) \tag{14c}$$

Equation (14a) says a stock's current price reflects its expected discounted dividend and resale price next year. To arrive at Equation (14b), you iterate forward, replacing $\mathbb{E}_t[\text{Price}_{t+h}]$ with $\frac{\mathbb{E}_t[\text{Div}_{t+(h+1)}]}{1+\text{r}} + \frac{\mathbb{E}_t[\text{Price}_{t+(h+1)}]}{1+\text{r}}$ for $h = 1, 2, 3, \ldots$

Gordon's insight was that, with constant expected dividend growth, $(1+\text{g})^h = \frac{\mathbb{E}_t[\text{Div}_{t+h}]}{\text{Div}_t}$, it is possible to replace an infinite sum with a forward multiple

$$\left(\frac{1}{\text{r}-\text{g}}\right) = \left(\sum_{h=1}^{\infty} \frac{\mathbb{E}_t[\text{Div}_{t+h}]}{(1+\text{r})^h}\right) \Big/ \mathbb{E}_t[\text{Div}_{t+1}] \tag{15}$$

In other words, to capture the present value of dividends from year $(t+h)$ onward, just multiply next year's dividend forecast by PD $= \left(\frac{1}{\text{r}-\text{g}}\right)$.

The Gordon model gives a clean closed-form expression. But it assumes that discount rates and dividend growth rates are both constant over time. In the real world, neither is. The Gordon model handles this time variation poorly. As we shall see, even small changes in $(\text{r}-\text{g})$ produce large swings in the forward multiple $\left(\frac{1}{\text{r}-\text{g}}\right)$, and these swings are highly nonlinear.

Campbell and Shiller (1988a)'s contribution is to provide a linear approximation to the time-varying version of this relationship. If we exponentiate Equation (13), we see that it is just a fancy time-varying version of $\left(\frac{1}{\text{r}-\text{g}}\right)$

$$\left(\frac{1}{e^{\sum_{h=1}^{\infty} \rho^{h-1} \cdot \{\mathbb{E}_t[\text{R}_{t+h}] - \mathbb{E}_t[\Delta \log \text{Div}_{t+h}]\}}}\right) \tag{16}$$

Instead of working directly with the Gordon model, they start from the definition of stock returns and use a first-order Taylor expansion.

2 Using Realized Data

The Campbell-Shiller framework takes two per-share time series as inputs, price (Price_t) and dividends (Div_t), and constructs three objects from them: the log dividend yield $\ln(\text{Div}_t/\text{Price}_t)$, the log dividend growth rate $\Delta \log \text{Div}_{t+1} = \ln(\text{Div}_{t+1}/\text{Div}_t)$, and the log return $\text{R}_{t+1} = \ln\left(\frac{\text{Price}_{t+1}+\text{Div}_{t+1}}{\text{Price}_t}\right)$. The number of shares outstanding, $\#\text{Shares}_t$, appears nowhere. This omission is the source of serious problems when the formula is applied to firms that issue new equity or repurchase their own shares.

2.1 Repurchases Break the Formula

To see why, consider two firms that are identical at time t. Each has $\#\text{Shares}_t = 100\text{M}$ shares outstanding, a share price of $\text{Price}_t = \$5.00/\text{sh}$, a market capitalization of \$500M, total earnings of \$50M, and EPS of \$0.50/sh. During the period, each firm earns \$50M and distributes \$25M to shareholders. Last year, both firms paid a \$0.25/sh dividend. Each retains \$25M and invests the money to generate a 10% ROE (return on equity). Both firms discount the future cash flows at 8% annually.

The two firms differ only in the form of distribution and financing. Baseline Inc (BSLN) pays \$25M in cash dividends. No issuance, no repurchase. Repurchase Corp (REPC) pays no cash dividends and instead repurchases \$25M worth of shares at the ex-distribution price. Baseline Inc and Repurchase Corp are directly comparable: same investment, same total payout, same fundamental value. The only difference is the form in which the \$25M is returned to shareholders.

Let's start with Baseline Inc. Its \$25M investment at 10% ROE generates \$2.5M per year in perpetuity, which is valued at $\frac{\$2.5\text{M}}{8\%} = \$31.25\text{M}$, for an NPV of $\$31.25\text{M} - \$25\text{M} = \$6.25\text{M}$. Given the firm's initial market cap and subsequent earnings, its cum-distribution value is $\$550\text{M} = \$500\text{M}+\$50\text{M}$. Firm A pays out a \$25M dividend and adds \$6.25M in NPV from investing, yielding a \$531.25M ex-distribution value. With $\#\text{Shares}^{\text{BSLN}}_{t+1} = 100\text{M}$ shares, Baseline Inc's dividend and ex-distribution share price are $\text{Div}^{\text{BSLN}}_{t+1} = \$0.25/\text{sh}$ and $\text{Price}^{\text{BSLN}}_{t+1} = \$5.3125/\text{sh}$.

Now consider Repurchase Corp. This company has the same assets and investments as Baseline Inc. It also distributes the same amount of money. Hence, both firms have the same ex-distribution valuation, \$531.25M. The only difference is that Repurchase Corp switched from paying a \$25M dividend to spending \$25M on buybacks. The firm started off the year with $\#\text{Shares}_t = 100\text{M}$, so we can solve for Repurchase Corp's ex-distribution price using the formula

$$\text{Price}_{t+1}^{\text{REPC}} = \frac{\$531.25\text{M}}{100\text{M} - \left(\frac{\$25\text{M}}{\text{Price}_{t+1}^{\text{REPC}}}\right)} \tag{17}$$

This gives $\text{Price}_{t+1}^{\text{REPC}} = \$5.5625/\text{sh}$ and $\#\text{Shares}_{t+1}^{\text{REPC}} = 100\text{M} - \frac{\$25\text{M}}{\$5.5625/\text{sh}} = 95.506\text{M}$.

For Baseline Inc, the Campbell-Shiller inputs are all well-defined. The company's return is given by $\text{R}_{t+1}^{\text{BSLN}} = \log\left(\frac{\$5.3125/\text{sh}+\$0.25/\text{sh}}{\$5.00/\text{sh}}\right) = \log(1.1125) = 10.65\%$. Since the company paid \$25M in dividends last year, its log dividend growth rate is $\Delta \log \text{Div}_{t+1}^{\text{BSLN}} = 0\%$. The firm's PD ratio is $\text{PD}_{t+1}^{\text{BSLN}} = \left(\frac{\$5.3125/\text{sh}}{\$0.25/\text{sh}}\right) = 21.25\times$.

We can compute Repurchase Corp's return, $\text{R}_{t+1}^{\text{REPC}} = \log\left(\frac{\$5.5625/\text{sh}+\$0.00/\text{sh}}{\$5.00/\text{sh}}\right) = \log(1.1125) = 10.65\%$, and it is identical to that of Baseline Inc. This is exactly as Modigliani-Miller requires: the two firms have the same assets, the same investment, and the same total payout, so they must deliver the same return. The \$25M is simply delivered in different ways.

This is a good start. But look at what happens to Repurchase Corp's other Campbell-Shiller inputs. Since the company has switched to distributing all its cash via repurchases, Repurchase Corp's per-share dividend next year will be $\text{Div}_{t+1}^{\text{REPC}} = \$0.00/\text{sh}$. As a result, the firm's log dividend growth rate is undefined

$$\Delta \log \text{Div}_{t+1}^{\text{REPC}} = \log\left(\frac{\$0.00/\text{sh}}{\$0.25/\text{sh}}\right) = \log(0) = -\infty \tag{18}$$

We get the opposite problem when trying to compute Repurchase Corp's multiple price-to-dividend ratio, $\text{PD}_{t+1}^{\text{REPC}} = \left(\frac{\$5.5625/\text{sh}}{\$0.00/\text{sh}}\right) = +\infty$. There is simply no way to evaluate the Campbell-Shiller formula for Repurchase Corp even though it is an economically identical firm to Baseline Inc with the same assets, investments, total shareholder payout, and return.

This is not a contrived example. The majority of large US public companies now return cash primarily through repurchases rather than dividends (Boudoukh et al., 2007; Ben-David and Chinco, 2026). A firm that switches from paying dividends to repurchasing shares does not change its fundamental value or its expected return, but it does make the Campbell-Shiller framework inapplicable. Moreover, the problem runs deeper than the $\log(0)$ issue.

Even for firms that pay some dividends alongside their buybacks, repurchases introduce share termination risk. The return definition $\mathrm{R}_{t+1} = (\mathrm{Price}_{t+1} + \mathrm{Div}_{t+1})/\mathrm{Price}_t$ assumes that a shareholder who holds one share at time t still holds that share at time $(t+1)$. With Baseline Inc, this is guaranteed since the firm only pays dividends. With Repurchase Corp, it is not. There is a 4.5% chance that a randomly selected share at time t will get repurchased, in which case the owner will no longer hold a claim on the firm at time $(t+1)$.

Put another way, Campbell-Shiller gives a log-linear approximation to

$$\mathrm{Price}_t \;=\; \mathbb{E}_t\left[\sum_{h=1}^{\infty}\left(\prod_{i=1}^{h}\frac{1}{1+\mathrm{R}_{t+i}}\right)\cdot \mathrm{Div}_{t+h}\right] \tag{19}$$

This expression represents the value of a claim conditional on never being repurchased, but that is not what investors get when they buy a share of Repurchase Corp. An investor in this company owns a claim that, each period, might be terminated early with a cash payment. Structurally, this is closer to a callable bond than to a perpetual coupon stream.

Proposition 2.1. *Suppose Baseline Inc and Repurchase Corp both distribute $X > \$0$. If Repurchase Corp distributes the entire amount via buybacks, then*

(a) Both firms deliver the same per-share return: $\mathrm{R}^{\mathrm{BSLN}}_{t+1} = \mathrm{R}^{\mathrm{REPC}}_{t+1}$.

(b) For Baseline Inc, $\Delta \log \mathrm{Div}^{\mathrm{BSLN}}_{t+1}$ *and* $\mathrm{PD}^{\mathrm{BSLN}}_{t+1}$ *are both well-defined.*

(c) For Repurchase Corp, neither is: $\Delta \log \mathrm{Div}^{\mathrm{REPC}}_{t+1} = -\infty$ *and* $\mathrm{PD}^{\mathrm{REPC}}_{t+1} = +\infty$.

If Repurchase Corp distributes $\omega \cdot X$ via buybacks and the rest as a cash dividend for some $\omega \in (0,1)$, each share will still face a positive termination probability of

$$\omega \cdot \left(\frac{X}{\#\mathrm{Shares}^{\mathrm{REPC}}_t \cdot \mathrm{Price}^{\mathrm{REPC}}_{t+1}}\right) > 0\% \tag{20}$$

A correct present-value formula for Repurchase Corp would need to account for any dividends paid on shares that were not repurchased as well as the expected buyout payment for shares that were repurchased. Under Modigliani-Miller with fair-value repurchases and frictionless reinvestment, these two components together would be worth the same amount as Baseline Inc. After all, the only difference between Repurchase Corp and Baseline Inc is how shareholders receive their money. Yet Campbell-Shiller cannot speak to this difference because it lacks any notion of share count.

2.2 New Issuance Can Be Misleading

From a firm's perspective, new issuance is the opposite of repurchasing shares. Instead of spending $X > \$0$ reducing the share count, a company raises $X > \$0$ of equity capital at the cost of increasing its share count. However, while the policies are perfectly symmetric, their impact on Campbell-Shiller is not. New issuance does not break the log-linear approximation in the same way that repurchases do. Instead, it introduces a different problem.

To illustrate the problem, let's compare the Baseline Inc to another firm, Issuance NV (NVIS). This new company has all the same financials as the first two firms. Issuance NV starts off with $\#\text{Shares}_t = 100\text{M}$ shares outstanding, a share price of $\text{Price}_t = \$5.00/\text{sh}$, a market capitalization of \$500M, total earnings of \$50M, and EPS of \$0.50/sh. During the period, the company earns \$50M and distributes \$25M to shareholders. Last year, the firm paid a \$0.25/sh dividend.

Issuance NV pays a \$25M cash dividend, just like Baseline Inc. However, the company also raises \$25M of equity capital by issuing 5M shares at $\text{Price}_t = \$5.00/\text{sh}$ and uses the extra funds to invest \$50M at 10% ROE. By doubling its investment, the company generates \$5.0M per year in perpetuity, for an NPV of $\left(\frac{\$5.0\text{M}}{8\%}\right) - \$50\text{M} = \$12.5\text{M}$. Thus, Issuance NV has an ex-distribution valuation of

$$\$562.5\text{M} = \$500\text{M} + \underbrace{\$50\text{M}}_{\text{Earnings}} - \underbrace{\$25\text{M}}_{\text{Div}} + \underbrace{\$25\text{M}}_{\text{Issuance}} + \underbrace{\$12.5\text{M}}_{\text{Investment}} \tag{21}$$

giving the company an ex-distribution $\text{Price}_{t+1}^{\text{NVIS}} = \frac{\$562.5\text{M}}{105\text{M}} = \$5.357/\text{sh}$.

Unlike with Repurchase Corp, the Campbell-Shiller inputs for Issuance NV are all finite. The company has a return of $R^{NVIS}_{t+1} = \log\left(\frac{\$5.357/\text{sh}+\$0.25/\text{sh}}{\$5.00/\text{sh}}\right) = \log(1.1214) = 11.47\%$. The firm's log dividend growth rate is $\Delta \log \text{Div}^{NVIS}_{t+1} = \log\left(\frac{\$0.25/\text{sh}}{\$0.25/\text{sh}}\right) = 0\%$. And its PD ratio is $\text{PD}^{NVIS}_{t+1} = \left(\frac{\$5.357/\text{sh}}{\$0.25/\text{sh}}\right) = 21.43\times$. The framework does not produce undefined quantities. In this sense, new issuance is compatible with Campbell-Shiller.

But the compatibility is misleading. Issuance NV earns a higher per-share return than Baseline Inc (11.47% vs. 10.65%) because it invests twice as much in positive-NPV projects ($50M vs. $25M). In subsequent periods, however, Issuance NV will experience lower dividend growth on a per-share basis because the firm's earnings are now split across 105M shares rather than 100M. Issuance dilutes per-share dividend growth even as it accelerates total firm growth.

Proposition 2.2. *Baseline Inc and Issuance NV have identical financials at time t and pay the same per-share dividend at time $(t+1)$. If Issuance NV also issues $\Delta\#\text{Shares}_{t+1} > 0$ new shares and invests the proceeds, then the company will have*

(a) $\Delta \log \text{Div}^{NVIS}_{t+1}$ and PD^{NVIS}_{t+1} that are well-defined;

(b) A return that is higher than Baseline Inc's, $R^{NVIS}_{t+1} > R^{BSLN}_{t+1}$; and

(c) Future per-share dividend growth below Baseline Inc

$$\mathbb{E}_t[\Delta \log \text{Div}^{NVIS}_{t+h}] \;<\; \mathbb{E}_t[\Delta \log \text{Div}^{BSLN}_{t+h}] \qquad \text{for } h \geq 2 \tag{22}$$

Issuance is mathematically compatible with Campbell-Shiller: existing shares persist, and all log quantities remain finite. The problem is that the Campbell-Shiller decomposition always interprets lower per-share dividend growth as bad news. It has no way to distinguish dilution caused by value-creating issuance from genuinely weak fundamentals. A researcher who runs a variance decomposition on Issuance NV's data will see lower "cash-flow news" in periods when the firm issues shares, even if the issuance finances positive-NPV projects that increase total firm value. The formula treats the per-share dividend slowdown the same way it would treat a collapse in earnings.

3 Taking Expectations

Set aside the issues raised in Section 2 and suppose that the Campbell-Shiller formula holds exactly in realized data. Even so, if you want to take expectations and apply the formula to forward-looking data, then you need investors to consistently price stocks using present-value logic. Investors must be thinking in these terms. In Subsection 3.1, we document that many are not. In Subsection 3.2, we then show why even small departures from perfect present-value reasoning are a problem. For Campbell-Shiller to be considered an accounting "identity", investors must apply present-value logic with a degree of precision that far exceeds what anyone can realistically achieve.

3.1 Not Using Present-Value Logic

When the Campbell-Shiller formula is applied to realized data, there is no behavioral assumption. A realized return is defined as $R_{t+1} = (\text{Price}_{t+1} + \text{Div}_{t+1})/\text{Price}_t$. This is the definition. It holds regardless of why the investor bought the stock or what she thought it was worth. But the forward-looking version of the formula is different. It requires investors to actually price stocks using present-value logic: $\text{Price}_t = \frac{\mathbb{E}_t[\text{Div}_{t+1}]}{1+r} + \frac{\mathbb{E}_t[\text{Price}_{t+1}]}{1+r}$. Nothing forces this to be true. An investor is free to value a share at \$5/sh regardless of what $\mathbb{E}_t[\text{Div}_{t+1}]$ and $\mathbb{E}_t[\text{Price}_{t+1}]$ happen to be.

For the Campbell-Shiller formula to hold in expectation, investors must be thinking in present-value terms. They must be reliably applying the logic embedded in Equations (14a), (14b), and (14c). It is not enough for them to have reasonable beliefs about future cash flows. Investors must impose a specific constraint on the numbers in their heads: a stock's current price should equal the expected discounted value of all future dividends. And there is a lot of evidence to suggest that many investors do not think this way.

Sell-side equity analysts are professional forecasters whose price targets move billions of dollars of trading volume. Ben-David and Chinco (2025) documents that these market participants typically value companies using a backward-

looking multiple rather than a forward-looking one like the Gordon model says to do. Over 3/4 of analyst reports explicitly state that they used some version of the pricing rule below

$$\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE} \tag{23}$$

where TrailingPE = Price/EPS is the trailing twelve-month PE ratio. Analysts take their short-term EPS forecast, $\mathbb{E}[\text{EPS}]$, and scale it by the stock's trailing PE. They do not pick a PE ratio that reflects the present value of the company's expected earnings stream in years $t = 2, 3, 4, \ldots$ The result is a price target anchored to the stock's recent valuation history, not to a discounted cash flow model. There is no discount rate. There is no terminal value. There is no infinite sum over future cash flows. The entire pricing exercise is a two-variable calculation: expected earnings times a backward-looking multiple.

It is not just sell-side analysts, either. In February 2013, David Einhorn of Greenlight Capital Management gave a public presentation to fellow Apple shareholders in an effort to convince the company to distribute its cash hoard.[1] Figure 1 shows a key slide from this presentation. Einhorn's argument was that an \$84B cash-financed repurchase program would mechanically boost Apple's per-share earnings from \$45/sh to \$52.80/sh, and that this \$7.80/sh EPS gain should be valued at Apple's existing 10× PE ratio. His pitch was that repurchases would increase share prices by $\$528/\text{sh} - \$450/\text{sh} = \$7.80/\text{sh} \times 10 \approx +\$78/\text{sh}$.

Notice the logic. Einhorn did not discount the future earnings stream that the repurchase would generate. He did not estimate how the buyback would change Apple's cost of capital or long-run growth rate. He took his EPS forecast and multiplied it by a backward-looking multiple, exactly the same pricing rule that Ben-David and Chinco (2025) documents among sell-side analysts. Einhorn used Equation (23). This is a sophisticated investor, managing billions of dollars, making a public case to the shareholders of the world's most valuable company, and the valuation framework he reached for was $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$, not a discounted cash flow model.

[1] David Einhorn. "iPrefs: Unlocking Value." Conference Call. February 21, 2013. [link]

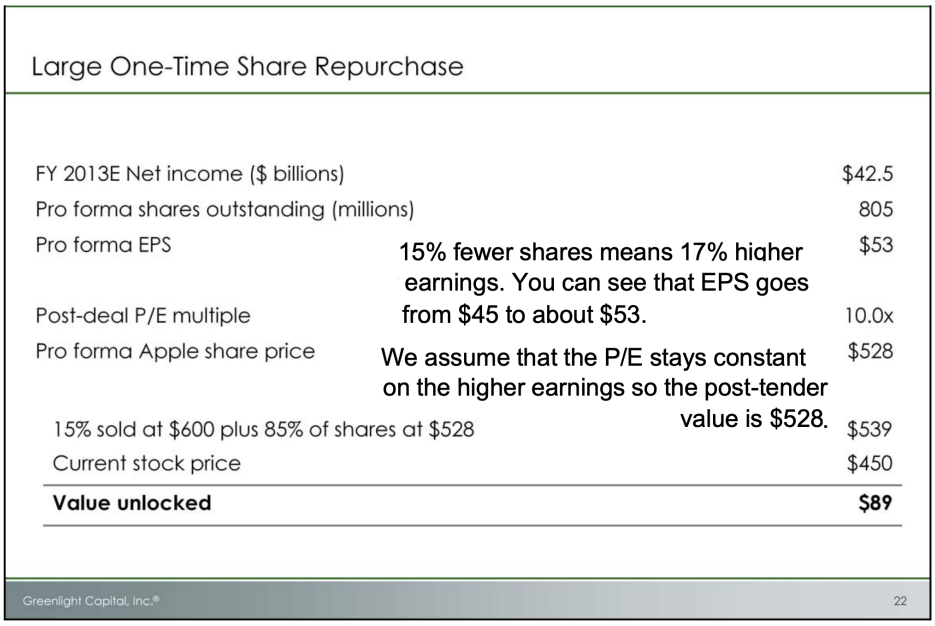


Figure 1. Slide from a February 2013 presentation by David Einhorn of Greenlight Capital Management to Apple (AAPL) shareholders. Einhorn calculates that an $84B cash-financed repurchase program would generate a +7.80/sh EPS pop and a share price increase of $\$528/\text{sh} - \$450/\text{sh} = \$7.80/\text{sh} \times 10 \approx +\$78/\text{sh}$.

Hartzmark and Sussman (2025) provide perhaps the most direct test of whether investors apply present-value logic. If people reliably set prices by discounting expected future payoffs, then you should be able to hand someone information about a company's fundamentals and have them give you back the correct price. This is exactly what the authors do. They show online participants, MBA students, and asset-management professionals information about fundamentals and ask them to identify the market price. Across all three groups, subjects cannot detect even extreme deviations from actual prices. When asked directly, both professionals and large language models say that it is not possible to reliably estimate the level of the market to within 10% using fundamental information alone, without first seeing the prevailing price.

One reaction to this evidence is that a lot of people are making dumb mistakes. That could be. But it also might not be unreasonable to deviate from present-value logic. Consider the housing market. If you own a home, which discount

rate did you use when deciding how much to pay for it? We are guessing you did not calculate a discounted cash flow model. Almost no one does. Homebuyers look at comparable sales, recent transaction prices for similar homes in similar neighborhoods. The pricing rule is backward-looking: what did this kind of house sell for recently? Not what is the discounted present value of the rents this house will generate over the next thirty years?

This backward-looking approach is so deeply embedded that the dominant methodology for measuring home prices, the Case-Shiller repeat-sales index, is explicitly built on it (Case and Shiller, 1987, 1989). This does not make you a bad economist. It just means that not every market operates according to present-value principles. Bond markets do. Housing markets clearly do not. The evidence above suggests that, in many cases, neither do equity markets.

3.2 Unrealistic Levels of Precision

Some equity investors clearly do apply present-value reasoning. But as the previous subsection documented, many do not. You can debate the exact fraction, but the number of investors who are not applying present-value logic is not zero. How much does this matter? You might think that approximate present-value reasoning is good enough.

If most investors are roughly in the right ballpark, then Campbell-Shiller should still work as a useful approximation, right? Wrong. The Campbell-Shiller formula is approximating the Gordon model, and the Gordon model's price-dividend ratio is highly nonlinear, which means that even small deviations from exact present-value reasoning produce large pricing errors. The rest of this subsection quantifies how small "small" really is.

For something to qualify as an accounting identity, you have to be able to reliably count the relevant terms. In an income statement, every line item can be measured, and the bottom line is the sum of all the entries above it. A present-value relationship ought to work in the same way. Each line item is an expected discounted cash flow, one for each future period.

We can make this analogy precise. In the Gordon model, the single largest term is next year's expected dividend, $\frac{\mathbb{E}[\text{Div}_1]}{1+r}$. Every subsequent term is smaller by a factor of $\left(\frac{1+g}{1+r}\right) < 1$. So the minimum standard for calling the formula an identity is straightforward: estimation errors in the model's parameters had better produce price changes that are smaller than the present value of next year's dividend. If they do not, we cannot reliably count even the first term.

The Gordon model says $\text{Price}_0 = \mathbb{E}[\text{Div}_1] \times \left(\frac{1}{r-g}\right)$. A small change in the cap rate, $\Delta(r-g)$, shifts the price by approximately

$$\Delta\text{Price}_0 \approx -\mathbb{E}[\text{Div}_1] \cdot \left(\frac{1}{r-g}\right)^2 \times \Delta(r-g) \tag{24}$$

We want this to be no larger than $\frac{\mathbb{E}[\text{Div}_1]}{1+r}$, the present value of next year's expected dividend. Setting these equal and solving gives a maximum tolerable cap-rate error on the order of DY^2, where $\text{DY} = \frac{\mathbb{E}[\text{Div}_1]}{\text{Price}_0}$ is the forward dividend yield.

Proposition 3.2. *Suppose* $\text{Price}_0 = \mathbb{E}[\text{Div}_1] \times \left(\frac{1}{r-g}\right)$ *as in* (14c)*. If the cap rate is perturbed by more than* DY^2*, where* $\text{DY} = \frac{\mathbb{E}[\text{Div}_1]}{\text{Price}_0}$ *is the forward dividend yield, then the resulting price change will exceed the present value of next year's expected dividend*

$$|\Delta(r-g)| > \text{DY}^2 \qquad \rightsquigarrow \qquad |\Delta\text{Price}_0| > \frac{\mathbb{E}[\text{Div}_1]}{1+r} \tag{25}$$

The S&P 500 has a forward dividend yield of $\text{DY} \approx 2\%$. So $\text{DY}^2 \approx 0.04\%$. That is the maximum amount of parameter noise you can tolerate before the price impact of your estimation error exceeds the present value of the very next dividend payment. This is preposterously small. Standard estimates of the equity risk premium range from 4% to 8%, a spread that is 100× larger than the required precision.

If researchers cannot pin down the discount rate to within 0.04%, then errors in r and g will move the price by more than the value of next year's dividend, and it makes no sense to call the result a present-value "identity." You cannot claim to be counting terms in a sum if you cannot keep track of whether the first term has been included.

4 Subjective Beliefs

The Campbell-Shiller decomposition has become the standard tool for studying the asset-pricing implications of biased beliefs. The conventional wisdom is that Equation (26) will approximately hold no matter which $\widetilde{\mathbb{E}}_t[\cdot]$ you choose. De la O and Myers (2021) write that subjective expectations "do not need to be rational. Given any set of dividend growth and return expectations, [the present-value identity] will hold." Bordalo et al. (2024) similarly assert that "because this decomposition is derived from an identity, it holds even if expectations differ from the full-information rational expectations benchmark."

However, this is not true. Even if there are no changes in share count and investors reliably enforce present-value relationships, you cannot plug any arbitrary subjective beliefs into Campbell-Shiller. Subsection 4.1 shows that viable subjective beliefs must satisfy a specific error-aggregation condition. Subsection 4.2 proves that only beliefs that come from a well-posed probability distribution will satisfy the condition. Subjective beliefs that involve one-off errors about a single moment will not. For example, extrapolating $\mathbb{E}[\mathrm{R}]$ alone will not work. Subsection 4.3 looks at three applications from the literature.

4.1 Error-Aggregation Condition

The standard approach works as follows. Start from the Campbell-Shiller approximation in realized values

$$\log \mathrm{PD}_t \approx \kappa + \Delta \log \mathrm{Div}_{t+1} - \mathrm{R}_{t+1} + \rho \cdot \log \mathrm{PD}_{t+1} \tag{8}$$

Next, apply any subjective operator to both sides, $\widetilde{\mathbb{E}}_t[\cdot]$, and iterate forward

$$\log \mathrm{PD}_t \approx \sum_{h=1}^{\infty} \rho^{h-1} \cdot \{ \widetilde{\mathbb{E}}_t[\Delta \log \mathrm{Div}_{t+h}] - \widetilde{\mathbb{E}}_t[\mathrm{R}_{t+h}] \} \tag{26}$$

This formula is just a subjective-belief analog to Equation (13). Just like before, we are going to suppress the constant term, $\kappa \cdot \left(\frac{1}{1-\rho}\right)$, to minimize clutter.

Even if a) there is no new issuance or repurchases and b) investors apply present-value logic with a jeweler's attention to detail, Equation (26) will still not hold for any arbitrary choice of subjective beliefs. Two additional conditions must be satisfied. First, agents must not subjectively expect the PD ratio to grow persistently at rate $1/\rho$ or faster

$$\lim_{H\to\infty} \rho^H \cdot \widetilde{\mathbb{E}}_t[\log \mathrm{PD}_{t+H}] = 0 \tag{27}$$

This is a restriction on the tail behavior of beliefs: agents' views about the distant future cannot be too explosively optimistic about relative valuations.

The second condition is more subtle and more important. Most behavioral models specify the structure of belief distortions $\widetilde{\mathbb{E}}_t[\cdot]$ and separately derive equilibrium prices. For the subjective-belief version of Campbell-Shiller to hold, the errors implied by the belief specification must exactly cancel out when summed across all future horizons. This is not automatically satisfied.

Proposition 4.1. *Suppose the Campbell-Shiller formula holds under the correct expectations, $\mathbb{E}_t[\cdot]$. Further suppose that the transversality condition is satisfied under the subjective beliefs, $\widetilde{\mathbb{E}}_t[\cdot]$. Define the belief error at horizon h as*

$$\tilde{\varepsilon}_h \;=\; (\widetilde{\mathbb{E}}_t - \mathbb{E}_t)[\Delta \log \mathrm{Div}_{t+h}] \;-\; (\widetilde{\mathbb{E}}_t - \mathbb{E}_t)[\mathrm{R}_{t+h}] \tag{28}$$

The subjective Campbell-Shiller formula in Equation (26) holds if and only if

$$\sum_{h=1}^{\infty} \rho^{h-1} \cdot \tilde{\varepsilon}_h \;=\; 0 \tag{29}$$

We refer to Equation (29) as the "error-aggregation condition". In words, the discounted sum of all belief errors across every future horizon must be exactly zero. There is an important conceptual distinction at play here. It is one thing to require that beliefs satisfy a moment condition. For example, you might want subjective expectations about future returns to match certain unconditional averages. It is quite another thing to require that errors in beliefs satisfy such moment conditions. Equation (29) is the latter.

The Campbell-Shiller framework does not merely ask that investors hold biased beliefs; it asks that investors impose a precise adding-up constraint on the errors themselves. An agent who is too optimistic about dividend growth at short horizons must be correspondingly pessimistic about dividend growth or optimistic about discount rates at other horizons. Moreover, the factor of ρ^{h-1} implies that the offsetting forces must occur in exactly the right proportions.

4.2 A Specific Kind of Mistake

This is a strange kind of irrationality. If investors understood the structure of their own errors well enough to ensure that those errors aggregate correctly, netting out across variables and horizons in precisely the way the formula requires, they would have enough information to simply correct the errors and stop making them. The adding-up constraint thus describes agents who are simultaneously irrational (they have biased beliefs) and extraordinarily sophisticated (they calibrate the biases to satisfy a cross-equation restriction). This combination is difficult to motivate psychologically.

And yet beliefs like these are exactly what the behavioral asset-pricing literature produces. This is not a coincidence. It is an artifact of how behavioral economists defend their (often very reasonable) assumptions. For example, Rabin (2013) advocates an approach to incorporating psychology into economics in which the researcher takes an otherwise standard model, makes one targeted change to preferences or beliefs, and then ensures that the single change is the only non-standard element.

You do not simply assert that agents are biased; you embed the bias inside a framework where the agent is otherwise solving a fully rational problem. Most of Laibson (1997)'s paper on hyperbolic discounting, for instance, is devoted to showing that a present-biased agent solves a forward-looking optimal-consumption problem using a Bellman equation. The bias is the point of the paper; the Bellman equation is the price of admission.

What kinds of subjective beliefs satisfy the error-aggregation condition? Proposition 4.2 shows that the answer is tightly linked to whether the beliefs

come from a well-posed probability distribution. To build intuition, consider what the error-aggregation condition demands at just a single horizon. Suppose a subjective expectation operator, $\widetilde{\mathbb{E}}_t[\cdot]$, satisfies the Campbell-Shiller approximation at $h = 1$

$$\widetilde{\mathbb{E}}_t[\mathrm{R}_{t+1}] \;=\; \kappa + \widetilde{\mathbb{E}}_t[\Delta \log \mathrm{Div}_{t+1}] + \rho \cdot \widetilde{\mathbb{E}}_t[\log \mathrm{PD}_{t+1}] - \log \mathrm{PD}_t \tag{30}$$

This must rule out some kinds of biased beliefs. Which ones exactly? What sorts of distortions are permissible if we still want the Campbell-Shiller to hold?

Proposition 4.2. *Let $\widetilde{\mathbb{E}}_t[\cdot]$ denote an arbitrary subjective-belief specification, which distorts the conditional expectations of R_{t+1}, $\Delta \log \mathrm{Div}_{t+1}$, and $\log \mathrm{PD}_{t+1}$ by*

$$\widetilde{\mathbb{E}}_t[\mathrm{R}_{t+1}] = \mathbb{E}_t[\mathrm{R}_{t+1}] + \delta_{\mathrm{R}} \tag{31a}$$

$$\widetilde{\mathbb{E}}_t[\Delta \log \mathrm{Div}_{t+1}] = \mathbb{E}_t[\Delta \log \mathrm{Div}_{t+1}] + \delta_{\mathrm{Div}} \tag{31b}$$

$$\widetilde{\mathbb{E}}_t[\log \mathrm{PD}_{t+1}] = \mathbb{E}_t[\log \mathrm{PD}_{t+1}] - \delta_{\mathrm{PD}} \tag{31c}$$

These subjective beliefs satisfy the Campbell-Shiller formula at horizon $h = 1$ if and only if

$$\delta_{\mathrm{R}} = \delta_{\mathrm{Div}} - \rho \cdot \delta_{\mathrm{PD}} \tag{32}$$

Under any probability measure this condition holds automatically. By contrast, it generically fails for subjective beliefs that specify δ_{R}, δ_{Div}, and δ_{PD} independently, even if they are all generated by the same formula applied variable by variable.

The logic behind Proposition 4.2 is straightforward. Under any probability measure, expectations are linear

$$\mathbb{E}[\, a \cdot X + b \cdot Y \,] \;=\; a \cdot \mathbb{E}[X] + b \cdot \mathbb{E}[Y] \tag{33}$$

This is because expectation is integration, and integration is a linear operator. Since the Campbell-Shiller formula holds in realized values, taking expectations under any measure preserves equality. That includes the true measure, any equivalent martingale measure, or any incorrect-but-coherent subjective measure. A Bayesian agent with the wrong model (wrong priors, wrong likeli-

hood, or wrong state space) still has a coherent posterior distribution, and their expectations are still linear. So the adding-up constraint must be satisfied.

A belief distortion that is not derived from a probability measure has no reason to preserve linearity. To see why, consider three random variables X, Y, and Z satisfying $Z = a \cdot X + b \cdot Y + c$ in realized values. Under a coherent measure, $\mathbb{E}[Z] = a \cdot \mathbb{E}[X] + b \cdot \mathbb{E}[Y] + c$ automatically. Now apply an arbitrary distortion $\delta(\cdot)$ to each expectation independently: $\widetilde{\mathbb{E}}[X] = \mathbb{E}[X] + \delta_X$, $\widetilde{\mathbb{E}}[Y] = \mathbb{E}[Y] + \delta_Y$, and $\widetilde{\mathbb{E}}[Z] = \mathbb{E}[Z] + \delta_Z$. The adding-up constraint requires $\delta_Z = a \cdot \delta_X + b \cdot \delta_Y$. If the distortions are specified independently, this linear relationship among the distortions is a knife-edge condition. It holds if and only if the distortion functional is itself linear, which is equivalent to saying the distortion corresponds to integration under an alternative measure.

In other words: checking whether biased beliefs satisfy the Campbell-Shiller adding-up constraint is essentially checking whether the belief distortion corresponds to a probability distribution. The class of behavioral models that automatically satisfy the adding-up constraint is precisely the class where the agent is a Bayesian with the wrong model: the agent has incorrect priors, an incorrect likelihood, or an incorrect state space, but still performs coherent probabilistic inference. Only these "wrong-model Bayesian" agents have beliefs that integrate to a probability measure, ensuring linearity and hence consistency with the Campbell-Shiller formula.

The error-aggregation condition is therefore not an exotic corner case. It is a binding constraint on any attempt to use the Campbell-Shiller decomposition with subjective beliefs. To date, no one has put the problem in these terms. Researchers have not consciously responded to the error-aggregation condition; they have navigated around it without seeing it, settling on approaches that happen to avoid the problem for different reasons. The next subsection examines two such approaches and explains why each one works, or fails to work, in light of the constraint identified here.

4.3 Three Applications

Researchers using the Campbell-Shiller decomposition with subjective beliefs have, without articulating the error-aggregation condition, gravitated toward two approaches. Some build full equilibrium models in which the agent's beliefs come from a coherent probability distribution, which automatically satisfies the constraint. Others use survey data and leave one of the three Campbell-Shiller inputs unspecified, backing it out as a residual so that the formula holds by construction. Neither approach was adopted as a response to the error-aggregation condition. Examining each in light of this constraint clarifies what the literature has actually achieved and where it falls short.

Engineering coherence. Jin and Sui (2022) study a representative-agent model in which investors have extrapolative beliefs about stock-market returns: they overweight recent returns when forming forecasts of future returns. The behavioral assumption is a distortion to return expectations:

$$\delta_{\mathrm{R}} \;=\; \widetilde{\mathbb{E}}_t[\mathrm{R}_{t+1}] - \mathbb{E}_t[\mathrm{R}_{t+1}] \;\neq\; 0 \tag{34}$$

But this is not the only thing the model specifies.

Given these return beliefs and Epstein-Zin preferences, market clearing pins down the equilibrium price as a function of the state. This price function determines the agent's implied beliefs about dividend growth and future valuations. The agent does not separately decide what to believe about $\Delta \log \mathrm{Div}_{t+1}$ and $\log \mathrm{PD}_{t+1}$. Those expectations are restricted in equilibrium: the agent forms beliefs over the state, and the equilibrium price function maps those beliefs into implied expectations of all three Campbell-Shiller inputs.

Crucially, the agent's beliefs over the state constitute a well-defined probability measure. Thus, expectations are linear: $\mathbb{E}[a \cdot X + b \cdot Y] = a \cdot \mathbb{E}[X] + b \cdot \mathbb{E}[Y]$. Since the Campbell-Shiller formula holds in realized values, taking expectations under this measure preserves equality. The distortions δ_{R}, δ_{Div}, and δ_{PD} are not three free parameters. They are all determined by the single change of measure, and the adding-up condition $\delta_{\mathrm{R}} = \delta_{\mathrm{Div}} - \rho \cdot \delta_{\mathrm{PD}}$ is automatically satisfied.

But this comes at a cost. The behavioral assumption, return extrapolation, is simple. The machinery needed to make Campbell-Shiller work is not. The model requires a full specification of preferences, a consumption process, and market-clearing conditions. The agent's implied beliefs about cash flows and valuations are not extrapolative in the same sense that their return beliefs are; they are whatever the equilibrium demands. The simplicity of the original behavioral insight is buried under the apparatus needed to ensure that the error-aggregation condition holds.

Backing out residuals. De la O and Myers (2021) take a different approach. Rather than building an equilibrium model, they measure subjective expectations directly from survey data: cash-flow expectations from I/B/E/S analyst forecasts and return expectations from the Graham-Harvey CFO survey. In the notation of Proposition 4.2, this gives independent measurements of two of the three distortions: δ_{Div} from analyst forecasts and δ_{R} from the CFO survey. They then use the Campbell-Shiller approximation to decompose the variance of the price-dividend ratio into contributions from each component.

The justification for this exercise is that the Campbell-Shiller formula holds under any subjective beliefs. De la O and Myers (2021) write that subjective expectations "do not need to be rational. Given any set of dividend growth and return expectations, [the present-value identity] will hold." But the approximation has three inputs on the right-hand side: subjective expectations of returns, dividend growth, and future valuations. De la O and Myers (2021) independently measure the first two from surveys. The third, subjective expectations of the future price-dividend ratio, is not measured.

Since δ_{Div} and δ_{R} are measured independently, the adding-up condition in Equation (32), $\delta_{\text{R}} = \delta_{\text{Div}} - \rho \cdot \delta_{\text{PD}}$, pins down δ_{PD}: it is set to whatever value makes the equation balance. The logic is therefore circular. The premise of the exercise is that Campbell-Shiller holds under any beliefs, and this premise is what justifies plugging survey data into the decomposition. But the exercise only works because the researcher forces the formula to hold by leaving one input unspecified.

If Campbell-Shiller genuinely held for any subjective beliefs, this step would be unnecessary: you could measure all three inputs independently and the formula would be satisfied. The fact that δ_{PD} must be backed out is implicitly conceding that the formula does not hold when you bring independent measurements to the table. This has consequences for interpretation.

A variance decomposition attributes fluctuations in the price-dividend ratio to either cash-flow expectations or discount-rate expectations. But that attribution depends on where the residual lands. When δ_{Div} and δ_{R} are measured independently, the valuation component δ_{PD} inherits all the slack from the error-aggregation residual. If instead the researcher measured δ_{Div} and δ_{PD} independently and backed out δ_{R}, the return component would inherit the slack. The debate about which variable explains more price variation is likely an artifact of which survey a researcher happened to start with.

Diagnostic expectations. Bordalo et al. (2019) use the Campbell-Shiller framework to study the asset-pricing implications of diagnostic expectations. This bias distorts the conditional expectation of a random variable, X, by overweighting recent news

$$\widetilde{\mathbb{E}}_t[X] \;=\; \mathbb{E}_t[X] \;+\; \theta \cdot \left\{ \mathbb{E}_t[X] - \mathbb{E}_{t-1}[X] \right\} \tag{35}$$

where $\theta > 0$ is the diagnostic parameter. The paper measures diagnostic cash-flow expectations from analyst forecasts and backs out the implied return expectations from the Campbell-Shiller formula. Like De la O and Myers (2021), the equation is imposed rather than tested.

What would happen if it were tested? The adding-up condition in Equation (32) requires $\delta_{\text{R}} = \delta_{\text{Div}} - \rho \cdot \delta_{\text{PD}}$. Consider the simplest possible case: the agent has diagnostic beliefs about returns only, with correct beliefs about dividend growth and future valuations. That is, $\delta_{\text{Div}} = 0$ and $\delta_{\text{PD}} = 0$, but

$$\delta_{\text{R}} \;=\; \theta \cdot \left\{ \mathbb{E}_t[\text{R}_{t+1}] - \mathbb{E}_{t-1}[\text{R}_{t+1}] \right\} \;\neq\; 0 \tag{36}$$

The adding-up condition requires $\delta_{\text{R}} = 0 - \rho \cdot 0 = 0$. Any nonzero distortion to

return expectations alone, no matter how small, violates the error-aggregation condition. A single biased input is enough to break the formula.

The lesson is not that diagnostic expectations are wrong, or that they cannot be used with the Campbell-Shiller decomposition. The lesson is that you cannot simply plug diagnostic beliefs about $\mathbb{E}_t[\mathrm{R}_{t+1}]$ and call it a day. Diagnostic expectations, applied variable by variable, do not produce a coherent probability measure. So the adding-up condition is not automatically satisfied. To make it work, a researcher would need to do what Jin and Sui (2022) did: embed the error in a more elaborate model involving preferences, endowments, and market-clearing conditions. That way, the agent's beliefs over all three Campbell-Shiller inputs would stem from a single distorted probability distribution. The behavioral assumption itself is a one-line formula. But it takes a village to make that formula compatible with Campbell-Shiller.

Conclusion

The Campbell and Shiller (1988a) log-linear approximation is one of the most influential results in empirical asset pricing. The original paper is impeccable in its analysis and forthright about its limitations. The formula itself is elegant, tractable, and genuinely useful for the questions it was designed to answer. This paper is not a critique of Campbell-Shiller. It is an accounting of how many qualifications are needed to apply it.

There are many. The formula does not hold in realized data for firms that repurchase shares, because repurchases make the dividend growth rate undefined. Even setting repurchases aside, the formula does not automatically hold in expectation: the marginal investor must price stocks using present-value logic, and must know the relevant discount rates and growth rates to within a fraction of a percent. And even granting all of this, replacing the rational expectation operator with subjective beliefs requires that the *errors* in those beliefs satisfy a precise adding-up constraint, $\delta_{\mathrm{R}} = \delta_{\mathrm{Div}} - \rho \cdot \delta_{\mathrm{PD}}$.

Each of these qualifications is a separate result, and each one narrows the domain in which the formula can be applied. By the time you have navigated

all of them, the set of permissible applications is considerably smaller than the folk wisdom suggests. Campbell-Shiller holds for a representative agent with rational expectations who prices stocks using a dividend-discount model and who lives on a planet where companies never repurchase shares. It also holds if you replace rational expectations with beliefs from a coherent but incorrect probability measure. Beyond that, every extension requires additional assumptions, additional infrastructure, or both.

The error-aggregation condition is a case in point. To use Campbell-Shiller with subjective beliefs that do not come from a coherent probability measure, a researcher must either build a full equilibrium model, with preferences, endowments, and market-clearing conditions, so that the agent's beliefs over all three inputs derive from a single distorted distribution, or leave one of the three inputs unspecified and back it out as a residual so that the equation holds by construction. Both approaches work, in the sense that the equation is satisfied. But neither is free. The first buries a simple behavioral assumption under a mountain of modeling infrastructure. The second produces a variance decomposition whose conclusions depend on which input the researcher chooses to measure and which one absorbs the slack.

None of this means that Campbell-Shiller should be abandoned. It means that the formula is not a panacea. The log-linear approximation does not automatically hold in any setting, for any beliefs, without further assumptions. Researchers who want to use it outside its original domain should be explicit about what those further assumptions are and what they cost.

References

Adam, K., A. Marcet, and J. Beutel (2017). Stock price booms and expected capital gains. American Economic Review 107(8), 2352–2408.

Adam, K., A. Marcet, and J. Nicolini (2016). Stock-market volatility and learning. Journal of Finance 71(1), 33–82.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer (2015). X-capm: An extrapolative capital asset pricing model. Journal of Financial Economics 115(1), 1–24.

Ben-David, I. and A. Chinco (2025). Expected EPS × Trailing PE: Pricing without discounting. Working paper.

Ben-David, I. and A. Chinco (2026). EPS-maximizing payout policy. Working paper.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2019). Diagnostic expectations and stock returns. Journal of Finance 74(6), 2839–2874.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2024). Finance without (exotic) risk. Working paper.

Bordalo, P., N. Gennaioli, and A. Shleifer (2018). Diagnostic expectations and credit cycles. Journal of Finance 73(1), 199–227.

Boudoukh, J., R. Michaely, M. Richardson, and M. Roberts (2007). On the importance of measuring payout yield: Implications for empirical asset pricing. Journal of Finance 62(2), 877–915.

Campbell, J. (1991). A variance decomposition for stock returns. The Economic Journal 101(405), 157–179.

Campbell, J. (2017). Financial decisions and markets: a course in asset pricing. Princeton University Press.

Campbell, J. and R. Shiller (1988a). The dividend-price ratio and expectations of future dividends and discount factors. Review of Financial Studies 1(3), 195–228.

Campbell, J. and R. Shiller (1988b). Stock prices, earnings, and expected dividends. Journal of Finance 43(3), 661–676.

Case, K. E. and R. J. Shiller (1987). Prices of single-family homes since 1970: New indexes for four cities. New England Economic Review, 45–56.

Case, K. E. and R. J. Shiller (1989). The efficiency of the market for single-family homes. American Economic Review 79(1), 125–137.

Cho, T., L. Kremens, D. Lee, and C. Polk (2024). Scale or yield? A present-value identity. Review of Financial Studies 37(3), 950–988.

Cochrane, J. (2011). Presidential address: Discount rates. Journal of finance 66(4), 1047–1108.

Cochrane, J. (2017). Macro-finance. Review of Finance 21(3), 945–985.

De la O, R. and S. Myers (2021). Subjective cash flow and discount rate expectations. Journal of Finance 76(3), 1339–1387.

Engsted, T., T. Q. Pedersen, and C. Tanggaard (2012). The log-linear return approximation, bubbles, and predictability. Journal of Financial and Quantitative Analysis 47(3), 643–665.

Fuster, A., D. Laibson, and B. Mendel (2010). Natural expectations and macroeconomic fluctuations. Journal of Economic Perspectives 24(4), 67–84.

Gabaix, X. and R. Koijen (2024). In search of the origins of financial fluctuations: The inelastic markets hypothesis. Working paper.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. Review of Financial Studies 27(3), 714–746.

Hartzmark, S. and A. Sussman (2025). Price agnostic demand. Working paper.

Jin, L. J. and P. Sui (2022). Asset pricing with return extrapolation. Journal of Financial Economics 145(2), 273–295.

Koijen, R. and M. Yogo (2019). A demand-system approach to asset pricing. Journal of Political Economy 127(4), 1475–1515.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112(2), 443–478.

Lettau, M., S. Ludvigson, and P. Manoel (2018). Characteristics of mutual-fund portfolios: Where are the value funds? Working paper.

Nagel, S. and Z. Xu (2022). Asset pricing with fading memory. Review of Financial Studies 35(5), 2190–2245.

Rabin, M. (2013). An approach to incorporating psychology into economics. American Economic Review 103(3), 617–622.

Vuolteenaho, T. (2002). What drives firm-level stock returns? Journal of Finance 57(1), 233–264.

Technical Appendix

Proof. **(Proposition 2.1)**

Both firms have the same assets, investment, and total payout, so by Modigliani-Miller, total ex-distribution firm value is the same:

$$V_{t+1}^{\text{BSLN}} \;=\; V_{t+1}^{\text{REPC}} \;\equiv\; V_{t+1}$$

Part (a). Baseline Inc pays $\text{Div}_{t+1}^{\text{BSLN}} = X/\#\text{Shares}_t$ *per share and has ex-dividend price* $\text{Price}_{t+1}^{\text{BSLN}} = V_{t+1}/\#\text{Shares}_t$*. Its per-share return is*

$$R_{t+1}^{\text{BSLN}} \;=\; \frac{\text{Price}_{t+1}^{\text{BSLN}} + \text{Div}_{t+1}^{\text{BSLN}}}{\text{Price}_t} \;=\; \frac{V_{t+1} + X}{\#\text{Shares}_t \cdot \text{Price}_t}$$

Repurchase Corp pays no dividend, so $R_{t+1}^{\text{REPC}} = \text{Price}_{t+1}^{\text{REPC}}/\text{Price}_t$*, where the ex-distribution price satisfies*

$$\text{Price}_{t+1}^{\text{REPC}} \;=\; \frac{V_{t+1}}{\#\text{Shares}_t - X/\text{Price}_{t+1}^{\text{REPC}}}$$

Multiplying both sides by the denominator and solving:

$$\text{Price}_{t+1}^{\text{REPC}} \cdot \#\text{Shares}_t - X \;=\; V_{t+1} \qquad\Longrightarrow\qquad \text{Price}_{t+1}^{\text{REPC}} \;=\; \frac{V_{t+1} + X}{\#\text{Shares}_t}$$

Therefore $R_{t+1}^{\text{REPC}} = (V_{t+1} + X)/(\#\text{Shares}_t \cdot \text{Price}_t) = R_{t+1}^{\text{BSLN}}$*.*

Part (b). For Baseline Inc, $\text{Div}_{t+1}^{\text{BSLN}} = X/\#\text{Shares}_t > 0$ *and* $\text{Price}_{t+1}^{\text{BSLN}} = V_{t+1}/\#\text{Shares}_t > 0$*, so both*

$$\Delta \log \text{Div}_{t+1}^{\text{BSLN}} \;=\; \ln\!\left(\text{Div}_{t+1}^{\text{BSLN}}/\text{Div}_t\right) \qquad \textit{and} \qquad \log \text{PD}_{t+1}^{\text{BSLN}} \;=\; \ln\!\left(\text{Price}_{t+1}^{\text{BSLN}}/\text{Div}_{t+1}^{\text{BSLN}}\right)$$

are finite.

Part (c). For Repurchase Corp, $\text{Div}_{t+1}^{\text{REPC}} = 0$*, so*

$$\Delta \log \text{Div}_{t+1}^{\text{REPC}} \;=\; \ln(0/\text{Div}_t) \;=\; -\infty \qquad \textit{and} \qquad \log \text{PD}_{t+1}^{\text{REPC}} \;=\; \ln(\text{Price}_{t+1}^{\text{REPC}}/0) \;=\; +\infty$$

Repurchase Corp buys back $X/\text{Price}_{t+1}^{\text{REPC}}$ *shares. Each of the* $\#\text{Shares}_t$ *outstanding shares faces a termination probability of* $\omega \cdot X/(\#\text{Shares}_t^{\text{REPC}} \cdot \text{Price}_{t+1}^{\text{REPC}}) > 0$*.* □

Proof. **(Proposition 2.2)**

Part (a). Issuance NV pays $\text{Div}_{t+1}^{\text{NVIS}} > 0$ *and has* $\text{Price}_{t+1}^{\text{NVIS}} > 0$*, so both* $\Delta \log \text{Div}_{t+1}^{\text{NVIS}}$ *and* $\log \text{PD}_{t+1}^{\text{NVIS}}$ *are finite.*

Part (b). Let $I_{\text{NVIS}} > I_{\text{BSLN}}$ *denote the two firms' investment levels, with the difference financed by issuance:*

$$I_{\text{NVIS}} - I_{\text{BSLN}} \;=\; \Delta\#\text{Shares}_{t+1} \cdot \text{Price}_{t+1}$$

If the investment earns a return exceeding the cost of capital, then

$$V_{t+1}^{\text{NVIS}} - V_{t+1}^{\text{BSLN}} \;>\; \Delta\#\text{Shares}_{t+1} \cdot \text{Price}_{t+1}$$

The per-share price of Issuance NV is

$$\text{Price}_{t+1}^{\text{NVIS}} \;=\; \frac{V_{t+1}^{\text{NVIS}}}{\#\text{Shares}_t + \Delta\#\text{Shares}_{t+1}}$$

Combining these inequalities gives $\text{R}_{t+1}^{\text{NVIS}} > \text{R}_{t+1}^{\text{BSLN}}$ *for positive-NPV investment.*

Part (c). Total firm earnings grow faster for Issuance NV (more investment), but per-share earnings grow slower because the denominator increases:

$$\#\text{Shares}_t + \Delta\#\text{Shares}_{t+1} \;>\; \#\text{Shares}_t$$

Since per-share dividends are a function of per-share earnings,

$$\mathbb{E}_t[\Delta \log \text{Div}_{t+h}^{\text{NVIS}}] \;<\; \mathbb{E}_t[\Delta \log \text{Div}_{t+h}^{\text{BSLN}}] \qquad \textit{for } h \geq 2$$

□

Proof. **(Proposition 3.2)**

Consider a cap-rate decrease of $\xi > 0$*. The Gordon-model price moves from* $\text{Price}_0 = \mathbb{E}[\text{Div}_1] \times \left(\frac{1}{\text{r}-\text{g}}\right)$ *to* $\widetilde{\text{Price}}_0 = \mathbb{E}[\text{Div}_1] \times \left(\frac{1}{(\text{r}-\text{g})-\xi}\right)$*, so*

$$\Delta\text{Price}_0 \;=\; \mathbb{E}[\text{Div}_1] \cdot \left(\frac{1}{\text{r}-\text{g}}\right) \times \left[\frac{\xi}{(\text{r}-\text{g})-\xi}\right] \tag{A.1}$$

We equate this price change with $\frac{\mathbb{E}[\text{Div}_1]}{1+r}$*, cancel out the* $\mathbb{E}[\text{Div}_1]$*, and solve for* ξ*. This gives us the following result*

$$\xi \;>\; \frac{(r-g)^2}{(r-g)+(1+r)} \;=\; DY^2 \times \left(\frac{1}{1+r+DY}\right) \tag{A.2}$$

where the second equality uses $DY = (r-g)$*. Since* $r > 0\%$ *and* $DY \geq 0\%$*, we know that* $1 > \frac{1}{1+r+DY}$*. Thus,* $\xi > DY^2$ *is sufficient since* $DY^2 > DY^2 \cdot \left(\frac{1}{1+r+DY}\right)$*. For the S&P 500,* $DY \approx 2\%$*, so* $DY^2 \approx 0.04\%$*.*

The above considers a cap-rate decrease. For a cap-rate increase of the same magnitude, the analogous calculation yields the slightly tighter threshold $\xi > DY^2/(1+g)$*. Both exact thresholds are below* DY^2*, so the condition stated in the proposition is sufficient regardless of the direction of the perturbation.* □

Proof. **(Proposition 4.1)**

Both decompositions are derived from the same one-period realized-data formula (Equation 8), using the same linearization constants κ *and* ρ*. Both therefore have the same left-hand side,* $\log PD_t$*, and the same constant term,* $\kappa/(1-\rho)$*. Subtracting the rational-expectations version (Equation 13) from the subjective version (Equation 26) gives*

$$0 = \sum_{h=1}^{\infty} \rho^{h-1} \cdot \left\{ \begin{array}{r} \left(\widetilde{\mathbb{E}}_t[\Delta \log \text{Div}_{t+h}] - \widetilde{\mathbb{E}}_t[R_{t+h}]\right) \\ -\left(\mathbb{E}_t[\Delta \log \text{Div}_{t+h}] - \mathbb{E}_t[R_{t+h}]\right) \end{array} \right\} \tag{A.3a}$$

$$= \sum_{h=1}^{\infty} \rho^{h-1} \cdot \varepsilon_h \tag{A.3b}$$

establishing necessity. For sufficiency, if $\sum_{h=1}^{\infty} \rho^{h-1} \cdot \varepsilon_h = 0$*, then adding the error-aggregation condition to the rational-expectations decomposition recovers the subjective decomposition directly.* □

Proof. **(Proposition 4.2)**

The Campbell-Shiller formula in realized values is $R_{t+1} \approx \kappa + \Delta \log \text{Div}_{t+1} + \rho \cdot \log PD_{t+1} - \log PD_t$*. Under rational expectations* $\mathbb{E}_t[\cdot]$*, taking expectations gives*

$$\mathbb{E}_t[R_{t+1}] \;=\; \kappa + \mathbb{E}_t[\Delta \log \text{Div}_{t+1}] + \rho \cdot \mathbb{E}_t[\log PD_{t+1}] - \log PD_t \tag{A.4}$$

This holds because $\mathbb{E}_t[\cdot]$ *is integration under a probability measure, hence linear.*

Under subjective beliefs $\widetilde{\mathbb{E}}_t[\cdot]$, the analog would be

$$\widetilde{\mathbb{E}}_t[\mathrm{R}_{t+1}] \stackrel{?}{=} \kappa + \widetilde{\mathbb{E}}_t[\Delta \log \mathrm{Div}_{t+1}] + \rho \cdot \widetilde{\mathbb{E}}_t[\log \mathrm{PD}_{t+1}] - \log \mathrm{PD}_t \tag{A.5}$$

Substituting the distortions from Equation (31) into both sides:

$$\begin{aligned} LHS &= \mathbb{E}_t[\mathrm{R}_{t+1}] + \delta_{\mathrm{R}} \\ RHS &= \kappa + \mathbb{E}_t[\Delta \log \mathrm{Div}_{t+1}] + \delta_{\mathrm{Div}} + \rho \cdot \left(\mathbb{E}_t[\log \mathrm{PD}_{t+1}] - \delta_{\mathrm{PD}}\right) - \log \mathrm{PD}_t \\ &= \mathbb{E}_t[\mathrm{R}_{t+1}] + \delta_{\mathrm{Div}} - \rho\, \delta_{\mathrm{PD}} \end{aligned} \tag{A.6}$$

The last line uses Equation (A.4). Equality holds if and only if $\delta_{\mathrm{R}} = \delta_{\mathrm{Div}} - \rho\, \delta_{\mathrm{PD}}$.

Under a probability measure Q, both $\mathbb{E}[\cdot]$ and $\mathbb{E}_Q[\cdot]$ are linear. Since the Campbell-Shiller formula holds in realized values, taking expectations under Q preserves it, so the subjective version holds automatically and the adding-up condition is satisfied.

Under a belief distortion that is not derived from a measure, the three distortions δ_{R}, δ_{Div}, δ_{PD} are specified independently (or generated by a common nonlinear formula applied to each variable). The condition $\delta_{\mathrm{R}} = \delta_{\mathrm{Div}} - \rho\, \delta_{\mathrm{PD}}$ is a single linear equation relating three free parameters. It defines a two-dimensional hyperplane in $\mathbb{R}^3$, which is a set of measure zero in the space of possible distortion triples. Generic distortions violate it. □