Why People Disagree About What Drives Stock Prices*

Andrew Atkeson Jonathan Heathcote Fabrizio Perri
UCLA Federal Reserve Bank of Minneapolis

July 2026

Abstract

How much of stock-price variation reflects movements in fundamental value, as defined by Shiller (1981), rather than in expected returns? To a first-order approximation, the answer depends only on forecasts of expected returns: researchers using different cash-flow and value measures disagree only because different cash-flow-to-value ratios imply different return forecasts. Quantifying the predictability of cash-flow growth does not help answer Shiller's question. Valuations are driven primarily by expected cash flows unless expected returns vary substantially over very long horizons. Whether they do cannot be precisely identified from a century of annual data.

Keywords: Cash flows, Excess Volatility, Long Run Returns, Trading Strategies

JEL classification codes: G12, G32, G35

*The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. We would like to thank Stefan Nagel for a very helpful discussion of a previous draft of this paper and Charles Engel, Ellen McGrattan, Moto Yogo, Stavros Panageas, Yueran Ma, Andrea Eisfeldt, Stijn Van Nieuwerburgh, Zhengyang Jiang, Dan Greenwald, Sydney Ludvigson, and Martin Lettau for many helpful conversations on this topic.

1 Introduction

Shiller (1981) famously asked, "Do stock prices move too much to be justified by subsequent movements in dividends?" Since the work of Campbell and Shiller (1988a,b), this question has been reframed as a question about what drives the dynamics of cash flow to price ratios in the context of a first-order approximation to a present value identity. In this reframing, under the hypothesis that cash flow to price ratios are stationary, their movements should be driven either by changes in expected future returns or by changes in expected future cash flows.

Most analyses of cash flow to price ratio movements for the aggregate stock market use dividends-per-share and price-per-share for broad stock indices. These analyses have argued that most of the volatility of the ratio is driven by movements in future expected returns. See Cochrane (2011) and Campbell (2014, 2018) for summaries of this argument. But others, using alternative measures of value and cash flows, have found a larger role for cash flows, with the paper by Larrain and Yogo (2008) being a prominent early example of such research.[1]

Here, we aim to make a theoretical and an applied contribution to the question of what drives volatility in equity prices. Our theoretical point is that quantifying the extent to which stock prices are excessively volatile, as defined by Shiller (1981), requires directly quantifying time variation in future expected returns. Quantifying the predictability of cash flow growth for different cash flow measures does not help answer Shiller's question. Our applied point is that valuations will primarily be driven by fluctuations in expected cash flows unless there is significant time variation in expected returns over very long horizons. These conclusions are, to a first-order approximation, independent of the measures of value and cash flows used. Hence, we conclude that different researchers arrive at different conclusions about what drives the stock market because they have different forecasts of future long-horizon returns.

To understand these points, we first emphasize, as did Marsh and Merton (1986), that while there is a unique sequence of one-period total returns to investing in the aggregate equity market, this path for returns can be implemented by many possible paths for equity value and cash flow, each of which corresponds to a different way to split the same sequence of returns into a capital gain component and a cash flow component. Firms choose whether returns take the form of dividends or repurchase-generated capital gains; Miller and Modigliani (1961) show that returns are independent of this choice. Similarly, investors can choose whether or not to adjust their cash flow from equity investment in response to initial public offerings (IPOs) or myriad other corporate actions that change total equity market capitalization.

[1] This is a large literature surveyed in Koijen and Van Nieuwerburgh (2011). Other examples include Baker and Wurgler (2000), Boudoukh et al. (2007), Chen (2009), Van Binsbergen and Koijen (2010), Bansal and Yaron (2011), Chen, Da, and Priestley (2012), Eaton and Paye (2017) and Pruitt (2025).

In fact, investors can choose from an infinite number of possible “trading strategies”, each of which corresponds to a different sequence for the amount of funds invested relative to aggregate market capitalization. Different trading strategies correspond to different sequences for equity value P_t, cash flows D_t, and the cash flow to price ratio D_t/P_t. Does one's answer to the question of what drives the stock market depend on the trading strategy used in constructing one's data on price and cash flows?

To answer this question, following Shiller (1981), we define a *fundamental price* $P_t^\star$ as equal to the present value of expected future cash flows discounted at a constant rate. Our first result is that the ratio of this fundamental price to the observed price $P_t^\star/P_t$ is equivalent to a discounted sum of dollar-weighted deviations of expected returns from a constant benchmark divided by current price. Note that the trading strategy shows up in the ratio of fundamental price to observed price only through the dollar weights applied to future expected returns.

We next show that, to a first-order approximation, the dynamics of this ratio depend only on expected returns without consideration of their dollar weighting. With this simplification, the problem of estimating the dynamics of Shiller's fundamental price relative to observed price comes down to estimating the dynamics of the familiar discounted sum of expected log returns from a Campbell-Shiller log-linear approximation. Hence, given the results of Miller and Modigliani (1961) that total returns are not altered by the choice of trading strategy, estimates of the dynamics of $P_t^\star/P_t$ should also not be impacted by the trading strategy.

How does this result on $P_t^\star/P_t$ relate to familiar Campbell-Shiller-style decompositions of D_t/P_t? It is an identity that the ratio of cash flow to observed price, D_t/P_t, is equal to the ratio of fundamental price to observed price, $P_t^\star/P_t$, times the ratio of cash flow to fundamental price, $D_t/P_t^\star$, where the last term is the analogue to the expected dividend growth term in the standard Campbell-Shiller decomposition. While different trading strategies imply different sequences for D_t/P_t, the results above indicate that they should imply the same sequence (up to the first-order approximation) for $P_t^\star/P_t$. It follows that different trading strategies must imply different time paths for the expected cash flow growth term $D_t/P_t^\star$.

We conclude that quantifying whether fluctuations in cash flows account for a large or a small share of fluctuations in the dividend-price ratio under a particular trading strategy is not informative about what drives fluctuations in the price itself. Different trading strategies will imply different decompositions for the dynamics of the cash flow to price ratio, but those different strategies should imply the same path for expected returns – trading strategies do not alter returns – and thus approximately the same path for $P_t^\star/P_t$. If researchers working with different cash flow measures end up estimating different paths for $P_t^\star/P_t$, that must be because they are making different forecasts for expected returns. Thus, disagreements around what drives valuations can only be resolved by deciding which return forecasts are more reasonable and which are less reasonable. In making this point, we follow Nagel (2024).

These results lead to our third point: the question of the role of variation in expected returns in driving fluctuations in the ratio of fundamental price to observed price is really a question about the dynamics of expected returns over long horizons.[2]

To illustrate this point, we consider a model in which expected log real returns follow an AR(1) process. In this case, the dynamics of the log ratio of fundamental price to observed price are determined by two statistics. The first is a sequence of estimates of one-period ahead expected log real returns, $\mathbb{E}_t r_{t+1}$. The second is the persistence of these expected returns. .

We first show that one can obtain a remarkably high R^2 coefficient forecasting aggregate stock market returns one-year ahead with a cash flow to value ratio based on a broad measure of payouts to investors.[3] But we also show that fluctuations in this return predictor based on aggregate cash flow are quite transitory, suggesting only transitory fluctuations in expected returns, and thus only small fluctuations in our estimate of $\log(P_t^\star/P_t)$. These estimates therefore imply that expected cash flows are the dominant driver of stock valuations.

Next, we contrast these results to those obtained following the same procedure, but using as our return-predictor the ratio of dividends per share to price per share. This ratio is a much poorer predictor of one-year ahead returns than the ratio using our broader cash flow measure. But this per share dividend yield measure is much more persistent, and as a result, implies much larger fluctuations in expected long run returns, and in our estimate for $\log(P_t^\star/P_t)$.

While these two measures imply different degrees of persistence in their expected-return forecasts, neither identifies the key persistence parameter with much precision. As a result, we cannot definitively quantify the amount of excess volatility in stock prices.

Our results also have implications for interpreting the stock-market boom of the past 25 years. In the 1990s there was a persistent downward shift in the per-share dividend yield. One interpretation is that this signals persistently lower expected real returns. But the aggregate cash flow to value ratio did not decline over the same time period, cautioning against the view that booming valuations reflect lower future returns. Which signal is correct? Only time will tell, but the downward shift in the per share ratio mechanically reflects a trend break in net equity issuance — rising equity repurchases and declining new firm entry — that may be largely orthogonal to changes in expected returns.

[2]We see this as the central point of Larrain and Yogo (2008).

[3]See Baker and Wurgler (2000) and Boudoukh et al. (2007) and related papers.

2 A Decomposition of Asset Price Movements

Consider data on the price of some asset $\{P_t\}$ and the cash flows paid to the owner of that asset $\{D_t\}$. These data provide us with sequences of realized returns with dividends R^{wd}_{t+1} and without dividends R^{nd}_{t+1}:

$$R^{wd}_{t+1} \equiv \frac{P_{t+1} + D_{t+1}}{P_t} \qquad \text{and} \qquad R^{nd}_{t+1} \equiv \frac{P_{t+1}}{P_t}.$$

We wish to make use of the framework laid out in Shiller (1981) to estimate the drivers of stock price volatility. To do so, we study the following identity decomposing the asset price P_t into two components:

$$P_t = P^\star_t + \Phi_t, \tag{1}$$

where

$$P^\star_t \equiv \sum_{k=1}^{\infty} \beta^k \mathbb{E}_t D_{t+k} \tag{2}$$

for some $\beta < 1$. In this decomposition, $P^\star_t$ is the expected present value of dividends discounted at a constant rate β. We label $P^\star_t$ the "fundamental" component of price. We assume that β is small enough that the infinite sum in equation (2) is defined. The component Φ_t is simply the difference between the observed price P_t and the fundamental component $P^\star_t$. We label Φ_t the "residual" component of price.[4]

Our goal is to estimate the role of the fundamental price $P^\star_t$ in driving observed price P_t. In particular, we seek to estimate the dynamics of

$$\frac{P^\star_t}{P_t} = \sum_{k=1}^{\infty} \beta^k \mathbb{E}_t \frac{D_{t+k}}{P_t}. \tag{3}$$

If this ratio is constant over time, then we must have that fluctuations in observed price are proportional to fluctuations in fundamental price $P^\star_t$. To the extent that this ratio $P^\star_t/P_t$ does vary over time, then fluctuations in observed price P_t differ from fluctuations in fundamental price, with this difference accounted for by fluctuations in the residual term Φ_t/P_t.

To implement this estimation, we make use of the following proposition.

Proposition 1: Assume that for all t, $\lim_{k\to\infty} \beta^k \mathbb{E}_t \frac{P_{t+k}}{P_t} = 0$. Then the dynamics of the ratio of the fundamental price to observed price can be restated in terms of the dynamics of

[4]We discuss the interpretation of this residual Φ_t in Appendix A.

dollar-weighted expected returns as

$$\frac{P_t^\star}{P_t} = 1 + \sum_{k=0}^{\infty} \beta^{k+1} \mathbb{E}_t \left[\left(R_{t+k+1}^{wd} - \frac{1}{\beta} \right) \frac{P_{t+k}}{P_t} \right]. \tag{4}$$

Proof: See Appendix B.

3 Relationship to Campbell-Shiller Decompositions

Following Campbell and Shiller (1988a), there is a large literature in finance that examines the dynamics of cash flow to price ratios D_t/P_t, expected returns, and changes in future cash flows using a first-order approximation to the joint dynamics of these series. Our decomposition of the dynamics of price in equation (4) relates to this literature through the identity

$$\frac{D_t}{P_t} = \left(\frac{P_t^\star}{P_t} \right) \left(\frac{D_t}{P_t^\star} \right), \tag{5}$$

where $P_t^\star/P_t$ is given by equation (4) and

$$\frac{D_t}{P_t^\star} = \left(\sum_{k=1}^{\infty} \beta^k \mathbb{E}_t \frac{D_{t+k}}{D_t} \right)^{-1}. \tag{6}$$

Equation (5) gives an exact decomposition of the dynamics of the cash flow to price ratio into a component due to fluctuations in the ratio of fundamental price to observed price and a component due to expected growth of cash flows, where the latter is defined in equation (6). Taking logs of both sides of equation (5), we can see that the log-linear approximation made popular by Campbell and Shiller (1988a) and the many papers that follow is based on the claim that

$$\log \left(\frac{P_t^\star}{P_t} \right) \approx \sum_{k=0}^{\infty} \rho^k \mathbb{E}_t \left[r_{t+k+1}^{wd} - \bar{r} \right] \tag{7}$$

where $r_{t+1}^{wd} = \log(R_{t+1}^{wd})$ and ρ and $\bar{r}$ are constants of approximation.[5] Thus, to the extent that this first-order approximation in equation (7) is accurate, the dynamics of the log ratio of fundamental price to observed price are a function only of the dynamics of expected log returns without reference to the dollar-weighting of returns (other than through the choice of the constant ρ).[6]

[5]We derive this first-order approximation and these constants of approximation in Appendix C.

[6]We do not take a stand on whether this first-order approximation is sufficiently accurate to be useful. In Appendix F.1 we consider an example in which a firm returns cash to its investors entirely through share repurchases and we show that in this case, this approximation is not useful.

This observation brings us to our main result. To the extent that the first-order approximation in equation (7) is accurate, different researchers using different measures of value and cash flows corresponding to different trading strategies arrive at different conclusions as to what drives the stock market in terms of estimates of the dynamics of $P_t^\star/P_t$ because these researchers are arriving at different conclusions about the dynamics of expected returns. All other considerations are second-order.

Nagel (2024) in his equation (28) derives our equation (7) directly from a standard Campbell and Shiller (1988a) log-linear approximation to the dividend price ratio. He uses that equation to make the point that the measure of excess volatility on the right side of equation (7) depends only on the information set used to forecast returns. In contrast, he shows that the expected-return variance *share* in a Campbell–Shiller decomposition can be made close to anything by changing the cash flow measure in the denominator. Hence, he recommends reporting estimates of excess volatility in terms of estimates of the variance of expected returns in equation (7). We follow him in this respect. We next explore the difficulties with this estimation.

What do we know about the dynamics of expected returns? Unfortunately, there is considerable disagreement in the literature regarding the dynamics of expected returns, as discussed, for example, by Goyal and Welch (2008) and Goyal, Welch, and Zafirov (2023). Boudoukh, Richardson, and Whitelaw (2008) and others suggest that this disagreement is particularly severe for estimates of the volatility of long-horizon expected returns. Hence, there is considerable room for disagreement regarding the question of what drives the stock market.

4 Trading Strategies and Return Forecasting

One important reason for this disagreement is that different researchers are using different measures of cash flow and price, and are therefore forecasting returns using different sequences for the cash flow to price ratio, D_t/P_t. We now discuss why the data on price $\{P_t\}$, cash flow $\{D_t\}$, and the ratio $\{D_t/P_t\}$ are not uniquely defined for a given asset or portfolio of assets. Instead, these measures depend on what we refer to as the *trading strategy* which defines the amount invested at each date.

To understand this point, consider the sequences of portfolio value (which we call price) and cash flow recorded for an investor who follows a trading strategy of holding a time-varying fraction $\{S_t\}$ of total market capitalization of an asset, with S_t denoting the investor's holdings at the end of period t after cash flows are received. We let $\{\bar{P}_t\}$ and $\{\bar{D}_t\}$ be the sequences for portfolio value and cash flow recorded for an investor who holds a constant fraction $S_t = 1$ of the total market capitalization of the asset, and we denote the corresponding sequences of value and cash flow for an investor following trading strategy $S = \{S_t\}$ by $\{P_t(S)\}$ and $\{D_t(S)\}$ with

$$P_t(S) = S_t \bar{P}_t \tag{8}$$

and

$$D_t(S) = S_{t-1}\bar{D}_t + (S_{t-1} - S_t)\bar{P}_t. \tag{9}$$

These equations indicate that the trading strategy changes the time paths for P_t and D_t. But the time path for realized (and expected) total returns is invariant to the trading strategy, since

$$R^{wd}_{t+1}(S) = \frac{P_{t+1}(S) + D_{t+1}(S)}{P_t(S)} = \frac{S_{t+1}\bar{P}_{t+1} + S_t\bar{D}_{t+1} + S_t\bar{P}_{t+1} - S_{t+1}\bar{P}_{t+1}}{S_t\bar{P}_t} = \frac{\bar{P}_{t+1} + \bar{D}_{t+1}}{\bar{P}_t}.$$

Note, however, that the trading strategy *does* change the dynamics of the cash-flow-to-price valuation ratio, $D_t(S)/P_t(S)$. That is important, because these ratios are commonly used to forecast returns. In particular,

$$\begin{aligned} \frac{D_t(S)}{P_t(S)} &= \frac{S_{t-1}}{S_t}\frac{\bar{D}_t}{\bar{P}_t} + \frac{S_{t-1} - S_t}{S_t}, & (10)\\ &= \frac{\bar{D}_t}{\bar{P}_t} - s_t\left(1 + \frac{\bar{D}_t}{\bar{P}_t}\right), & (11) \end{aligned}$$

where $s_t = (S_t - S_{t-1})/S_t$ denotes the growth rate for S_t.

Thus, if the trading strategy implies a rising ownership share ($s_t > 0$), then the cash flow to price ratio will be smaller than the corresponding ratio with a constant ownership share.[7] One can also show that when the trading strategy features $s_t > 0$, expected growth in cash flows for that strategy will exceed expected cash flow growth given a constant ownership share. That is the natural counterpart to the cash flow to price ratio being smaller.

To illustrate these points, suppose the asset in question is a diversified portfolio of all equities in the CRSP Value-Weighted Total Market Index, and consider two alternative trading strategies. The first, which we label the strategy for an "aggregate investor" corresponds to maintaining a constant fraction of the aggregate market capitalization of the stocks in the CRSP index at the start of each period. The data on the value and cash flows associated with this trading strategy are $\{\bar{P}_t\}$ and $\{\bar{D}_t\}$.

The second trading strategy we consider corresponds to that of a passive index investor. We refer to this trading strategy as that of a "per-share investor." The per-share investor

[7] See Appendix F for more details.

holds a time-varying fraction $\{S_t\}$ of the total market capitalization of the CRSP index, where S_t is equal to the ratio of the level of this index to the total market capitalization of the stocks in the index at the end of period t. The inverse of this ratio $1/S_t$ is referred to as the *index divisor.* The per-share investor's ownership share S_t will change every day as a result of what are called "corporate actions." We let $\{\tilde{P}_t\}$ denote *price per share* for the per-share investor, and $\{\tilde{D}_t\}$ denote *dividends per share.*

There are many different sorts of corporate actions that lead the level of the index to diverge from the total market capitalization of the stocks in the index, but two important examples are Initial Public Offerings (IPOs) and equity repurchases. Following an IPO, S_t for an index investor declines: if a newly public firm accounts for one percent of aggregate market capitalization, the index investor sees her share of aggregate market capitalization decline by approximately one percent because the level of the index does not change with the IPO. Conversely, if firms in the index collectively repurchase one percent of publicly-held equity, the index investor sees her share of aggregate market capitalization S_t rise by approximately one percent because, again, the level of the index does not change with this corporate action. In both cases, it is the divisor of the index that is revised.[8]

Much of the literature examining the dynamics of the stock market has used the data on price per share and dividends per share for a per-share investor. But others, including Larrain and Yogo (2008) and Pruitt (2025), have used broader measures of value and cash flows closer to $\bar{P}_t$ and $\bar{D}_t$. And other researchers have considered intermediate cases, for example adjusting dividends to incorporate the impact of repurchases and other corporate actions of incumbent firms, but excluding the effects of firm entry and exit.

In the remainder of this paper we use data on value and cash flows for an aggregate investor and for a per-share investor to illustrate the dramatic difference in return forecasts that one obtains with these alternative measures of the ratio of cash flow to value.

5 Data from Two Trading Strategies

We use data from the CRSP Value-Weighted Annual File on WRDS for all three exchanges (NYSE, Amex, and NASDAQ) from 1929-2025 to construct series for value and cash flow for a per-share and an aggregate investor.[9]

For the per-share investor, we use the return series with and without dividends directly to construct our measures of price per share $\tilde{P}_t$ and dividends per share $\tilde{D}_t$.

For the aggregate investor, value $\bar{P}_t$ equals Total Market Capitalization TMC_t. We follow

[8]We discuss the details of index construction in Appendix D.

[9]Further details are provided in Appendix E

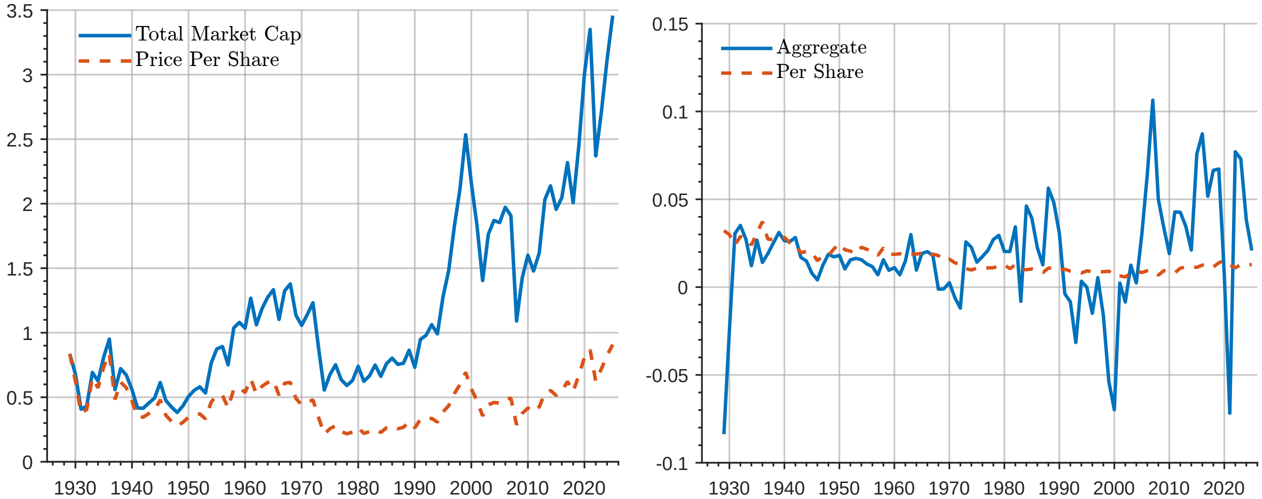

(a) End-of-year total market capitalization divided by nominal Personal Consumption Expenditure (PCE) (solid) and price per share divided by PCE (dashed). Price per share in 1929 is normalized to equal Total Market Capitalization in that year. We scale by PCE as a way to strip out the impact of trend growth, which improves visual legibility.

(b) Total cash flows to an aggregate investor, $\bar{D}_t$, relative to aggregate PCE (solid) and dividends per share, $\tilde{D}_t$, relative to aggregate PCE (dashed). Scaling of these variables by PCE is for legibility.

Figure 1: Aggregate versus per-share valuation and cash-flow measures.

the methodology in Dichev (2007) to recover cash flow $\bar{D}_t$. In particular, because the trading strategy does not impact realized returns, cash flow for the aggregate investor is the solution to

$$R_t^{wd} = \frac{TMC_t + \bar{D}_t}{TMC_{t-1}}. \tag{12}$$

The differences in the evolution of total market capitalization $\bar{P}_t$ and price per share $\tilde{P}_t$ for the CRSP Total Market Value-Weighted Index are illustrated in Figure 1a. There we plot the ratio of each series to nominal aggregate Personal Consumption Expenditure (PCE). The two series co-move, but there is much more long term growth in the market capitalization series.

Figure 1b shows total cash flows for the aggregate investor (solid) and dividends for a per-share investor (dashed). Dividends per share are very smooth. In contrast, cash flow to an aggregate investor is quite volatile and often negative.[10] Negative cash flow reflects net equity issuance, and is often associated with a wave of entry of firms into public markets.

From equations (8) and (9), we know that the differences in the dynamics of value and

[10]See also Davydiuk et al. (2023) for a similar calculation of aggregate cash flow to equity with similar results regarding the volatility of these cash flows.

cash flows for the aggregate investor versus the index investor that are apparent in Figure 1 reflect the dynamics of the trading strategy S_t for the index investor. In Figure 2a, we show the log of the ratio of price per share to total market capitalization corresponding to $\log(S_t)$.

One can see that from 1929 through 1999, $\log(S_t)$ declined steadily over time. The slope of the trendline in this figure from 1929 to 1999 is -1.88%. More recently, $\log(S_t)$ appears to level off, or even increase, reflecting a combination of increased share repurchases by incumbent firms and a reduced rate of net entry of firms into public markets. The slope of the trendline in this figure from 2000 through 2025 is 0.24%.

In Figure 2b, we see that this change in trend in the per share trading strategy coincided with a significant increase in trend growth in real dividends per share.

Figure 2c plots the dynamics of the ratio of cash flow to value for an aggregate investor, while Figure 2d shows the ratio of dividends per share to price per share. The dynamics of these two cash flow to value ratios are quite different, at both high and low frequencies. The ratio of cash flow to value for an aggregate investor has large transitory fluctuations. But it appears fairly stable at long horizons. Its mean over the period 1929-1999 is 2.14%, compared to 1.78% over the period 2000-2025. In contrast, the ratio of dividends per share to price per share shows large changes at low frequencies. Its mean over the period 1929-1999 is 4.2%, while it is only 1.92% over the period 2000-2025. This stark difference between the low frequency dynamics of the cash flow to value ratio for an aggregate and an index investor reflects the two percentage points change in the trend growth rate of S_t shown in Figure 2a.[11]

To understand this, note that until 2000, new net equity issuance every year – by existing firms and via IPOs – was equal to about two percent of market capitalization. Given $s_t = -0.0188$, equation (11) indicates that the per share dividend yield $\tilde{D}_t/\tilde{P}_t$ was approximately two percent higher than the aggregate yield. The per share yield was boosted relative to the aggregate yield, because the per-share investor anticipated weaker cash flow growth, as her share of total market ownership was gradually diluted. Around 2000, net equity issuance dropped to near zero and remained there for the past 25 years, with equity issuance henceforth almost exactly offset by equity withdrawal via share repurchases. Equation (11) then accounts for why the per share yield then fell by about two percent to converge closer to the aggregate yield.

[11] In Appendix G, we show that the trading strategy and cash flow to price ratios for an aggregate and a per-share investor in the S&P 500 show similar patterns.

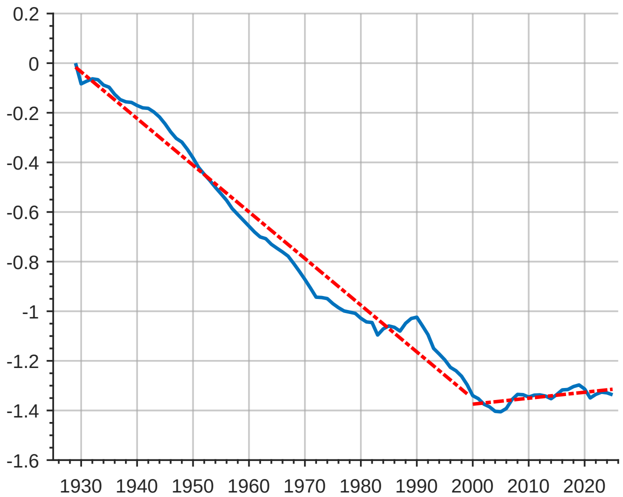

(a) Log ratio of price per share to total market capitalization corresponding to the index investor trading strategy $\{\log(S_t)\}$. Dotted lines are linear trends for 1929–1999 and 2000–2025.

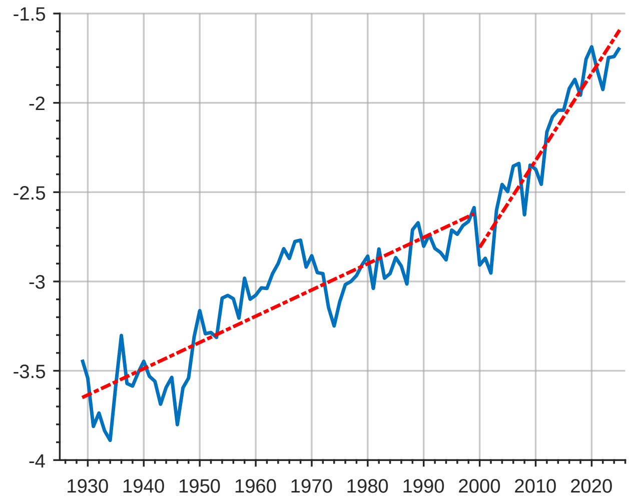

(b) Log of real dividends per share (deflated by the CPI). Dotted lines are linear trends for 1929–1999 and 2000–2025.

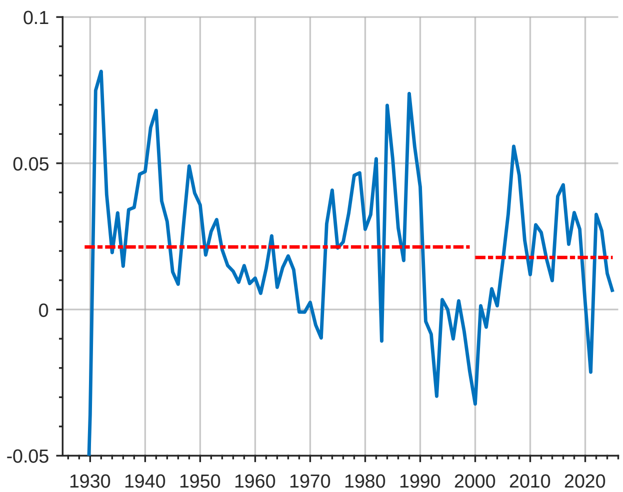

(c) Cash flow to price for an aggregate investor, $\bar{D}_t/\bar{P}_t$. Dotted lines are means for 1929–1999 and 2000–2025.

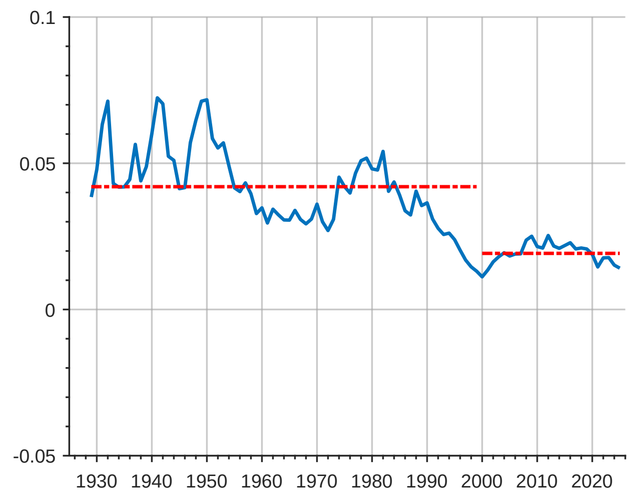

(d) Cash flow to value for a per-share investor, $\tilde{D}_t/\tilde{P}_t$. Dotted lines are means for 1929–1999 and 2000–2025.

Figure 2: Trading strategies and cash-flow ratios.

6 Estimating $P_t^\star/P_t$

We take a simple approach to estimating the dynamics of the ratio of fundamental price to observed price in equation (7). We base our estimation on the observation that if one-period ahead expected real returns $\mathbb{E}_t\left[r_{t+1}^{wd} - \bar{r}\right]$ follow an AR(1) process with persistence ψ, then equation (7) can be written

$$\log\left(\frac{P_t^\star}{P_t}\right) \approx \frac{1}{1-\rho\psi}\mathbb{E}_t\left[r_{t+1}^{wd}-\bar{r}\right]. \tag{13}$$

If equation (13) holds, then to estimate the dynamics of fundamental price to observed price, we need only estimate a series for one-period ahead expected log real returns, $\mathbb{E}_t\left[r_{t+1}^{wd}-\bar{r}\right]$ and the persistence of those expected real returns, ψ. We set the constant of approximation ρ to 0.96 as is standard in the literature.[12]

We next assume that investors use some measure of cash flow relative to value to forecast log real returns. We will compare and contrast forecasts based on two versions of that ratio corresponding to the two trading strategies we have discussed: the ratio for the aggregate investor, and the ratio for the per-share investor. The hypothesis that the ratio of cash flows to price for the aggregate investor is a good proxy for expected returns goes back to the findings of Baker and Wurgler (2000) and Boudoukh et al. (2007). Many others have used the per share version of the ratio to forecast returns.

We therefore run two OLS forecasting regressions of the form

$$r_{t+1}^{wd} = \alpha_A + \gamma_A\frac{\bar{D}_t}{\bar{P}_t} + \epsilon_{A,t+1}, \tag{14}$$

$$r_{t+1}^{wd} = \alpha_P + \gamma_P\frac{\tilde{D}_t}{\tilde{P}_t} + \epsilon_{P,t+1}, \tag{15}$$

where r_{t+1}^{wd} denotes the log CRSP return net of growth in the Consumer Price Index between year t and year $t+1$, and where we use the subscripts A and P to denote the coefficients for the aggregate and per share regressions. Results are reported in Table 1 with more details in Appendix I.

We use these two regressions to construct two alternative time series for $\log(P_t^\star/P_t)$. In each case, we use the corresponding regression coefficient estimates to construct a time series for $\mathbb{E}_t\left[r_{t+1}^{wd}-\bar{r}\right]$. That series is then substituted into equation (13), where we set the persistence parameter ψ equal to the persistence of the aggregate or per-share dividend-price ratio, depending on the predictor being used.

Given an estimated time series for $\log(P_t^\star/P_t)$, we can immediately construct a corresponding time series for the aggregate fundamental price:

$$\log\bar{P}_t^\star = \log(P_t^\star/P_t) + \log\bar{P}_t. \tag{16}$$

[12]We examine the sensitivity of our estimates to ρ in Appendix K.

6.1 Estimation Results

The estimated series for $\log(P_t^\star/P_t)$ from this procedure are shown in Figure 3a. Note that the two series are simply rescaled versions of the corresponding valuation ratios, D_t/P_t, where the scaling factor is equal to $\gamma/(1-\rho\psi)$. The plot indicates more volatility in the log of $P_t^\star/P_t$ when per share data is used to forecast returns. And that series indicates a step down in $\log(P_t^\star/P_t)$ in the 1990s, mirroring the step down in the $\tilde{D}_t/\tilde{P}_t$ series visible in Figure 2d. Thus, the per share data points to a persistent decline in the discounted sum of future expected returns, and an associated decline in $P_t^\star/P_t$ of around 40 log points. Conversely, when the aggregate ratio is used to forecast returns, the implied ratio of $P_t^\star/P_t$ fluctuates around its historical mean in the post 2000 period, reflecting the fact that the aggregate valuation ratio $\bar{D}_t/\bar{P}_t$ is also near its historical mean.

Figure 3b shows the ratio of aggregate market capitalization relative to PCE (dotted), against two estimated series for the corresponding fundamental price $\bar{P}_t^\star$ also scaled by PCE. The solid series is constructed using the series for $\log(P_t^\star/P_t)$ estimated from aggregate data, while the dashed series uses the series for $\log(P_t^\star/P_t)$ estimated using per share data. In both cases, a large portion of observed fluctuations in valuations are attributed to time variation in the fundamental price. But a key distinction between the two models is apparent in recent decades. When the aggregate yield is used to forecast returns, almost the entire run-up in valuations since the 1980s is attributed to growth in the fundamental component of value. In contrast, when the per share yield is used a sizable gap opens up between observed valuations and the estimated fundamental component, where this gap reflects low expected returns going forward.

We see that the two series in Figure 3a differ substantially, even though $\log(P_t^\star/P_t)$ is invariant to the trading strategy under the first-order approximation in equation (7) when the return forecast is held fixed. There is no contradiction. The two series here are constructed from different return forecasts and hence yield different estimates of $\log(P_t^\star/P_t)$. We explore this point in greater detail next.

6.2 Why Different Trading Strategies Give Different Results

Why are the two paths for the fundamental price in Figure 3 so different? First, the explanation for why the dashed series for $\log(P_t^\star/P_t)$ estimated with per share data is more volatile than the solid series estimated with aggregate data is *not* that the per share dividend-price ratio generates more volatility in one-period-ahead expected log returns, $\mathbb{E}_t r_{t+1}^{wd}$. To the contrary, the per share ratio is a poor predictor of log returns one year ahead: the R^2 from regression (15) is only 0.026 (see Table 1). Rather, the high volatility reflects the fact that

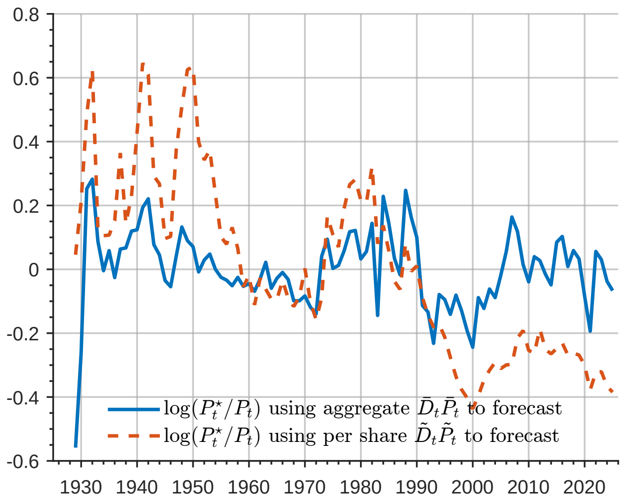

(a) Estimate of $\log(P_t^\star/P_t)$ using D_t/P_t to forecast returns for an aggregate investor (solid) and a per-share investor (dashed).

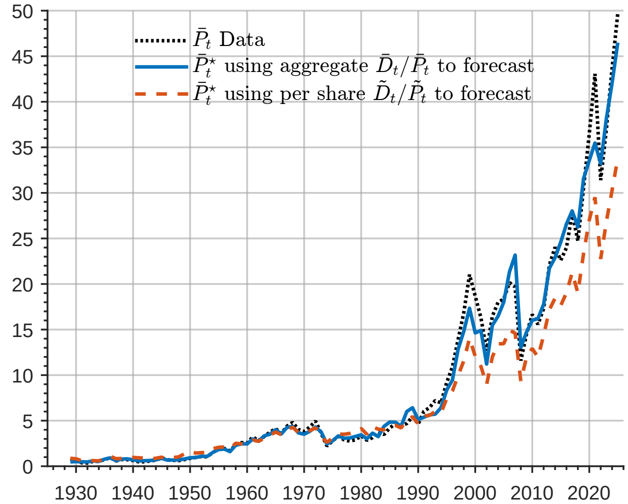

(b) The ratio of total market capitalization (dotted) and the estimated fundamental price for an aggregate investor (solid) and a per-share investor (dashed), all relative to aggregate personal consumption expenditure (PCE). Scaling by PCE is for legibility.

Figure 3: Estimated deviations from fundamental valuation.

the per share dividend price ratio is very persistent, with an estimated autocorrelation ψ of 0.923. Small movements in expected returns that are expected to be very persistent generate large deviations in equilibrium valuations relative to the expected present value of future cash flows.

When the aggregate cash flow to price ratio is used to forecast returns, in contrast, the model is much more successful at forecasting near-term returns, consistent with prior findings in the literature. Regression (14) for predicting log one-year real returns yields an impressive $R^2 = 0.119$ But the aggregate cash flow to price ratio $\bar{D}_t/\bar{P}_t$ is much less persistent than the per share version, with an autocorrelation ψ of only 0.481. Because movements in expected returns are expected to be relatively transitory in this specification, they do not generate large movements in valuations.

To better understand the role of persistence, note that the variance of $\log(P_t^\star/P_t)$ can be expressed in terms of the R^2 of the regression as:

$$\mathrm{Var}\left(\log\frac{P_t^\star}{P_t}\right) = \left(\frac{1}{1-\rho\psi}\right)^2 \times R^2 \times \mathrm{Var}(r_{t+1}^{wd}). \tag{17}$$

This expression indicates that the implied variance of $\log(P_t^\star/P_t)$ is much more sensitive to ψ than to the regression fit statistic. In particular, given $\rho = 0.96$, an increase in persistence ψ from 0.481 to 0.923 is equivalent to reducing the R^2 for the one period return forecast by

Table 1: Return forecasts and implied valuation volatility

	R^2	ψ OLS	ψ bias corrected	$\mathrm{Std}\left(\log \frac{P_t^\star}{P_t}\right)$
Aggregate $\bar{D}_t/\bar{P}_t$	0.119	0.481	0.509	0.123
95% confidence int.	(0.012 , 0.220)	(0.268 , 0.619)	n.a.	(0.037 , 0.173)
Per-share $\tilde{D}_t/\tilde{P}_t$	0.026	0.923	0.965	0.277
95% confidence int.	(0.001 , 0.125)	(0.754 , 0.960)	n.a.	(0.052 , 0.551)

Column (1) reports R^2 from one-year-ahead log return forecasting regressions (equations 14 and 15). Column (2) reports the OLS estimate of the predictor autocorrelation ψ. Column (3) reports the bootstrap bias-corrected estimate of the predictor autocorrelation ψ. Column (4) reports the standard deviation of the estimated $\log(P_t^\star/P_t)$. Column (4) is computed from the OLS estimate in column (2) via equation (17). Bootstrap 95 percent confidence intervals where reported are below corresponding estimates. See online Appendix J for the bootstrap methodology.

a factor of 22.3! Simply put, large but transitory fluctuations in expected returns have a minimal impact on equilibrium valuations.[13]

The bootstrap confidence intervals reported in Table 1 indicate that the persistence parameters ψ are estimated with considerable uncertainty. And yet the 95% intervals for the estimates of ψ for the aggregate predictor and the per-share predictor do not overlap: the two conditioning variables have genuinely different persistence. Propagated through equation (17), however, the implied 95% intervals for $\mathrm{sd}(\log P_t^\star/P_t)$ overlap substantially. The magnitude of the volatility of $\log\left(P_t^\star/P_t\right)$ is thus far less well identified than the difference in persistence between the two predictors. With fewer than one hundred annual observations, the data are consistent both with a world in which prices deviate little from fundamental value and with one in which they deviate substantially.

Why is the per share dividend-price ratio so much more persistent than the aggregate one? One reason is that firms deliberately smooth dividend payments. But another reason is that the trend break around 2000 in the path for the per-share investor's share of aggregate market capitalization translates into a very persistent level shift downward in the per share ratio, as discussed in Section 5. Suppose this shift reflects an increase in the rate of share repurchases (a decline in the rate of equity issuance via existing firms or via IPOs would work similarly). Then one would expect a decline in dividends per share (equation 11), but faster growth in dividends per share moving forward. Figure 2b illustrates precisely this dynamic for dividends per share. We take that as a signal that one should interpret the downward shift observed in the series for dividends per share relative to price per share as reflecting a persistent change in the rate of net equity issuance and thus in the expected cash flow growth

[13]In Online Appendix L we consider a VAR specification in which both the aggregate cash flow to price ratio and the per share version are used to forecast returns.

rate for a per share investor, rather than as a signal of lower expected returns.

In Appendix H and Appendix I we consider an alternative trading strategy that is identical to the per-share strategy, except that under this alternative trading strategy the investor maintains the same trend decline in $\log(S_t)$ post 2000 as in the pre-2000 period. We show that fluctuations in the dividend-price ratio under this trading strategy successfully predict 32 percent of the realized variance in future 10-year returns. And yet the implied volatility of $\log(P_t^\star/P_t)$ is smaller than this strong predictability might suggest. The reason is that the autocorrelation of the dividend-price ratio under this strategy is only 0.840, so expected returns revert most of the way to their mean within a decade. This illustrates that documenting significant predictability in 10-year returns does not suffice to answer Shiller's question about what drives the stock market. In Appendix M we show that the persistence of the return predictor remains critical for quantifying the impact of fluctuations in expected returns on valuations even under the Cochrane (2008) view that fluctuations in dividend-price ratios entirely reflect fluctuations in expected returns.[14]

7 Conclusion

We have argued that to answer Shiller (1981)'s question of whether stock prices move too much to be justified by expected movements in cash flows, one must measure the extent to which the discounted value of dollar-weighted expected returns fluctuates relative to its unconditional mean. To a first-order approximation, this estimation problem can be simplified to ignore the effect of dollar-weighting of returns. Thus, to a first-order approximation, the estimate that one obtains for the dynamics of the ratio of fundamental price to observed price should not depend on the trading strategy used to define cash flows and price.

Our second point is more quantitative. We argued, in a simple model for estimation, that the determinants of the volatility of the ratio of fundamental price to observed price, $P_t^\star/P_t$, are the volatility of expected returns one period ahead and the persistence of those expected returns, with the persistence being much the more important of these two parameters. In particular, we argue that for movements in expected returns to generate significant movements in valuations, they must involve substantial movements in expected returns over very long horizons. Given fewer than 100 years of stock return data, it is very difficult to settle

[14] We have focused on alternative measures of dividend-price ratios as return forecasting variables. Some have instead advocated for using corporate earnings as a cash flow measure. In Atkeson, Heathcote, and Perri (2026) we show that a trend decline in labor's share of value-added coupled with relatively stable investment has translated into much faster growth in corporate free cash flow than in corporate earnings in recent decades, with the implication that the free cash flow yield has remained relatively stable even as the earnings yield has declined.

the question of whether investors expect enough variation in returns over very long horizons for that variation to be the primary driver of stock market fluctuations.

How might research into this question of what drives the stock market progress? It seems likely that information beyond aggregate time series is needed. Cross-sectional present-value methods that pool information across many portfolios (Kelly and Pruitt (2013)) exploit the restriction that a common discount-rate factor prices many assets at once. The term structure of equity claims, recovered from dividend strips and dividend futures (for example Van Binsbergen, Brandt, and Koijen (2012) and Van Binsbergen et al. (2013)), prices expected returns horizon by horizon and so speaks directly to the persistence of the expected-return process. Long international panels (for example Rangvid, Schmeling, and Schrimpf (2014) and Jordà et al. (2019)) add cross-country variation in ψ. Each of these could sharpen the estimate of long-horizon expected-return variation far more than several additional decades of U.S. data.

References

Andrews, Donald W. K. 1993. "Exactly Median-Unbiased Estimation of First Order Autoregressive/Unit Root Models." *Econometrica* 61 (1):139–165.

Atkeson, Andrew, Jonathan Heathcote, and Fabrizio Perri. 2026. "A Macroeconomic Perspective on Stock Market Valuation Ratios." Working Paper 34748, NBER.

Baker, Malcolm and Jeffrey Wurgler. 2000. "The Equity Share in New Issues and Aggregate Stock Returns." *Journal of Finance* 55 (5):2219–2257.

Bansal, Ravi and Amir Yaron. 2011. "The Asset Pricing-Macro Nexus and Return-Cash Flow Predictability." Tech. Rep. 1786144, SSRN.

Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael R. Roberts. 2007. "On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing." *Journal of Finance* 62 (2):877–915.

Boudoukh, Jacob, Matthew Richardson, and Robert F. Whitelaw. 2008. "The Myth of Long-horizon Predictability." *The Review of Financial Studies* 21 (4):1577–1605.

Campbell, John Y. 1991. "A Variance Decomposition for Stock Returns." *The Economic Journal* 101 (405):157–179.

———. 2014. "Empirical Asset Pricing: Eugene Fama, Lars Peter Hansen, and Robert Shiller." *Scandinavian Journal of Economics* 116 (3):593–634.

———. 2018. *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton University Press.

Campbell, John Y. and Robert J. Shiller. 1987. "Cointegration and Tests of Present Value Models." *Journal of Political Economy* 95 (5):1062–1088.

———. 1988a. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *The Review of Financial Studies* 1 (3):195–228.

———. 1988b. "Stock Prices, Earnings, and Expected Dividends." *Journal of Finance* 43 (3):661–676.

Chen, Long. 2009. "On the reversal of return and dividend growth predictability: A tale of two periods." *Journal of Financial Economics* 92 (1):128–151.

Chen, Long, Zhi Da, and Richard Priestley. 2012. "Dividend Smoothing and Predictability." *Management Science* 58 (10):1834–1853.

Cochrane, John. 2008. "The Dog That Did Not Bark: A Defense of Return Predictability." *The Review of Financial Studies* 21 (4):1533–1575.

———. 2011. "Presidential Address: Discount Rates." *Journal of Finance* 66 (4):1047–1108.

Davydiuk, Tetiana, Scott Richard, Ivan Shaliastovich, and Amir Yaron. 2023. "How Risky Are U.S. Corporate Assets?" *Journal of Finance* 78 (1):141–208.

Dichev, Ilia D. 2007. "What Are Stock Investors' Actual Historical Returns? Evidence from Dollar-Weighted Returns." *American Economic Review* 97 (1):386–401.

Eaton, Gregory W. and Bradley S. Paye. 2017. "Payout Yields and Stock Return Predictability: How Important Is the Measure of Cash Flow?" *The Journal of Financial and Quantitative Analysis* 52 (4):1639–1666.

Goyal, Amit and Ivo Welch. 2008. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies* 21 (4):1455–1508.

Goyal, Amit, Ivo Welch, and Athanasse Zafirov. 2023. "A Comprehensive 2022 Look at the Empirical Performance of Equity Premium Prediction." Research Paper 21-85, Swiss Finance Institute.

Jagannathan, Ravi and Binying Liu. 2019. "Dividend Dynamics, Learning, and Expected Stock Index Returns." *The Journal of Finance* 74 (1):401–448.

Jordà, Òscar, Katherina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor. 2019. "The Rate of Return on Everything, 1870-2015." *Quarterly Journal of Economics* 134 (3):1225–1298.

Kelly, Bryan and Seth Pruitt. 2013. "Market Expectations in the Cross-Section of Present Values." *The Journal of Finance* 68 (5):1721–1756.

Kendall, M. G. 1954. "Note on Bias in the Estimation of Autocorrelation." *Biometrika* 41 (3-4):403–404.

Koijen, Ralph S.J. and Stijn Van Nieuwerburgh. 2011. "Predictability of Returns and Cash Flows." *Annual Review of Financial Economics* 3:467–491.

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 (1):200–226.

Lettau, Martin and Stijn Van Nieuwerburgh. 2008. "Reconciling the Return Predictability Evidence." *Review of Financial Studies* 21 (4):1607–1652.

Marsh, Terry A. and Robert C. Merton. 1986. "Dividend Variability and Variance Bounds Tests for the Rationality of Stock Market Prices." *American Economic Review* 76 (3):483–498.

Miller, Merton H. and Franco Modigliani. 1961. "Dividend Policy, Growth, and the Valuation of Shares." *Journal of Business* 34 (4):411–433.

Nagel, Stefan. 2024. "Stock Market Valuation: Explanations, Non-Explanations, and Some Open Questions."

Pruitt, Seth. 2025. "Dogs and Cats Living Together: A Defense of Cash-Flow Predictability." Working Paper 4343101, SSRN.

Rangvid, Jesper, Maik Schmeling, and Andreas Schrimpf. 2014. "Dividend Predictability Around the World." *Journal of Financial and Quantitative Analysis* 49 (5/6):1596–1618.

Shiller, Robert J. 1981. "Do Stock Prices Move Too Much to be Justified by Subsequent Movements in Dividends?" *American Economic Review* 71 (3):421–436.

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

Van Binsbergen, Jules, Wouter Hueskes, Ralph S.J. Koijen, and Evert Vrugt. 2013. "Equity Yields." *Journal of Financial Economics* 110 (3):503–519.

Van Binsbergen, Jules and Ralph S.J. Koijen. 2010. "Predictive Regressions: A Present-Value Approach." *Journal of Finance* 65 (4):1439–1471.

Internet Appendices

A Interpretation of Residual Component of Value

Here we discuss the interpretation of the residual component of value Φ_t.

Let $M_{t,t+k}$ denote the stochastic discount factor used to discount between periods t and $t+k$. The observed price P_t can then be expressed as

$$P_t = \sum_{k=1}^{\infty} \mathbb{E}_t \left[M_{t,t+k} D_{t+k} \right] + \Theta_t,$$

where the term Θ_t reflects all other influences on price other than the expected discounted present value of dividends. Given equations (1) and (2), the residual component of price is given by

$$\Phi_t = \sum_{k=1}^{\infty} \mathbb{E}_t \left[\left(M_{t,t+k} - \beta^k \right) D_{t+k} \right] + \Theta_t.$$

Thus, the residual component of price captures fluctuations in discounting done by $M_{t,t+k}$ relative to the constant rate β^k and all other influences on price as captured by Θ_t.

B Proof of Proposition 1:

We present two proofs of Proposition 1. The first proof uses direct substitution of each term in the definition of the ratio of fundamental price to observed price. The second proof uses an argument similar to that in Campbell and Shiller (1987).

Our first proof of Proposition 1 begins as follows.

First, observe that

$$D_{t+k+1} = R^{wd}_{t+k+1} P_{t+k} - P_{t+k+1}$$

This gives us that for any finite K

$$\sum_{k=0}^{K} \beta^{k+1} \frac{D_{t+k+1}}{P_t} = \sum_{k=0}^{K} \beta^{k+1} R^{wd}_{t+k+1} \frac{P_{t+k}}{P_t} - \sum_{k=0}^{K} \beta^{k+1} \frac{P_{t+k+1}}{P_t}$$

Taking the conditional expectation of both sides of this equation gives us

$$\sum_{k=0}^{K} \beta^{k+1} \mathbb{E}_t \frac{D_{t+k+1}}{P_t} = \sum_{k=0}^{K} \beta^{k+1} \mathbb{E}_t R^{wd}_{t+k+1} \frac{P_{t+k}}{P_t} - \sum_{k=0}^{K} \beta^{k+1} \mathbb{E}_t \frac{P_{t+k+1}}{P_t}$$

Adding and subtracting 1 from the right side of this equation gives

$$\sum_{k=0}^{K} \beta^{k+1} \mathbb{E}_t \frac{D_{t+k+1}}{P_t} = 1 + \sum_{k=0}^{K} \beta^{k+1} \mathbb{E}_t \left(R^{wd}_{t+k+1} - \frac{1}{\beta} \right) \frac{P_{t+k}}{P_t} - \beta^{K+1} \mathbb{E}_t \frac{P_{t+K+1}}{P_t}$$

Taking limits of both sides of this equation as $K \to \infty$ and imposing our assumption that the sum defining $P_t^\star$ on the left side of this equation converges together with our assumption that $\lim_{K\to\infty} \beta^{K+1}\mathbb{E}_t \frac{P_{t+K+1}}{P_t} = 0$ gives us our desired result.

Our second proof follows from application of the Law of Iterated Expectations. We have

$$\beta \mathbb{E}_t P_{t+1}^\star - P_t^\star = \mathbb{E}_t \sum_{k=2}^{\infty} \beta^k \mathbb{E}_{t+1} D_{t+k} - \sum_{k=1}^{\infty} \beta^k \mathbb{E}_t D_{t+k} = -\beta \mathbb{E}_t D_{t+1}$$

where the last equality follows from the Law of Iterated Expectations. Given our definition of Φ_t, we then have

$$\beta \mathbb{E}_t P_{t+1} - P_t = \beta \mathbb{E}_t \Phi_{t+1} - \Phi_t - \beta \mathbb{E}_t D_{t+1}$$

This gives us that

$$\beta \mathbb{E}_t \Phi_{t+1} - \Phi_t = \beta \mathbb{E}_t \left(R_{t+1}^{wd} - \frac{1}{\beta} \right) P_t$$

Now observe that if the infinite sum defining $P_t^\star$ converges, then, by the Law of Iterated Expectations,

$$\lim_{K\to\infty} \beta^K \mathbb{E}_t P_{t+K}^\star = 0$$

This observation together with our assumption that $\lim_{K\to\infty} \beta^{K+1}\mathbb{E}_t P_{t+K+1} = 0$ implies $\lim_{k\to\infty} \mathbb{E}_t \Phi_{t+k} = 0$. Hence, solving this difference equation for the residual component of price forward gives us that

$$-\Phi_t = \sum_{k=0}^{\infty} \beta^{k+1} \mathbb{E}_t \left[\left(R_{t+k+1}^{wd} - \frac{1}{\beta} \right) P_{t+k} \right]$$

or

$$1 - \frac{\Phi_t}{P_t} = \frac{P_t^\star}{P_t} = 1 + \sum_{k=0}^{\infty} \beta^{k+1} \mathbb{E}_t \left[\left(R_{t+k+1}^{wd} - \frac{1}{\beta} \right) \frac{P_{t+k}}{P_t} \right]$$

C First-Order Approximation

We can obtain the result that the ratio $P_t^\star/P_t$, to a first-order approximation, depends only on expected returns in two ways. One is with a linear approximation and another is with a log-linear approximation. We present both here.

We repeat equation (4)

$$\frac{P_t^\star}{P_t} = 1 + \sum_{k=0}^{\infty} \beta^{k+1} \mathbb{E}_t \left[\left(R_{t+k+1}^{wd} - \frac{1}{\beta} \right) \frac{P_{t+k}}{P_t} \right]$$

Note that the sum on the right side of this equation can be written

$$\sum_{k=0}^{\infty} \beta^{k+1} \mathbb{E}_t \left[R_{t+k+1}^{wd} - \frac{1}{\beta} \right] \mathbb{E}_t \frac{P_{t+k}}{P_t} + \sum_{k=0}^{\infty} \beta^{k+1} \operatorname{Cov}_t \left(\left[R_{t+k+1}^{wd} - \frac{1}{\beta} \right], \frac{P_{t+k}}{P_t} \right)$$

We will assume that these covariance terms are constant over time

$$H = \sum_{k=0}^{\infty} \beta^{k+1} \operatorname{Cov}_t \left(\left[R^{wd}_{t+k+1} - \frac{1}{\beta} \right], \frac{P_{t+k}}{P_t} \right)$$

With this assumption we have that

$$\frac{P_t^\star}{P_t} = 1 + H + \sum_{k=0}^{\infty} \beta^{k+1} Z_t^{(k)} G_{P,t}^{(k)} \tag{18}$$

where

$$Z_t^{(k)} \equiv \mathbb{E}_t \left[R^{wd}_{t+k+1} - \frac{1}{\beta} \right]$$

and

$$G_{P,t}^{(k)} \equiv \mathbb{E}_t \frac{P_{t+k}}{P_t}$$

If we take a first-order Taylor approximation to equation (18) around the vectors $Z_t^{(k)} = 0$ (corresponding to $\mathbb{E}_t R^{wd}_{t+k+1} = 1/\beta$) and $G_{P,t}^{(k)} = \bar{G}_P^k$, we get

$$\frac{P_t^\star}{P_t} \approx 1 + H + \sum_{k=0}^{\infty} \beta^{k+1} \bar{G}_P^k Z_t^{(k)}$$

which we can write as

$$\frac{P_t^\star}{P_t} \approx 1 + H + \beta \sum_{k=0}^{\infty} \rho^k Z_t^{(k)} = 1 + H + \sum_{k=0}^{\infty} \rho^k \left[\beta \mathbb{E}_t R^{wd}_{t+k+1} - 1 \right] \tag{19}$$

where

$$\rho \equiv \beta \bar{G}_P$$

Note that equation (19) gives us the result that the dynamics of the ratio of fundamental price to observed price are determined, to first order, only by the deviations of expected returns from the benchmark value of $1/\beta$.

A natural choice for β would be the reciprocal of the unconditional expectation of returns with dividends $1/\mathbb{E}R^{wd}_{t+1}$. A natural choice for $\bar{G}_P$ would be the unconditional expectation of returns without dividends $\mathbb{E}R^{nd}_{t+1}$. Hence, a natural choice for ρ would be the ratio in population of expected returns without dividends $\mathbb{E}R^{nd}_{t+1}$ and with dividends $\mathbb{E}R^{wd}_{t+1}$. In a deterministic model, this would be the ratio

$$\rho = \frac{P_{t+1}}{P_{t+1} + D_{t+1}} = \frac{1}{1 + D/P}$$

Note that this choice of ρ would depend on the trading strategy considered as this central dividend-price ratio depends on the growth of $\{S_t\}$. In particular, we see in Figure 2 panels (c) and (d) that the sample mean cash flow to value ratio for an aggregate investor and a per share investor are considerably different. This is simply the flip side of the observation that the cumulative price appreciation shown in Figure 1a for these two trading strategies is considerably different. Hence, one might wish to choose a value of $\rho = 0.98$ for an aggregate investor and $\rho = 0.96$ for a per share investor. In this paper we choose a single value of $\rho = 0.96$ for all of our calculations simply to be consistent with this choice in the literature. We report the sensitivity of our results to this parameter in Appendix K.

We now move to a log-linear approximation. Specifically, to obtain equation (7), we make further assumptions about second moments. If returns are lognormal with constant conditional variances, then we have

$$\mathbb{E}_t R^{wd}_{t+k+1} = \exp\left(\mathbb{E}_t r^{wd}_{t+k+1} + J^{(k)}_t\right)$$

where $J^{(k)}_t$ is one half the conditional variance of log returns at horizon k given information at t. Further assume that $J^{(k)}$ is a constant equal to J. (This would be true if returns were independent over time. It need not be true if there is persistent predictability of returns). With these assumptions, we can write equation (19) as

$$\frac{P^\star_t}{P_t} \approx 1 + H + \sum_{k=0}^{\infty} \rho^k \beta Z^{(k)}_t = 1 + H + \sum_{k=0}^{\infty} \rho^k \left[\beta \exp(\bar{r} + J) \exp\left(\mathbb{E}_t r^{wd}_{t+k+1} - \bar{r}\right) - 1\right]$$

A linear approximation to this equation around $\mathbb{E}_t r^{wd}_{t+k+1} = \bar{r}$ gives

$$\log\left(\frac{P^\star_t}{P_t}\right) \approx \log\left(1 + H + \frac{\beta \exp(\bar{r} + J) - 1}{1 - \rho}\right) + \frac{1}{1 + H} \sum_{k=0}^{\infty} \rho^k \beta \exp(\bar{r} + J) \left(\mathbb{E}_t r^{wd}_{t+k+1} - \bar{r}\right)$$

We then obtain equation (7) under the assumptions that $H = 0$ and $\beta \exp(\bar{r} + J) = 1$.

We do not take a stand on the validity of the assumptions that the covariance terms captured by H and $J^{(k)}$ are in fact approximately constant. These assumptions that second moments are constant are standard when one takes a first-order approximation.

Our derivation of equation (7) as a first order approximation of $\log(P^\star_t/P_t)$ does not require that cash flows be positive. In the event that cash flows D_t are always positive, then one can derive equation (7) directly as in Nagel (2024) by taking the standard Campbell and Shiller (1988a) first order approximation to $\log(P_t/D_t)$ and subtracting off a first order approximation to $\log(P^\star_t/D_t)$ derived from equation (5). Our derivation shows that equation (7) extends to cases in which the first order approximation to $\log(P^\star_t/D_t)$ is either not defined because cash flows are negative or inaccurate because cash flows are too close to zero. See Appendix F.1.

D Index Construction

The documentation for the CRSP Index Methodology is at `https://www.crsp.org/indexes/crsp-market-indexes-methodology/`. Here we give a heuristic summary of Chapter 8 of this methodology manual describing index maintenance.

To construct these indices, one chooses a universe of firms to be included in the index in period t. We can index these firms by subscript i and this set of firms by Ω_t. At the start of the accounting period we have a number of shares outstanding $s_{i,t}$ for each firm $i \in \Omega_t$ and an initial price $p^o_{i,t}$. On a daily basis, a security's start-of-day price is defined as the previous day's end-of-day price adjusted for all splits and non-ordinary distributions, such as special dividends, returns of capital, and spin-offs.

Given the start of day price and the number of shares outstanding at the start of the day, one can construct value weights $w_{i,t} = p^o_{i,t} s_{i,t} / \sum_{j\in\Omega_t} p^o_{j,t} s_{j,t}$ and compute index returns without and with dividends for the period t from

$$R^{nd}_t = \sum_{i\in\Omega_t} \frac{p_{i,t}}{p^o_{i,t}} w_{i,t} \qquad \text{and} \qquad R^{wd}_t = \sum_{i\in\Omega_t} \frac{p_{i,t} + d_{i,t}}{p^o_{i,t}} w_{i,t}$$

where $p_{i,t}$ is the price of share i at the end of period t and $d_{i,t}$ are regular dividends paid by firm i during the period. The level of the index is referred to as the *price per share*, which we denote by $\tilde{P}_t$. The change in the level of the index from $t-1$ to t is given by $\tilde{P}_t = R^{nd}_t \tilde{P}_{t-1}$. The *dividends per share* paid on the index in period t, which we denote by $\tilde{D}_t$, is the solution to $R^{wd}_t = R^{nd}_t + \tilde{D}_t / \tilde{P}_{t-1}$.

The entity constructing the index then follows a preset methodology to adjust the start of period t prices $p^o_{i,t}$ relative to end of period $t-1$ prices $p_{i,t-1}$ as well as the number of shares outstanding for each firm in the index to adjust for corporate actions. These include special dividends, share repurchases, mergers, and the entry and exit of firms from the index. As discussed above, however, since the growth of the index from the end of $t-1$ to the end of t is determined by R^{nd}_t for the index, the level of price per share does not change between the end of period $t-1$ and the beginning of period t as a result of these corporate actions. However, the total market capitalization of the shares in the index does change from the end of period $t-1$ to the beginning of period t as a result of these corporate actions. It changes from $TMC_{t-1} = \sum_{j\in\Omega_{t-1}} p_{j,t-1} s_{j,t-1}$ to $TMC^o_t = \sum_{j\in\Omega_t} p^o_{j,t} s_{j,t}$. Total market capitalization at the end of period t is then given by $TMC_t = R^{nd}_t TMC^o_t$. Thus, the ratio of the index to total market capitalization of the shares in the index at the end of period $t-1$, which we denote by S_{t-1}, differs from this same ratio S_t at the end of period t because total market capitalization incorporates the effects of corporate actions while the index level does not.

E Data Sources and Construction

We take the data on nominal aggregate Personal Consumption Expenditure from NIPA Table 2.3.5. We download the Consumer Price Index for all Urban Consumers (not seasonally adjusted) from FRED at `https://fred.stlouisfed.org/series/CPIAUCNS`. We use the level of the index for December of each year.

We work with three sets of three time series from the CRSP annual data files for the Value-Weighted Index. These three series are total returns with dividends, $R^{wd}_{j,t}$, total returns without dividends, $R^{nd}_{j,t}$, and total market capitalization of the stocks in the index, $TMC_{j,t}$. We take these series for $j = 1$ corresponding to the CRSP file for the NYSE, $j = 2$ corresponding to the CRSP file for the NYSE and Amex, and $j = 3$ corresponding to the CRSP file for the NYSE, Amex, and NASDAQ.

We construct the series for R^{wd}_t, R^{nd}_t, and TMC_t that we work with as follows. We use the series for returns R_t^{wd} and R_t^{nd} from the NYSE file ($j = 1$) through 1962, from the NYSE and Amex file ($j = 2$) for 1963 through 1972, and from the NYSE, Amex, and NASDAQ file ($j = 3$) from 1973 through 2024. We obtained legacy data through 2024 prior to the CRSP 2025 data format revision. We update these data through 2025 using the new monthly data for that year. We compounded monthly returns with and without dividends from the NYSE, Amex, and NASDAQ file for 2025 to construct annual returns for 2025. We used the total market capitalization of this index for December 2025 as our end of 2025 measure of this market cap.

We construct our series for TMC_t as follows. We take the series $TMC = TMC_{1,t}$ from the NYSE file ($j = 1$) for years through 1962. We define $\kappa_1 = \frac{TMC_{1,1962}}{TMC_{2,1962}}$ corresponding to the ratio of the total market capitalization of the stocks in the NYSE file at the end of 1962 to the total market capitalization of the stocks in the NYSE and Amex file for that year. We then set our measure of total market capitalization for 1963 through 1972 equal to $TMC_t = \kappa_1 TMC_{2,t}$. This adjustment is analogous to an adjustment of an index divisor so that our measure of total market capitalization does not jump due to the addition of the stocks listed on Amex during 1962. Likewise, we define $\kappa_2 = \frac{TMC_{2,1972}}{TMC_{3,1972}}$ corresponding to the ratio of the total market capitalization of the stocks in the NYSE and Amex file at the end of 1972 to the total market capitalization of the stocks in the NYSE, Amex, and NASDAQ file for that year. We then define our measure of total market capitalization as $TMC_t = \kappa_1 \kappa_2 TMC_{3,t}$ for 1973 through 2025.

We construct the measure of price per share that we use by setting $\tilde{P}_{1929} = TMC_{1929}$ and then we use for $t > 1929$, $\tilde{P}_t = R_t^{nd} \tilde{P}_{t-1}$. We construct the corresponding ratio of dividends per share to price per share from

$$\frac{\tilde{D}_t}{\tilde{P}_t} = \frac{R_t^{wd} - R_t^{nd}}{R_t^{nd}}. \tag{20}$$

We then construct our index of dividends per share from the product of $\tilde{P}_t$ and this dividend-price ratio. Our measure of the variable S_t representing the trading strategy of an index investor is simply

$$S_t = \frac{\tilde{P}_t}{TMC_t},$$

which corresponds to the reciprocal of the index divisor.

We construct our measure of total cash flows to an aggregate investor from the solution for $\bar{D}_t$ in the equation

$$\frac{TMC_t + \bar{D}_t}{TMC_{t-1}} = R_t^{wd}. \tag{21}$$

F Illustrative Examples: Alternative Trading Strategies

This numerical example illustrates how two investors can earn identical one-period returns while recording different cash flows and cash-flow-to-price ratios.

Suppose there is a firm with N_t shares outstanding at the end of period t. Let $\widetilde{P}_t$ denote price per share, and let $\bar{P}_t = \widetilde{P}_t N_t$ denote the firm's end-of-period market capitalization.

Firm-Level Accounting. Suppose that at the end of period t,

$$\begin{aligned}\bar{P}_t &= 100,\\ N_t &= 100,\\ \widetilde{P}_t &= 1.\end{aligned}$$

Suppose that in period $t+1$ the firm's total free cash flow is

$$\bar{D}_{t+1} = 20.$$

Assume that 10 is paid as cash dividends and 10 is used to repurchase shares:

$$\begin{aligned}N_t \widetilde{D}_{t+1} &= 10,\\ (N_t - N_{t+1})\widetilde{P}_{t+1} &= 10,\end{aligned}$$

where $\widetilde{D}_{t+1}$ denotes dividends per share.

Suppose that end-of-period market capitalization is

$$\bar{P}_{t+1} = \widetilde{P}_{t+1} N_{t+1} = 130.$$

The market value of the N_t shares outstanding before the repurchase is therefore

$$\begin{aligned}N_t \widetilde{P}_{t+1} &= (N_t - N_{t+1})\widetilde{P}_{t+1} + \widetilde{P}_{t+1} N_{t+1}\\ &= 10 + 130 = 140.\end{aligned}$$

It follows that

$$\begin{aligned}\widetilde{P}_{t+1} &= \frac{140}{100} = 1.40,\\ N_{t+1} &= \frac{\bar{P}_{t+1}}{\widetilde{P}_{t+1}} = \frac{130}{1.40},\\ \widetilde{D}_{t+1} &= \frac{10}{N_t} = 0.10.\end{aligned}$$

Trading Strategies. Define S_t as the investor's share of market ownership at the end of period t.

- The aggregate investor holds the entire firm, so $S_t = 1$ at every date.
- The per-share investor holds one share, so

$$S_t = \frac{1}{N_t} = 0.01, \qquad S_{t+1} = \frac{1}{N_{t+1}} = \frac{1.40}{130}.$$

Investor Outcomes. **Aggregate investor:**

$$\begin{aligned}
\text{Price at } t : \bar{P}_t &= 100,\\
\text{Cash flow at } t+1 : \bar{D}_{t+1} &= 20,\\
\text{Price at } t+1 : \bar{P}_{t+1} &= 130,\\
\text{Return} : \frac{\bar{P}_{t+1} + \bar{D}_{t+1}}{\bar{P}_t} &= \frac{130+20}{100} = 1.5,\\
\text{Cash-flow yield} : \frac{\bar{D}_{t+1}}{\bar{P}_{t+1}} &= \frac{20}{130}.
\end{aligned}$$

Per-share investor:

$$\begin{aligned}
\text{Price at } t : \widetilde{P}_t &= 1,\\
\text{Cash flow at } t+1 : \widetilde{D}_{t+1} &= 0.10,\\
\text{Price at } t+1 : \widetilde{P}_{t+1} &= 1.40,\\
\text{Return} : \frac{\widetilde{P}_{t+1} + \widetilde{D}_{t+1}}{\widetilde{P}_t} &= \frac{1.40+0.10}{1} = 1.5,\\
\text{Cash-flow yield} : \frac{\widetilde{D}_{t+1}}{\widetilde{P}_{t+1}} &= \frac{0.10}{1.40}.
\end{aligned}$$

Key Implication. Both investors earn identical one-period returns, but their cash-flow-to-price ratios differ:

$$\frac{\bar{D}_{t+1}}{\bar{P}_{t+1}} \neq \frac{\widetilde{D}_{t+1}}{\widetilde{P}_{t+1}}.$$

This example illustrates that total returns are invariant to the trading strategy, while cash-flow-to-price ratios depend mechanically on how payouts are split between dividends and repurchases.

Note also that equations (9) and (10) are satisfied. In particular, at $t+1$

$$\begin{aligned}
\widetilde{D}_{t+1} &= S_t \bar{D}_{t+1} + (S_t - S_{t+1})\bar{P}_{t+1},\\
0.10 &= 0.01(20) + \left(0.01 - \frac{1.40}{130}\right) 130.
\end{aligned}$$

F.1 A firm that does not pay dividends

We now explore the applicability of Proposition 1 and the first-order approximation (7) for a firm that does not pay dividends and instead returns cash to investors by changing the number of shares outstanding.

In particular, consider a firm that generates free cash flow $\{\bar{D}_t\}$ with value $\{\bar{P}_t\}$ given by

$$\bar{P}_t = \sum_{k=1}^{\infty} \mathbb{E}_t M_{t,t+k} \bar{D}_{t+k},$$

where $M_{t,t+k}$ is the stochastic discount factor between t and $t+k$.

Assume that $\bar{D}_t/\bar{P}_t > -1$ for all t. That assumes that the firm does not have negative free cash flow in any period in excess of its subsequent value.

Now assume that this firm starts with $N_0 = 1$ shares outstanding and it follows a policy of share issuance $\{N_t\}$ for $t \geq 1$ aimed at allowing it to pay no dividends and instead return all cash to its equity holders through net share repurchases. Define $S_t = 1/N_t$. We can solve from this trading strategy recursively from equation (9) as follows. Given $N_{t-1} = 1/S_{t-1}$, solve for S_t from equation (9) to set $D_t(S) = 0$. This solution is given by

$$S_t = S_{t-1} \left(\frac{\bar{D}_t}{\bar{P}_t} + 1 \right).$$

Note that our assumption above ensures that $S_t > 0$. Shares outstanding at the end of period t are then $N_t = 1/S_t$.

It is clear that, by construction, the fundamental value of a share of this firm is given by $P_t^*(S) = 0$ regardless of rates of return and free cash flows $\bar{D}_t$. The price appreciation per share is given by

$$\frac{P_{t+1}(S)}{P_t(S)} = R_{t+1}^{nd} = R_{t+1}^{wd}.$$

Assume further that β is chosen small enough so that our limiting condition for Proposition 1 given by $\lim_{K\to\infty} \beta^K \mathbb{E}_t P_{t+K}(S) = 0$ holds. This assumption holds when one choose β small enough such that

$$\lim_{K\to\infty} \beta^K \mathbb{E}_t \prod_{k=1}^{K} R_{t+k}^{wd} = 0$$

One can then verify that equation (3) holds with $P_t^\star = 0$. To see this point, review the steps of the proof of Proposition 1 in Appendix B. Given that $R_t^{wd} = R_t^{nd}$ for all realizations, the term $D_t(S)/P_t(S) = 0$ for all realizations. The remaining steps of the proof then carry through.

This example shows that there are trading strategies for which the approximation (7) is not sufficiently accurate to be useful. In particular, we have that the approximation is not defined as

$$\log \left(\frac{P_t^\star(S)}{P_t} \right) = \log(0)$$

By the same logic, this approximation should break down even with positive dividends as long as the bulk of cash flows to investors are in the form of share repurchases rather than dividends.

F.2 Impact of Trading Strategies on Dividend Growth

The identity in equation (5) can be applied to both the per-share dividend-price ratio and the aggregate one. Substituting the aggregate version into equation (11) gives

$$\frac{D_t(S)}{P_t(S)} = \frac{\bar{P}_t^*}{\bar{P}_t}\left(\frac{\bar{D}_t}{\bar{P}_t^*} - \frac{s_t\left(1+\frac{\bar{D}_t}{\bar{P}_t}\right)}{\frac{\bar{P}_t^*}{\bar{P}_t}}\right). \tag{22}$$

But we have argued that, up to a first-order approximation, the ratio of fundamental to observed price is independent of the trading strategy, implying $\frac{P_t^*(S)}{P_t(S)} \approx \frac{\bar{P}_t^*}{\bar{P}_t}$. Applying this approximation, the term in parentheses on the right-hand side of equation (22) must be an alternative expression for the inverse of expected cash flow growth defined as in equation (6), i.e.,

$$\left(\sum_{k=1}^{\infty}\beta^k\mathbb{E}_t\frac{D_{t+k}(S)}{D_t(S)}\right)^{-1} \approx \left(\sum_{k=1}^{\infty}\beta^k\mathbb{E}_t\frac{\bar{D}_{t+k}}{\bar{D}_t}\right)^{-1} - \frac{s_t\left(1+\frac{\bar{D}_t}{\bar{P}_t}\right)}{\frac{\bar{P}_t^*}{\bar{P}_t}}. \tag{23}$$

Note that when $s_t > 0$ (the per-share investor is increasing her ownership share) expected growth in dividends per share will exceed expected growth in aggregate cash flow. That is the natural counterpart to the previous point that when $s_t > 0$ the per-share dividend-price ratio is smaller than the aggregate one.

G Aggregate and Per Share Data for the S&P 500

The issues with the impact of trading strategies on measures of cash flows and price discussed in Section 5 and Figure 2 are pervasive in the construction of value-weighted stock market indices. To illustrate this point, we show the per-share trading strategy and alternative measures of the cash flow to price ratio for an aggregate and per share investor constructed using data for the S&P 500 Index. The S&P 500 Index is a value-weighted index of approximately 500 stocks that was launched in its modern form in March of 1957.

Monthly data on measures of returns with dividends R_t^{wd}, returns without dividends R_t^{nd}, and total market capitalization of the stocks in the index TMC_t are available in the CRSP data on WRDS. We download monthly data on these three series starting in April of 1957 through the end of 2025. We construct a measure of the level of the S&P 500 Index by setting it equal to total market capitalization in April of 1957. We construct the level of the index (price per share) for subsequent months using the formula $\tilde{P}_t = R_t^{nd}\tilde{P}_{t-1}$. We then construct the implied monthly trading strategy $\{S_t\}$ using the formula $S_t = \tilde{P}_t/TMC_t$.

One can use these data to construct monthly data on price and cash flows using the formulas (20) and (21), but the resulting series for cash flows from these monthly data are extremely volatile for both a per-share and an aggregate investor, likely because of a large seasonal component in corporate payouts. (This is the case as well if we use monthly data for the CRSP Total Value Weighted Index). To deal with this issue, we construct measures of the ratio of cash flows over the course of 12 months to current value as follows.

We construct our measure of total returns over a 12 month period by compounding monthly returns with dividends R_t^{wd}. This gives us a series for total returns over 12 months (for an investor who reinvests dividends received during the year in the index itself) that we denote by $R_{t,12}^{wd}$. We assume that this measure of total returns applies to both an aggregate and a per-share investor. We construct our measure of returns without dividends for a per share investor by compounding monthly returns without dividends, which is equivalent to $R_{t,12}^{nd} = \tilde{P}_t/\tilde{P}_{t-12}$. We construct our measure of returns without dividends over 12 months for an aggregate investor as TMC_t/TMC_{t-12}. We then use these 12 month measures of returns with and without dividends to construct our measures of the ratios of cash flows to price using equation (20) for both a per-share and an aggregate investor.

We show results from this exercise in Figure G.1. In the left panel of this figure, we show monthly data on the log of the per share trading strategy $\log S_t$ for the S&P 500 Index constructed as above. We see in this panel that the index grew more slowly than the total market capitalization of the stocks in the index (so that $\log(S_t)$ is declining) until the early 2000s. After that, we see that $\log(S_t)$ slightly increases over time. This changing trend growth in $\log(S_t)$ is similar to that for the CRSP Total Value Index. In the time period prior to 2000, $\log(S_t)$ declines at about 2 percent per year. Later, it appears to be flat or slightly rising. (Note that the S&P 500 does not have IPOs, but it does have firms that enter the index and firms that leave the index. A decline in S_t results when the firms that enter the index are larger in terms of their market cap than the firms that are removed from the index).

In the middle panel of Figure G.1 we show our measures of 12-month cash to price for an aggregate and a per-share investor in the S&P 500. The cash flow to price ratio for an aggregate investor is in blue. The ratio for a per share investor is in red. We see in this figure that the ratio of cash flows to price for an aggregate investor is volatile but does not show a trend over time. In contrast, the ratio of cash flow to price for a per share investor drops by about 2 percentage points from the early decades of the sample to the last two decades of the sample – from around 4% to under 2%. Again, these differences in the low frequency movements in the ratio of cash flow to price for an aggregate and per share investor are simply due to the change in trend in the per share trading strategy shown in the left panel of this figure.

In the right panel of Figure G.1, we show the log of the ratio of our measure of 12-month cash flows relative to aggregate PCE. As we saw for the CRSP Index in Figure 2, the change in the trend of the trading strategy for the S&P 500 per-share investor shown in the left panel of Figure G.1 mechanically generates a change in trend growth of dividends relative to PCE. This observation is again consistent with the hypothesis that the step down in the ratio of dividends per share to price per share around 2000 shown as the red line in the center panel of Figure G.1 is an artifact of the changing trend in the trading strategy, resulting in a changing trend growth of dividends per share, and not a long-term change in expected returns net of growth.

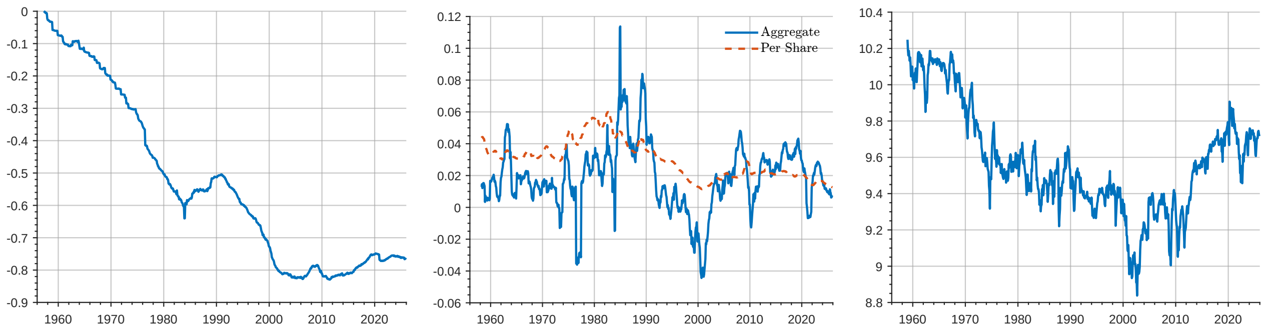

(a) $\log(S_t)$ for the S&P 500 Index.

(b) Ratios of 12-month cash flows to current price for an aggregate (blue) and per share (red) investor in the S&P 500.

(c) Log of 12-month cash flows relative to PCE expenditures for the S&P 500.

Figure G.1: Per Share Trading Strategy and Cash Flow to Value Ratios for the S&P 500.

H Alternative Trading Strategy: Have Expected Returns Declined?

We have shown that the persistent decline in the ratio of dividends per share to price per share in the 1990s coincides with a trend break in the rate at which the per-share investor's share of market ownership is diluted over time. Suppose this shift reflects an increase in the rate of share repurchases (a decline in the rate of equity issuance via existing firms or via IPOs would work similarly). Then one would expect a one-time decline in dividends per share, but faster growth in dividends per share moving forward. For example, if firms permanently raise the repurchase rate from zero to two percent of outstanding equity per year, price per share and dividends per share will grow two percent faster per year moving forward (see equations 8 and 9). Figure 2b illustrates precisely this sort of dynamic for dividends per share. We take that as a strong signal that one should interpret the downward shift observed in the series for dividends per share relative to price per share as reflecting a persistent change in the rate of net equity issuance and thus in the expected cash flow growth rate for a per share investor, rather than as a signal of lower expected returns.[15]

To further explore that hypothesis, we now consider yet another possible trading strategy. In particular, we construct time paths for D_t and P_t for a per-share investor who reduces the growth rate for S_t by a constant amount after 2000, so as to maintain the same downward trend in market ownership as that observed in Figure 2a for the pre-2000 period. This adjustment removes the post-2000 break in ownership growth that mechanically depresses the per-share dividend yield.

We construct our alternative series for trading strategy $\{S'_t\}$ for a per-share investor as

[15] Lettau and Van Nieuwerburgh (2008) have also recognized that return forecasting becomes difficult if there are very persistent changes at certain dates in the ratio of dividends per share to price per share. They propose a statistical procedure to identify those break dates and to adjust forecasting regressions accordingly. Our approach of comparing aggregate and per share ratios and plotting an explicit path for $\{S_t\}$ allows for a simple visual identification of a break date, and strongly suggests that this break reflects a secular change in corporate payout policy rather than a permanent change in expected returns.

follows. For dates $t = 1930$ through 1999 we set $S'_t = S_t$.

For dates $t = 2000$ through 2025, we set

$$S'_t = \exp\left(-(t - 1999) \times 0.0212\right) S_t.$$

The coefficient 0.0212 is the difference between the slopes of the two segments of the dotted red trend line for $\log(S_t)$ plotted in the left panel of Figure 2 $(0.0024 - (-0.0188) = 0.0212)$. With this construction, S'_t has the same downward trend in both sub-periods. We show this alternative trading strategy in Figure H.1a. This panel should be compared to Figure 2a.

To construct the corresponding alternative values of price per share $P_t(S')$, dividends per share $D_t(S')$, and the dividend-price ratio $D_t(S')/P_t(S')$, we use equations (8), (9), and (10) with our data on value and cash flow for an aggregate investor and our alternative series for the trading strategy $\{S'_t\}$.

Figure H.1b shows the resulting alternative dividend-price ratio. The dotted red lines correspond to the mean of this alternative dividend-price ratio over the periods 1929-1999 and 2000-2025. This panel should be compared to Figure 2d. We see from this comparison that the dividend-price ratio constructed from this alternative trading strategy does not show the stark change in mean shown by the dividend-price ratio for a per-share investor. As a result, Figure H.1d indicates that the model attributes almost all of the run-up in valuations in recent decades to higher expected future cash flows.

Table I.1 indicates that the dividend-price series associated with this strategy is a better predictor of one-period returns than the original per share series. But it is also less persistent than that series, and thus it implies a smaller variance for $\log(P^\star_t/P_t)$.

I Return Forecasts with Aggregate and Per Share Data

In this Appendix we report results from our return forecasting regressions. Note that, with an infinitely long sample of data one could evaluate equation (7) directly, by running direct return forecasting regressions at all horizons. In a finite sample, however, it is necessary to impose some structure on the path for expected returns. In what follows we evaluate our AR(1) assumption by comparing the predicted time paths for 10 year real returns implied by our AR(1) model against alternative forecasts from directly forecasting 10 year real returns using the current dividend to price ratio. We find that the two sets of forecasts track each other very closely.

We report results from OLS return forecasting regressions of the form

$$r^{wd}_{t+1} = \alpha_1 + \gamma_1 Z_t + error_{t+1}$$

and

$$\frac{1}{10}\sum_{k=1}^{10} r^{wd}_{t+k} = \alpha_{10} + \gamma_{10} Z_t + error_{t+10},$$

where Z_t are alternative return forecasting variables. The first regression estimates expected log real returns one year ahead. The second regression estimates annualized expected log real returns over the next ten years.

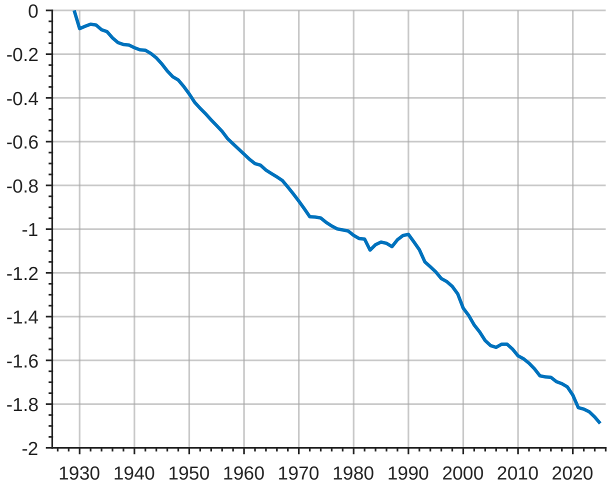

(a) Alternative trading strategy for a per-share investor, $\log(S')$.

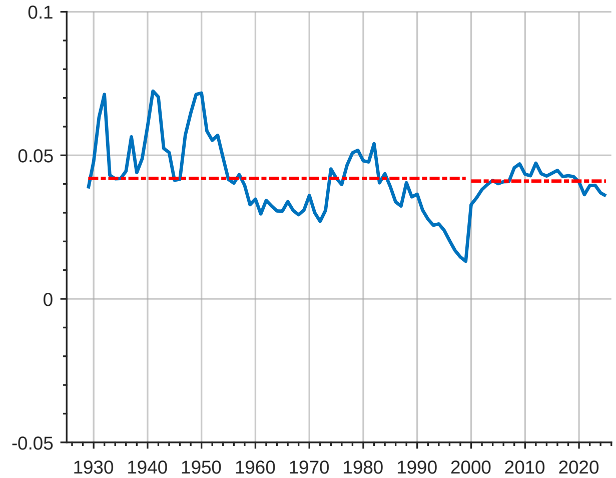

(b) Dividend–price ratio under the alternative trading strategy. Dotted red lines show means for 1929–1999 and 2000–2025.

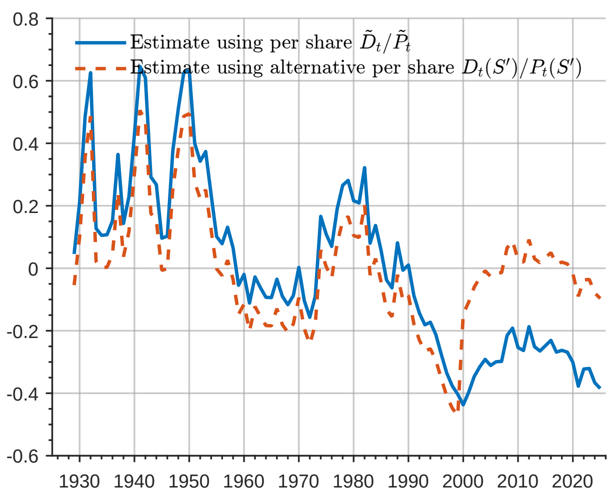

(c) Log ratio of fundamental to observed price estimated using per-share data and the alternative trading strategy.

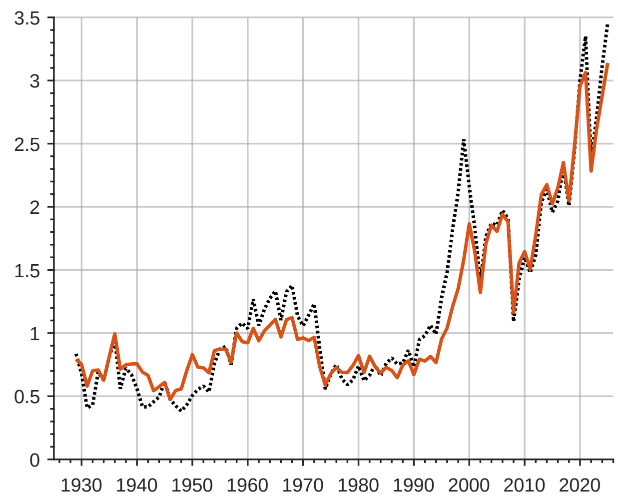

(d) Total market capitalization relative to PCE (dotted) and estimated fundamental price relative to PCE (solid) based on alternative strategy.

Figure H.1: Alternative per-share trading strategy and implied valuation decomposition.

We use three versions of Z_t to forecast log real returns one and ten years ahead. These are: (i) the ratio of cash flows to value for an aggregate investor $\frac{\bar{D}_t}{TMC_t}$, (ii) the ratio of dividends per share to price per share $\frac{\tilde{D}_t}{\tilde{P}_t}$, and (iii) our alternative measure of the dividend-price ratio for a per-share investor $D_t(S')/P_t(S')$ constructed as described in Appendix H.

Further results for these regressions are reported in Table I.1. We see in this table that the ratio of cash flows to value for an aggregate investor has strong predictive power for log real returns one year ahead. In contrast, both of our other measures of cash flow to value have low predictive power for log real returns over a one year horizon, but the R^2 of these

Table I.1: Return forecasting regressions using alternative valuation ratios

Forecasting Variable Z_t	Horizon	R^2	$\hat{\gamma}$	t-stat	ψ	$Var(\log(P_t^\star/P_t))$
Aggregate $\bar{D}_t/\bar{P}_t$	1 year	0.119	2.49	3.56	0.481	0.0151
	10 years	0.214	0.78			
Per-share $\tilde{D}_t/\tilde{P}_t$	1 year	0.026	2.01	1.59	0.923	0.0767
	10 years	0.222	1.44			
Alternative $D_t(S')/P_t(S')$	1 year	0.039	3.20	1.97	0.840	0.0387
	10 years	0.324	2.10			

The dependent variable is log CRSP real returns. The one-year horizon reports one-year-ahead returns. The ten-year horizon reports annualized log returns over the subsequent ten years. We do not report t-statistics for the ten year regressions as these regressions use overlapping data. See Boudoukh, Richardson, and Whitelaw (2008) for a discussion of the issues with such regressions.

regressions grows considerably at the ten year horizon.

We now use equation (13) to compute series for fundamental price relative to observed price implied by these regressions. To do so, we use estimates of ψ given by the autocorrelation of these three return forecasting variables. These give us an estimate of $\psi = 0.481$ when using the cash flow to value ratio for an aggregate investor, $\psi = 0.923$ when using the ratio for a per-share investor, and $\psi = 0.840$ when using the alternative ratio of cash flow to value for a per-share investor that we constructed in Appendix H.

We report the implied variance of $\log(P_t^\star/P_t)$ computed from these three estimates of expected returns one year ahead in Table I.1. We see that, despite the high R^2 of the one year return forecasting regression using the ratio of cash flows to value for an aggregate investor $\bar{D}_t/\bar{P}_t$, our estimate of the variance of the log ratio of fundamental price to observed price is the smallest of the three estimates. This result follows from the low persistence ψ of this valuation ratio when we apply equation (13). In contrast, the ratio of dividends per share to price per share for a per-share investor $\tilde{D}_t/\tilde{P}_t$ is a poor predictor of log real returns one year ahead, but, because it is very persistent, it results in the highest estimate of the variance of the log ratio of fundamental price to observed price. Note that this is true even though the alternative per share measure is a better predictor of 10 year returns.

Figure I.1 plots expected log real returns when these expected returns are based on regressions using the three different measures of cash flow to value described above. The top row of figures shows expected returns at a one year horizon, while the bottom row reports returns at a 10 year horizon. The forecasts based on per share measures indicate much more modest variability in one year expected returns, but much more volatility in expected returns over a 10 year horizon. Only the forecasts based on the original per share measure indicate a significant downward shift in long run expected returns in recent decades.

We can also examine the regression-implied time series estimates for the log ratio of fundamental price to observed price and for the level of $P_t^\star/PCE_t$. We have shown results using the first two measures of the ratio of cash flows to value in Figure 3. We show results using our alternative dividend-price ratio in Figure H.1.

In the left panel of this figure, we show the estimate of $\log(P_t^\star/P_t)$ derived from the dividend-price ratio of a per-share investor in blue and with our adjusted dividend-price

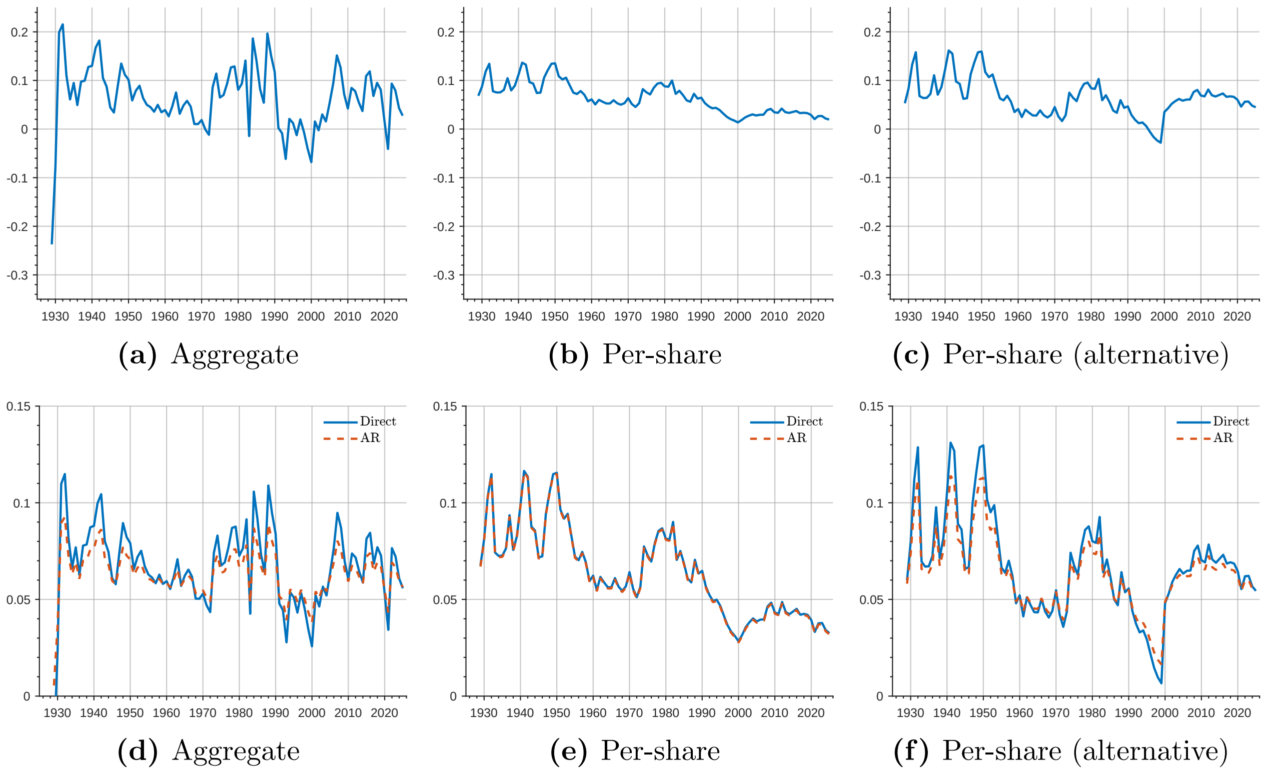

(a) Aggregate **(b)** Per-share **(c)** Per-share (alternative)

(d) Aggregate **(e)** Per-share **(f)** Per-share (alternative)

Figure I.1: Annualized expected log real returns net of PCE growth estimated using alternative cash-flow-to-value ratios. Top row: one-year-ahead forecasts from direct regressions. Bottom row: ten-year-ahead annualized forecasts from direct regressions (blue) and those implied by the one-year regression and estimated predictor persistence (red). Left panels: cash-flow-to-value for an aggregate investor. Middle panels: cash-flow-to-value for a per-share investor. Right panels: cash-flow-to-value for a per-share investor under the alternative trading strategy.

ratio for a per-share investor in red. We see in this figure that our adjustment of the trading strategy results in a very different picture of the relationship between fundamental price and observed price after 2000 — our estimate using the dividend-price ratio for a per-share investor shows fundamental price well below observed price (the log ratio is substantially negative) while this is not the case when we consider our adjusted dividend-price ratio (this log ratio is close to zero).

In Figure H.1d we show the ratio of total market capitalization to PCE in black and the ratio of fundamental price to PCE in red where we have estimated fundamental price using the dividend-price ratio for a per-share investor with our alternative trading strategy. We see here that the dynamics of fundamental price since 2000 appear to account for a very large portion of the dynamics of observed price. This is a direct consequence of our finding in Figure H.1c that the log ratio of fundamental price to observed price using this estimate is close to zero in recent decades. From equation (13), this finding is itself a direct result of the observation that our alternative dividend-price ratio is close to its historical mean during this time period and thus our forecast of expected returns is close to its historical mean over this time period.

From equation (16), one can decompose the variance of the observed aggregate price

Table I.2: Variance decomposition of log aggregate market capitalization (relative to PCE)

	$Var\,(\log P_t^\star)$	$Var\,(\log(P_t/P_t^\star))$	$Cov\,(\log P_t^\star, \log(P_t/P_t^\star))$
Aggregate	0.277	0.015	0.013
Per-share	0.107	0.077	0.067
Alternative	0.224	0.039	0.028

Decomposition is $Var(\log P_t) = Var(\log P_t^\star) + Var(\log(P_t/P_t^\star)) + 2Cov(\log P_t^\star, \log(P_t/P_t^\star))$. $Var(\log P_t) = 0.3178$.

relative to PCE as

$$\mathrm{Var}\left(\log \bar{P}_t\right) = \mathrm{Var}\left(\log \frac{\bar{P}_t}{\bar{P}_t^\star}\right) + \mathrm{Var}\left(\log \bar{P}_t^\star\right) + 2\,\mathrm{Cov}\left(\log \frac{\bar{P}_t}{\bar{P}_t^\star}, \log \bar{P}_t^\star\right).$$

Thus, observed price volatility can be attributed to time-varying expected returns, time variation in the fundamental valuation, and co-movement between those two sources of valuation movements. We report decompositions for our different models in Table I.2.

J Standard errors and confidence intervals

Table 1 reports, for each of the two return predictors, the one-year forecasting R^2, the persistence ψ, a bias-corrected estimate of ψ, and the implied standard deviation of $\log(P_t^\star/P_t)$, each with a 95% confidence interval. This appendix describes how those intervals are constructed.

Estimator

For predictor Z_t (the aggregate or per-share cash-flow-to-value ratio) we estimate two equations by OLS on the overlapping annual sample,

$$r_{t+1}^{wd} = a + \gamma\, Z_t + \varepsilon_{t+1}, \tag{24}$$

$$Z_{t+1} = c + \psi\, Z_t + u_{t+1}. \tag{25}$$

Equation 24 is the one-year return forecast. The fit R^2 is reported in column (1) of Table 1. Equation 25 is a first-order autoregression whose slope is the persistence ψ. We estimate ψ as the OLS slope of this regression (column (2) of Table 1), the convention in the return-predictability literature; this is the estimator for which the small-sample corrections discussed below are derived (see column (3) of Table 1). The implied volatility of the fundamental-to-observed price ratio follows from equation (13) as

$$\mathrm{Std}\Big(\log \tfrac{P_t^\star}{P_t}\Big) = \frac{\mathrm{Std}(\widehat{\mathbb{E}}_t[r_{t+1}^{wd}])}{1-\rho\psi}, \qquad \rho = 0.96,$$

which is column (4).

Bootstrap

Expected returns and the predictor are driven by shocks that are contemporaneously correlated: the innovation u_{t+1} that moves the valuation ratio at $t+1$ is the same price surprise that shows up in the realized return ε_{t+1}. To preserve that correlation we resample the two innovation series *as matched pairs*. Specifically, we form the fitted residual pairs $\{(\varepsilon_{t+1}, u_{t+1})\}$ from equation (24)–equation (25), draw from them with replacement, and use each draw to simulate a new predictor path from equation (25) and a matching return path from equation (24), of the same length as the data. On each simulated sample we re-estimate R^2, ψ, and $\mathrm{Std}(\log P^\star/P)$. The reported intervals are the 2.5th and 97.5th percentiles of the resulting bootstrap distributions, over 10,000 replications.[16]

Two points about the resulting intervals. First, the confidence interval reported for ψ is the percentile interval of the OLS estimates across bootstrap samples. Second, the interval for $\mathrm{Std}(\log P^\star/P)$ propagates uncertainty in both the forecasting slope γ and the persistence ψ, because each is re-estimated on every replication. This is what makes the ψ intervals for the two predictors non-overlapping while the $\mathrm{Std}(\log P^\star/P)$ intervals overlap substantially: the persistences are distinguishable, but the magnitude of the transitory component is not.

Bias correction

Least-squares estimates of an autoregressive slope are biased downward in finite samples, the more so the more persistent the series. We report a bootstrap bias-corrected estimate,

$$\psi_{\mathrm{bc}} = 2\hat{\psi} - \tfrac{1}{B}\textstyle\sum_b \hat{\psi}^*_b,$$

where $\hat{\psi}^*_b$ is the estimate on bootstrap sample b; this subtracts the bootstrap estimate of the bias from the point estimate. The correction is small for the aggregate predictor and larger for the more persistent per-share predictor, consistent with the direction of the finite-sample bias.

Other calculations not reported in the text

In the course of the analysis we computed several additional statistics and alternative corrections. They are consistent with what the paper reports and are omitted only to keep Table 1 compact; we summarize them here and they are reproducible from the code.

- *Alternative bias corrections for ψ.* Besides the bootstrap correction above, we computed the Kendall (1954) analytic approximation for an AR(1), $\psi_{\mathrm{bc}} \approx \hat{\psi} + (1+3\hat{\psi})/T$, and the Andrews (1993) median-unbiased estimator obtained by inverting the median of the OLS slope as a function of the true root. The three agree closely; the Andrews and Kendall values sit within the bootstrap confidence interval and tell the same story as the bootstrap correction, so we report only the latter. We also verified that return and

[16] The pairing is between ε_{t+1} and u_{t+1}, which share the same time stamp, rather than between ε_{t+1} and u_t. This preserves the contemporaneous return–predictor innovation correlation emphasized in the predictive-regression literature.

predictor innovations are substantially correlated; the matched-pair bootstrap preserves this feature.

- *Explosive draws.* Because the map $\psi \mapsto (1-\rho\psi)^{-1}$ diverges as $\rho\psi \to 1$, we track the fraction of bootstrap samples in the explosive region $\rho\hat{\psi}^* \geq 1$. Under the sample-autocorrelation estimator this region is unreachable, since an autocorrelation is bounded by one and $1 < 1/\rho$; under the OLS slope it is reachable in principle but essentially never occurs in our samples. We form the confidence interval for $\mathrm{Std}(\log P^\star/P)$ over the non-explosive draws; in practice this discards a negligible fraction and does not affect the reported endpoints.

- *Standard errors.* Alongside the percentile intervals we computed bootstrap standard errors (the standard deviation of each statistic across replications). We report intervals rather than standard errors in the table because the bootstrap distributions of ψ and of $\mathrm{Std}(\log P^\star/P)$ are noticeably asymmetric, so a symmetric standard-error band would misstate the uncertainty; the percentile interval is the more faithful summary.

Full details of every calculation in this appendix, including the alternative corrections we do not tabulate, are in the replication code.

K Sensitivity to the approximation constant ρ

Throughout the paper we set the Campbell–Shiller approximation constant to $\rho = 0.96$. This appendix reports how our estimates change when we instead set $\rho = 0.98$, a value that may be more appropriate for an aggregate investor with a lower dividend yield.

The constant ρ enters the estimates through a single channel: the factor $(1-\rho\psi)^{-1}$ in equation (13). It therefore leaves the forecasting fit R^2 and the persistence ψ—both the OLS estimate and its bias-corrected counterpart—completely unchanged; only the implied standard deviation of $\log(P_t^\star/P_t)$ responds. Table K.1 reports that statistic, and its confidence interval, at the two values of ρ; the other columns of Table 1 are identical to those reported there and are not repeated.

The effect is negligible for the aggregate predictor and sizeable for the per-share predictor: raising ρ from 0.96 to 0.98 moves the aggregate estimate from 0.123 to 0.125, but moves the per-share estimate from 0.277 to 0.331. This is the same convexity that makes ψ the decisive parameter in the paper. A change in ρ acts much like a change in ψ, and it matters where ψ is largest—for the highly persistent per-share predictor.

L A Vector-Autoregressive Specification

In the main text we forecast returns using one valuation ratio at a time. A natural question, and one raised in the literature since Campbell (1991), is whether a multivariate forecasting system—a vector autoregression in several predictors— would deliver a sharper answer. This appendix reports such a specification and explains why we do not think it changes the paper's conclusions. The short version is that the VAR does not resolve the identification problem at the center of the paper; it relocates it, from the persistence of a single predictor to the

Table K.1: Std($\log P_t^\star/P_t$) at $\rho = 0.96$ and $\rho = 0.98$

	Std$\left(\log \frac{P_t^\star}{P_t}\right)$	
	$\rho = 0.96$	$\rho = 0.98$
Aggregate $\bar{D}_t/\bar{P}_t$	0.123	0.125
95 percent c.i.	(0.037 , 0.173)	(0.038 , 0.177)
Per-share $\tilde{D}_t/\tilde{P}_t$	0.277	0.331
95 percent c.i.	(0.052 , 0.551)	(0.060 , 0.682)

Both columns use conventional real returns and the OLS estimate of ψ. The remaining statistics in Table 1 (R^2 and ψ) do not depend on ρ and are unchanged. Bootstrap 95% confidence intervals are reported below each estimate, computed as described in Appendix J.

dominant eigenvalue of the system, where it is estimated no more precisely and is harder to see.

Specification

Let $z_t = (\bar{D}_t/\bar{P}_t,\ \tilde{D}_t/\tilde{P}_t)'$ collect the aggregate and per-share valuation ratios. We estimate a first-order vector autoregression

$$z_{t+1} = c + \Psi\, z_t + u_{t+1}, \tag{26}$$

together with a one-year return forecast that uses both ratios,

$$r_{t+1}^{wd} = a + \gamma' z_t + \varepsilon_{t+1}. \tag{27}$$

The discount-rate object is then formed exactly as in the univariate case, but with the scalar autoregression replaced by the matrix Ψ. Iterating equation (26) forward and discounting,

$$\log \frac{P_t^\star}{P_t} \approx \gamma'(I - \rho\Psi)^{-1}(z_t - \bar{z}), \tag{28}$$

which is the multivariate counterpart of the univariate expression $\theta\,(Z_t - \bar{Z})$ with $\theta = \gamma/(1 - \rho\psi)$. In the univariate case the persistence ψ governs the amplification through $(1 - \rho\psi)^{-1}$; in the VAR the same role is played by the eigenvalues of Ψ through $(I - \rho\Psi)^{-1}$, and the amplification is dominated by the largest eigenvalue.

Results

Estimating equation (26)–equation (27) yields a dominant eigenvalue of Ψ of 0.94 and an implied standard deviation of $\log(P_t^\star/P_t)$ of 0.25. The implied standard deviation lies between the two univariate estimates, while the dominant eigenvalue is slightly above the per-share estimate.

Table L.1: Univariate specifications and the bivariate VAR

	Persistence (ψ or dominant eig.)	$\text{Std}\left(\log \frac{P_t^\star}{P_t}\right)$
Aggregate $\bar{D}_t/\bar{P}_t$ (univariate)	0.48	0.12
95 percent c.i.	(0.27 , 0.62)	(0.04 , 0.17)
Per-share $\tilde{D}_t/\tilde{P}_t$ (univariate)	0.92	0.28
95 percent c.i.	(0.75 , 0.96)	(0.05 , 0.54)
Bivariate VAR	0.94	0.25
95 percent c.i.	(0.77 , 0.97)	(0.03 , 0.27)

The persistence column reports the OLS AR(1) coefficient ψ for the univariate rows and the dominant eigenvalue of the companion matrix Ψ for the VAR row. Bootstrap 95% confidence intervals are in parentheses. $\rho = 0.96$ throughout. The dominant eigenvalue of the VAR is estimated no more precisely than the univariate per-share ψ.

How the persistent component enters the VAR

Why does the VAR's implied standard deviation lie in between the two univariate estimates? First, within the VAR the aggregate ratio does essentially all of the direct return forecasting: adding the per-share ratio to the one-year return regression leaves the R^2 unchanged (at 0.119) and enters with a coefficient near zero. Taken at face value, this might suggest the per-share ratio is irrelevant and the VAR should simply reproduce the aggregate univariate estimate of 0.12. Instead it produces a higher value of 0.25.

The reason is that the persistent per-share ratio enters expected returns indirectly through the cross-equation dynamics of the VAR. In the estimated system, the per-share ratio helps forecast the future aggregate ratio, and the aggregate ratio in turn forecasts returns. Thus, the persistent per-share ratio propagates into long-horizon expected returns along an indirect chain: per-share ratio today $\rightarrow$ aggregate ratio tomorrow $\rightarrow$ returns. Shutting down the cross-equation term through which the per-share ratio forecasts the aggregate ratio lowers the implied $\text{Std}(\log(P_t^\star/P_t))$ from 0.25 to 0.10. This indirect channel is therefore quantitatively important.

The amplification therefore turns on how persistent this propagating component is—which is exactly the dominant eigenvalue of Ψ. This is the multivariate form of the paper's central message: the variance of $\log(P_t^\star/P_t)$ is governed by the persistence of expected returns over long horizons, not by the fit of a short-horizon forecasting regression.

Why the VAR does not resolve the identification problem

The identification difficulty in the univariate analysis is that the magnitude of $\log(P_t^\star/P_t)$ depends on the persistence ψ through the convex factor $(1-\rho\psi)^{-1}$, and that ψ is not pinned down by a century of annual data: its 95% confidence interval for the per-share predictor runs from about 0.75 to 0.96, and the implied amplification therefore ranges from roughly 4 to 13.

The VAR does not escape this. In the multivariate system the amplification is governed by the dominant eigenvalue of Ψ through $(I - \rho\Psi)^{-1}$, and that eigenvalue is estimated with essentially the same imprecision as the univariate ψ. Its bootstrap 95% confidence interval runs from about 0.78 to 0.97— a width comparable to the univariate interval for ψ—and the implied amplification $(1 - \rho\lambda)^{-1}$ ranges from about 4 to 15. Substantial uncertainty about the magnitude of the valuation gap therefore remains in the VAR; it is expressed as uncertainty about an eigenvalue rather than about a scalar persistence coefficient.

M Implications of No Dividend Growth Forecastability

In this appendix, for illustrative purposes, we follow the argument in Cochrane (2008) that the lack of predictability of the growth in dividends per share gives us information on the dynamics of expected returns and hence of the log ratio of fundamental price to observed price. See Cochrane (2011), Campbell (2014), and Campbell (2018) for summaries of this argument. More recent research in this vein has found some evidence of predictability in the growth of dividends per share. See, for example, Van Binsbergen and Koijen (2010), Kelly and Pruitt (2013), and Jagannathan and Liu (2019).

Here, we demonstrate that, as under the regression-based approach to forecasting described above, the implications of this view for the question of what drives the stock market hinge entirely on the path for the dividend-price ratio, which in turn depends on the trading strategy. Specifically, using the original per share data, the Cochrane (2008) view suggests a rise in fundamental valuations of only around 50 percent between 1990 and 2025. In contrast, applying the Cochrane (2008) view to the alternative per share series attributes all of the surge in valuations over that period to the fundamental component (see Figure M.1). Again, the reason is that under the alternative per share trading strategy, the dividend-price ratio in the post 2000 period is close to its historical average, implying that the fundamental price must be close to the observed one, and that expected future returns must be close to average. Thus, even under the Cochrane view that cash flow growth is unpredictable, the alternative per-share series suggests that the large boom in observed valuations almost entirely reflects a rise in the expected present value of future cash flows.

To illustrate the method proposed by Cochrane (2008) consider again equations (5) and (6). If one takes as given that the growth of dividends per share in equation (6) cannot be forecast, then we should consider this term to be a constant over time. From taking logs of the terms in equation (5), we then conclude that the fluctuations in the log of the ratio of dividends per share to price per share on the left side of this equation correspond entirely to fluctuations in the log ratio of fundamental price to observed price. For illustrative purposes, we take the data on the deviations of the log ratio of dividends per share to price per share from its mean as an estimate of the log ratio of fundamental price to observed price, $\log(P_t^\star/P_t)$ (we impose that the mean of this series is zero). We refer to this estimate as our estimate obtained using the indirect method of Cochrane (2008).

Figure 2b shows a clear break in the trend growth of real dividends per share associated with the trend break in the per-share investor's trading strategy, $\log(S_t)$. This visible trend break makes the assumption that expected dividend growth is constant over the full sample unattractive. Hence, we will also construct an alternative indirect estimate of the log ratio of fundamental price to observed price using the dividend per share to price per share series that emerges from our alternative trading strategy described in Appendix H. We refer to this estimate as our estimate obtained using the indirect method of Cochrane (2008) with our alternative trading strategy.

We plot these series for $\log(P_t^\star/P_t)$ in the left panel of Figure M.1. Prior to 2000, both sets of estimates indicate large deviations between our estimates of fundamental price and observed price. After 2000, however, it appears that the gaps between observed price and fundamental price are highly sensitive to the data used in estimation. That is, the estimate of fundamental price in blue obtained using the data on ratio of dividends per share to price per share and this indirect method of Cochrane (2008) suggests that the fundamental price has been well below the observed price for the past 30 years. In contrast, if we use the ratio of dividends per share to price per share from our adjusted trading strategy to estimate fundamental price, as in the red line, we see that fundamental price is quite close to observed price over the past twenty years or more.

In the right panel of Figure M.1 we show the implications of these estimates for the predicted level of the fundamental component of total market capitalization, relative to aggregate PCE. Using the original per share data, this suggests a rise in fundamental valuations of only around 50 percent between 1990 and 2025. The ratio of dividends per share to price per share falls well below its historical mean in the post 2000 period (Figure 2d), which the Cochrane approach interprets as a large decline in the ratio $P_t^\star/P_t$, and (equivalently) as a large decline in future expected returns (see equation 7). Thus, the interpretation is that valuations have surged far above the present value of expected cash flows because expected returns have declined.

In contrast, applying the Cochrane view to the alternative per share series attributes all of the surge in valuations over that period to the fundamental component.

Following Campbell and Shiller (1988a), many researchers have focused on decomposing fluctuations in cash flow to value ratios into a component reflecting changes in expected returns versus a component reflecting changes in expected cash flow growth. But the discussion above highlights that this literature does not sufficiently address a more basic issue, which is that the *level* of the cash flow to value ratio depends on the trading strategy. The original dividends per share to price per share series has declined substantially, so if one focuses on that series the question of whether the decline is informative about changes in future returns or changes in future cash flow growth is live and important. But our alternative per share series has barely moved, so if one focuses on that series, the answer to the Campbell-Shiller question is inconsequential: even under the Cochrane view that cash flow growth is unpredictable, one is forced to conclude that the large boom in observed valuations almost entirely reflects a rise in the expected present value of future cash flows.

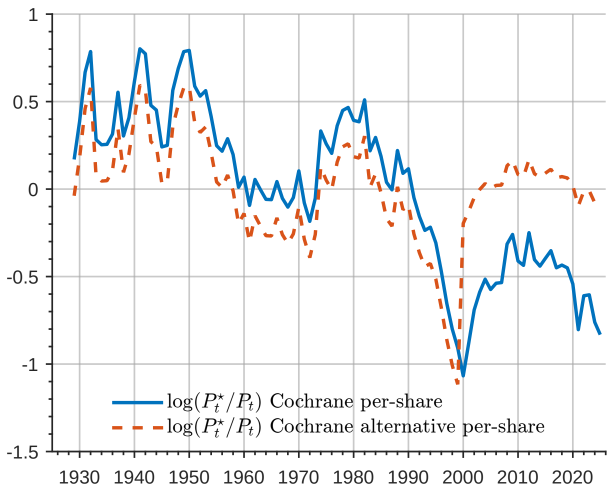

(a) Estimates of $\log(P_t^\star/P_t)$ using the Cochrane (2008) approach: $\log(P_t^\star/P_t) = \log(D_t/P_t) - \text{mean}(\log(D_t/P_t))$. Blue: original dividends per share relative to price per share. Red: alternative per-share series described in Appendix H.

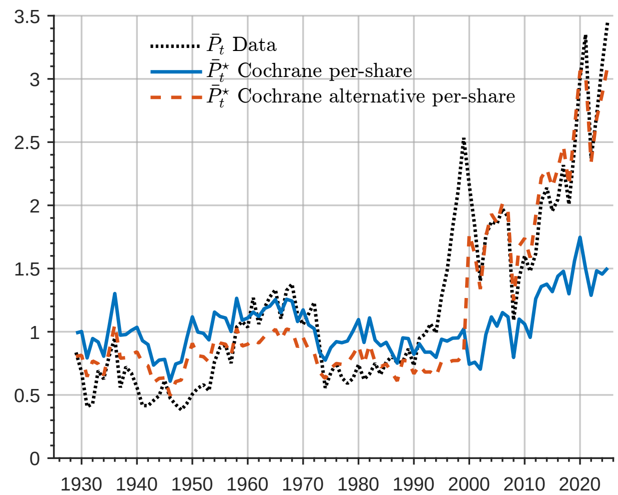

(b) Total market capitalization (black) and the implied fundamental component $\bar{P}_t^\star$ constructed from the two estimates for $\log(P_t^\star/P_t)$ shown in Panel (a). All these series are relative to aggregate PCE.

Figure M.1: Estimates of the fundamental price component using the Cochrane (2008) approach.