Dogs and Cats Living Together: A Defense of Cash-Flow Predictability*

Seth Pruitt

Arizona State University

July 21, 2023

Abstract

Present-value logic says that aggregate stock prices are driven by discount-rate and cash-flow expectations. Dividends and net repurchases are both cash flows between the firm and household sectors. Aggregate dividend-price ratios do not forecast dividend growth, but do robustly forecast future buybacks and issuance. Long-run variance decompositions say that discount-rate and cash-flow expectations contribute equally to aggregate dividend-price-ratio variation.

Keywords: stocks, dividend-price ratio, present value, predictability, variance decomposition

*seth.pruitt@asu.edu, Department of Finance, W.P. Carey School of Business, P.O. Box 873906, Tempe, AZ 85287-3906. Thanks to George Aragon, Michael Barnett, Hank Bessembinder, Oliver Boguth, John Campbell, Charlie Clarke, Maziar Kazemi, Sydney Ludvigson, William Murray, Jeffrey Pontiff, David Schreindorfer, Sunil Wahal, Jim Weston, Moto Yogo, and seminar participants at UMass (Isenberg), Rice (Jones), and ASU. All errors are mine alone.

1 Introduction

To analyze what drives aggregate stock prices, Campbell and Shiller (1988) developed a seminal decomposition with the following present-value intuition: the dividend-price ratio fluctuates due to varying expectations of future discount rates or future cash flows.[1] Thereafter, well-known regression results showed the dividend-price ratio significantly forecasts returns but not dividend growth. Hence, Cochrane (2005, 2008, 2011) and others have argued that aggregate dividend-price-ratio variation is almost entirely driven by discount-rate variation. This stylized fact has important economic consequences, for instance emphasizing discount-rate expectations over cash-flow expectations as the primary force driving asset prices.[2]

I argue the stylized fact is different: discount-rate and cash-flow expectations contribute about equally. First I show that the present-value identity of the well-known aggregate dividend-price ratio includes future stock buybacks and issuance, in addition to dividends, as cash flows between the household and firm sectors. Then I find these cash flows are robustly forecasted in the data. Those new facts change the variance decomposition across a variety of specifications, and the broad conclusion is about half of aggregate dividend-price fluctuations come from cash-flow expectations. Aggregate cash-flow expectations are alive and well and driving stock prices.[3]

To emphasize, the key point is that the same dividend-price ratio used by many preceding papers is related to future issuance and buybacks. This follows from the very definition of a return.[4] If we are investigating variation in the dividend-price ratio but ignoring buybacks and issuance, we have essentially imposed a constraint that those cash-flow expectations don't vary—we will see that the data reject that constraint.[5] Put another way, my estimates of return and dividend-growth forecast equations are essentially identical to Cochrane (2008, 2011) and Koijen and Van Nieuwerburgh (2011) because the predictor is the same. This

[1]In the absence of bubbles, which is assumed throughout the paper.

[2]For instance, Campbell and Cochrane (1999) emphasizes this sole channel, while Bansal and Yaron (2004) models this effect as stemming from cash-flow dynamics. As further examples of the impact of this stylized fact, see discussions in Koudijs and Voth (2016), Caballero and Simsek (2020), De La O and Myers (2021), and Dou et al. (2021), among others.

[3]Cochrane (2008) is titled "The Dog that Didn't Bark", dividends being the dog. Buybacks and issuance are cats. All live with the household.

[4]For instance, as given in Larrain and Yogo (2008) Section 5.1.

[5]My general point is reminiscent of Boudoukh et al. (2007)'s point that "all cash flow distributions to shareholders may have fundamental information about asset pricing", thus warning us to "be careful in using dividend yields alone." My twist on the idea is to use the dividend-price ratio alone, but show that "all cash flow distributions" matter to it.

paper's contribution is showing that buyback and issuance equations *also belong*, and finding robust empirical evidence of their predictability.[6] Because of this, the contribution of cash-flow news to dividend-price-ratio variation is larger than typically estimated. I point out that this observation implies dividend-price and excess-return variance decompositions more closely agree on the contribution of cash-flow expectations—which connects to the work of Campbell (1991), Campbell and Vuolteenaho (2004), and Campbell et al. (2018), amongst others.

I focus on the dividend-price ratio for at least three reasons. First, I aim to persuausively argue that stock prices respond to expected buybacks and issuance—showing they appear in the aggregate dividend-price ratio's present-value identity is a well-known way to do so. Second and relatedly, it is useful to directly connect to precursors like Campbell and Shiller (1988) and Koijen and Van Nieuwerburgh (2011)—this way, I am able show that the same variable delivers both familiar and new results.[7] Three, yearly total equity payout (i.e. dividends plus buybacks minus issuance) is negative many times in the data, even recently—therefore a log-linear payout-price-ratio decomposition is not always real-valued.

This paper connects most directly to Cochrane (2005, 2008, 2011) that use only CRSP data. As Campbell and Shiller (1988) note, "[t]he CRSP data incorporate careful corrections for stock splits, noncash distributions, mergers, delisting, and other potential problems". This means that one can measure a net repurchase (buyback minus issuance) between the firm and household sector whenever the cumulative-factor-adjusted number of outstanding shares changes (see Stephens and Weisbach, 1998): these are the non-stock distributions exactly implied by prices, returns, and shares outstanding.[8] Like Cochrane (2005, 2008, 2011), my main analysis requires only data from the Center for Research in Security Prices (CRSP).

[6]A caveat: the aggregate dividend-price ratio in Cochrane (2008, 2011) is not precisely the type constructed in Campbell and Shiller (1988), Welch and Goyal (2007), and Koijen and Van Nieuwerburgh (2011)—see Section 2. But the empirical difference is small, so I abstract from this in the main text. I am never arguing that the aggregate dividend-price ratio is mismeasured by anyone. I am arguing that cash flows beyond ordinary dividends should also be considered, in a point that is complementary to what Sabbatucci (2022) points out (as I discuss below).

[7]Another reason could also be related: the same variation that is part of its predictability also suggests net repurchases (i.e. buybacks minus issuance) could add noise to a scaled-stock-price predictor. The forward-looking part of any yield is the price, and smoothly-moving dividends impart stationarity with minimal backward-looking noise—which might be a statistical reason for sticking with the dividends as the cash scaling prices. Among others, see Lamont (1998), Campbell and Thompson (2008), and Kelly and Pruitt (2013) for discussions on the choice of price scalar.

[8]Intuitively, the implied issuance (if the cash flow is negative) or buyback (if the cash flow is positive) is measured by the yield that makes sense of the firm equity's starting market value, ending market value, and ex-dividend return (see Stephens and Weisbach, 1998; Bansal et al., 2005; Dichev, 2007; Welch and Goyal, 2007; Boudoukh et al., 2007; Bessembinder, 2018, amongst others for support of this view)—I detail this further in Sections 2 and 3.

It turns out that I find puzzling results regarding issuance in the CRSP data. The dividend-price ratio negatively predicts issuance, which is the opposite sign of what one would expect from the present-value relationship developed here.[9] To investigate further, I project the CRSP series on the aggregate amount of stock sales reported in Compustat—the predictive coefficient switches to be positive and significant. While my ultimate variance-decomposition results are robust to using either issuance measure, and I can remain agnostic about its sign and still accomplish my goals, future research is warranted to better understand what drives CRSP issuance.

Given the large related literature, further context for this paper's contribution is warranted. Koijen and Van Nieuwerburgh (2011) survey the research on return and cash-flow predictability, raising the issues of cash-flow reinvestment discussed by Binsbergen and Koijen (2010): so I use a zero-rate reinvestment strategy that Campbell and Shiller (1988) used and as advocated by Koijen and Van Nieuwerburgh (2011).[10] Boudoukh et al. (2007) construct on equity-payout yields that combine CRSP dividend data with net repurchase data from Computstat (starting in 1971), with Eaton and Paye (2017) following suit, and focus on the return predictability these ratios provide. Notably, Eaton and Paye (2017) find that their total-payout yields are positive after the Great Depression, which suggests that CRSP and Compustat may differ in what they measure as aggregate cash flows between the firm and household sectors; I leave exploration of this for future research, and use only CRSP data in my main empirical analysis. Larrain and Yogo (2008) focus on cash flows from both debt and equity, using both Compustat or Flow of Funds data, and therefore focuses on total firm value instead of only stock prices; also, for the purpose of investigating the total-payout yield, they present the return identity I too use. The contribution of this paper relative to Boudoukh et al. (2007), Larrain and Yogo (2008), and Eaton and Paye (2017) rests on choices, some already mentioned, of which I highlight two. First, I use the aggregate dividend-price ratio and show it has a present-value relationship with future buybacks and issuance. Second, I focus on results using CRSP data alone, and those data say that total-equity-payout yields cannot have a Campbell and Shiller (1988)-like log-linear decomposition because yearly equity payout is negative even in recent years (in fact, for almost all of 2021).

My conclusions connect to a number of other studies considering value-ratio predictability or arguing the importance of cash-flow news. Welch and Goyal (2007) find the aggregate dividend-price ratio does not forecast returns out-of-sample, which prompted Kelly and

[9]An observation, I should note, that appears due to my focus on the dividend-price ratio.

[10]My results are robust to instead using risk-free rate reinvestment. I also show my main conclusions are robust to structural breaks discussed in Koijen and Van Nieuwerburgh (2011) and analyzed fully in Lettau and Van Nieuwerburgh (2007).

Pruitt (2013) among others to use more sophisticated econometric methods that find greater return *and* dividend-growth predictability in book-to-market and dividend-price ratios; I find strong predictions of buyback and issuance cash flows using just the value-weighted dividend-price ratio and simple regressions, even out of sample. Chen and Zhao (2009) argue that decompositions are sensitive to the choice of target and predictors and find dividend-growth news is more important once this sensitivity is systematically addressed, whereas I use one predictor and show it robustly forecasts buyback and issuance cash flows. Chen et al. (2013) use analyst-forecast data, and Golez (2014) extracts dividend-growth expectations from the S&P500 using options prices,[11] and both find dividend-growth news an important driver of prices; I use only realized CRSP data, which starts in 1926, and predict buybacks and issuance. Pettenuzzo et al. (2020) put daily CRSP data into a Bayesian persistent-temporary-jump component model for dividend growth, and find the persistent component forecasts future dividend growth; I find that buybacks and issuance are even more strongly forecastable than is dividend growth, using only simple forecast equations and monthly data.[12] De La O and Myers (2021) use analyst expectations to argue that short-run dividend-growth expectations are the most important driver of the price-dividend ratio; I use only realized prices and cash flows and reach a similar conclusion by look at cash flows other than dividends, which could raise the question if analysts' expected dividend-growth is *only* showing up in future dividends or also shows up in future buybacks. Sabbatucci (2022) argues that M&A cash dividends are excluded from the standard measure of ordinary dividends, and once those are added back in then dividends are significantly predictable; I measure dividend-growth in the standard way (meaning my dividend-growth results are subject to Sabbatucci (2022)'s critique as well), but like him emphasize that non-ordinary-dividend cash flows (I focus on net repurchases) are important drivers of aggregate stock prices.[13]

The paper proceeds as follows. Section 2 presents a decomposition of the aggregate dividend-price ratio acknowledging that the number of shares outstanding can vary. Section 3 details the data and present summary statistics. Section 4 presents the benchmark results coming from simple forecasting equations estimated by OLS or GMM, and then presents a bat-

[11] His point being, take away the option-implied dividend-news and the adjusted dividend-price ratio better predicts returns.

[12] In fact, given that repurchase plans are often announced ahead of time (the precise timing of dividend announcements and realizations is something Pettenuzzo et al. (2020) take seriously), there is scope for future research to investigate whether net repurchase news enters into those daily stock price movements too. Pettenuzzo et al. (2022) may be moving in that direction.

[13] Somewhat related, Brogaard et al. (2022) use daily data and high-frequency TAQ data since 1990 to decompose news into firm-specific and market-wide components and find dividend-growth news is important to firms but idiosyncratic in nature, whereas I look only at aggregate data and find the market-wide buybacks and issuance drive aggregate stock prices.

tery of robustness checks considering different sample periods, non-overlapping observations, farther-ahead forecasts, different firm samples, and out-of-sample results. Section 5 uses the previous section's restricted-VAR results alongside unrestricted-VAR results, and provides dividend-price-ratio variance decompositions answering the paper's main question. I then conclude.

2 A dividend-price ratio decomposition

This section argues that buybacks and issuance appear in the identity of the aggregate dividend-price ratio, thus providing a motivation to empirically investigate their predictability. I start by considering a portfolio of all of one firm's stock because the algebra is a bit simpler and the same points appear there as in aggregate. We will obtain a decomposition by following Cochrane (2005)'s method, so start with the identity of a return.

2.1 A firm portfolio

For a stock n at the end of month t, let $P_{n,t}$ be the price per share, $D_{n,t}$ the dividend per share, and $S_{n,t}$ the number of shares outstanding. Then the definition of a firm's gross return is

$$
\begin{aligned}
R_{n,t+1} &= \left(\frac{P_{n,t+1}+D_{n,t+1}}{P_{n,t}}\right) \\
&= \left(\frac{S_{n,t+1}P_{n,t+1}+S_{n,t}D_{n,t+1}+(S_{n,t}-S_{n,t+1})P_{n,t+1}}{S_{n,t}P_{n,t}}\right) \qquad (1) \\
&= \frac{S_{n,t-1}D_{n,t}}{S_{n,t}P_{n,t}}\left(\frac{S_{n,t+1}P_{n,t+1}+S_{n,t}D_{n,t+1}+(S_{n,t}-S_{n,t+1})P_{n,t+1}}{S_{n,t-1}D_{n,t}}\right) \\
&= \frac{S_{n,t-1}D_{n,t}}{S_{n,t}P_{n,t}}\frac{S_{n,t}D_{n,t+1}}{S_{n,t-1}D_{n,t}}\left(1+\frac{S_{n,t+1}P_{n,t+1}}{S_{n,t}D_{n,t+1}}+\frac{(S_{n,t}-S_{n,t+1})P_{n,t+1}}{S_{n,t}D_{n,t+1}}\right). \qquad (2)
\end{aligned}
$$

Note that equation 1 is essentially what Larrain and Yogo (2008) use when discussing the equity payout yield (when I discuss the aggregate return below I even more closely relate to what they write).

At least three points are worth making. First, the fact that dividends-per-share, say $D_{n,t+1}$, is multiplied by the number of shares, $S_{n,t}$, of the previous period is consistent with how one calculates dividends in the data, as I'll discuss further below. Second, $S_{n,t}D_{n,t+1}$ is the total amount of dividends paid by the firm to the household (ultimately), and the firm's market capitalization $S_{n,t+1}P_{n,t+1}$ is the price of all the firm's equity—therefore $\frac{S_{n,t}D_{n,t+1}}{S_{n,t+1}P_{n,t+1}}$ and $\frac{S_{n,t-1}D_{n,t}}{S_{n,t}P_{n,t}}$ are dividend-price ratios. In fact, as the ratio of the price of a stock portfolio and the dividends paid on that portfolio, it is the dividend-price ratio described in Campbell and Shiller (1988). Third, such a dividend-price ratio leads to *two* cash-flow terms in the parenthesis: the gross growth rate of all dividend payments $S_{n,t}D_{n,t+1}/S_{n,t-1}D_{n,t}$, and net repurchases divided by all dividend payments $(S_{n,t} - S_{n,t+1})P_{n,t+1}/S_{n,t}D_{n,t+1}$. For this definition of net repurchases, see Stephens and Weisbach (1998), Bansal et al. (2005), Dichev (2007), Welch and Goyal (2007), Boudoukh et al. (2007), Larrain and Yogo (2008), and Bessembinder (2018), amongst others.

Obviously, when S_n is constant over time then the net-repurchase term in (2) is identically zero and S_n cancels out from the first three fractions. But when $S_{n,t}$ varies over time (in general it does), there are future net-repurchase cash flows related to the dividend-price ratio *by identity*. Nothing has been changed about the dividend-price ratio being considered—we are recognizing that it is driven by both ordinary dividends *and* net repurchases.

Crucially, we care about time-series variation in $S_{n,t}$ *that is not* a stock distribution like a split or stock-dividend. Those events change the number-of-shares-outstanding variable `SHROUT`, the price-per-share variable `PRC`, and the dividend-per-share variable in offsetting ways that do not matter to the economic question at hand.[14] For example: if the stock splits 2-for-1, then the number of shares doubles, while the price-per-share and dividend-per-share halve, and so this corporate event does not involve a cash flow between firm and household. Hence, it is helpful to describe CRSP variables `PRC` and `SHROUT` as directly pertaining to *exchange-traded* shares that are different than the shares S referred to by (1)—the variable S defines a single share's ownership stake $(\frac{1}{S})$ in the firm. Campbell and Shiller (1988) noted that "[t]he CRSP data incorporate careful corrections for stock splits, noncash distributions, mergers" and these corrections come via CRSP's cumulative adjustment factors. So view S, D, and P as referring to adjusted shares whose number *is not* altered by noncash distributions.

[14] I use `teletype font` to denote variable names in the CRSP database accessed via Wharton Research Data Services which, as I note below, I access via Wharton Research Data Services. Note there is no single variable in CRSP that delivers ordinary dividends per share. `DIVAMT` includes dividends both ordinary and not, and one needs to parse `DISTCD` to know which is which. The standard method of calculating ordinary dividends-per-share is (`RET`−`RETX`) times the previous period's price `PRC` (which agrees with `DIVAMT` whenever `DISTCD` says the dividend is ordinary). This standard method therefore agrees with the timing conventions shown in (2).

2.2 An aggregate portfolio

Now consider the value-weighted gross return

$$
\begin{aligned}
R_{t+1} &= \frac{\sum_n S_{n,t} P_{n,t} \left(\frac{P_{n,t+1} + D_{n,t+1}}{P_{n,t}} \right)}{\sum_n S_{n,t} P_{n,t}} \\
&= \frac{\sum_n S_{n,t} P_{n,t} \left(\frac{S_{n,t+1} P_{n,t+1} + S_{n,t} D_{n,t+1} + (S_{n,t} - S_{n,t+1}) P_{n,t+1}}{S_{n,t} P_{n,t}} \right)}{\sum_n S_{n,t} P_{n,t}} \qquad (3) \\
&= \frac{\sum_n S_{n,t-1} D_{n,t}}{\sum_n S_{n,t} P_{n,t}} \frac{\sum_n S_{n,t} D_{n,t+1}}{\sum_n S_{n,t-1} D_{n,t}} \left(1 + \frac{\sum_n S_{n,t+1} P_{n,t+1}}{\sum_n S_{n,t} D_{n,t+1}} + \frac{\sum_n (S_{n,t} - S_{n,t+1}) P_{n,t+1}}{\sum_n S_{n,t} D_{n,t+1}} \right) \qquad (4)
\end{aligned}
$$

Analogous to what (2) showed for a single firm, we are defining the aggregate dividend-price ratio as the total amount of paid dividends, divided by the total portfolio price (i.e. aggregate market capitalization). This is the aggregate dividend-price ratio constructed in Campbell and Shiller (1988), Welch and Goyal (2007), Koijen and Van Nieuwerburgh (2011), amongst others.[15] Once again obviously, if S_n is constant over time and firm then the fourth fraction in (4) is identically zero.

This aggregate dividend-price ratio is related to aggregate net repurchases *by identity* whenever a $S_{n,t}$ varies over time.[16] There is more than one cash-flow term linking the value-weighted gross return to the aggregate dividend-price ratio: the gross growth rate of dividends $\sum_n S_{n,t} D_{n,t+1} / \sum_n S_{n,t-1} D_{n,t}$, *and* net repurchases scaled by dividends $\sum_n (S_{n,t} - S_{n,t+1}) P_{n,t+1} / \sum_n S_{n,t} D_{n,t+1}$. This fact provides the simple economic motivation to empirically investigate whether or not the aggregate dividend-price ratio forecasts future buybacks and issuance, in addition to returns and dividend growth.

We need a variable to be positive for it to have a real-valued logarithm. With that in mind,

[15] In the case of Campbell and Shiller (1988), for the one calculated from CRSP data, not necessarily the one calculated from Cowles/S&P data which is described as using per-share data. My point here is not to argue that the two ratios are very different—it is to precisely state that it is for the former that we see net repurchases in the return identity.

Also, such yields typically sum dividends over a period of time, usually twelve months, to deal with seasonality—this practice is secondary to the point I am making here, and I will discuss yearly summing further below.

[16] I have seen two papers that come closest to what I'm pointing out here, but in both cases they are talking about the equity-payout ratio. Larrain and Yogo (2008) derive a log equity payout yield decomposition in their appendix, and note that outflow and inflow must be treated separately as I'm about to do in equation 5 below. Eaton and Paye (2017) also consider a log equity payout yield decomposition that is real-valued only when equity payout is positive.

rewrite (4) using $D_t \equiv \sum_n S_{n,t-1} D_{n,t}$ and $P_t \equiv \sum_n S_{n,t} P_{n,t}$:

$$
\begin{aligned}
R_{t+1} &= \frac{D_t}{P_t}\frac{D_{t+1}}{D_t}\left(1 + \frac{P_{t+1}}{D_{t+1}} + \frac{\sum_n (S_{n,t} - S_{n,t+1}) P_{n,t+1}}{D_{t+1}}\right) \\
&= \frac{D_t}{P_t}\frac{D_{t+1}}{D_t}\left(1 + \frac{P_{t+1}}{D_{t+1}} + BD_{t+1} - ID_{t+1}\right), \text{ where} \\
BD_{t+1} &\equiv \frac{\sum_n \left[(S_{n,t} - S_{n,t+1}) P_{n,t+1}\right]^+}{\sum_n D_{t+1}}, \; ID_{t+1} \equiv \frac{\sum_n \left[(S_{n,t} - S_{n,t+1}) P_{n,t+1}\right]^-}{D_{t+1}}. \qquad (5)
\end{aligned}
$$

Building on the return decomposition in section 5.1 of Larrain and Yogo (2008), I have broken net repurchases into the sum of its *firm-level* positive parts BD (scaled buybacks) and *firm-level* negative parts ID (scaled issuance), accomplishing two goals. First, we separate buybacks and issuance which may have different levels of predictability owing to different underlying economic forces. Two, we now have only positive-valued variables in (5), so the log of each can be taken.[17]

A log-linear decomposition of the dividend-price ratio follows in the usual way by adapting Cochrane (2005):

$$
\begin{aligned}
1 &= R_{t+1}^{-1} R_{t+1} \\
&= R_{t+1}^{-1} \frac{D_t}{P_t}\frac{D_{t+1}}{D_t}\left(1 + \frac{P_{t+1}}{D_{t+1}} + BD_{t+1} - ID_{t+1}\right) \\
\frac{P_t}{D_t} &= R_{t+1}^{-1} \frac{D_{t+1}}{D_t}\left(1 + \frac{P_{t+1}}{D_{t+1}} + BD_{t+1} - ID_{t+1}\right) \\
pd_t &= -r_{t+1} + \Delta d_{t+1} + \log\left(1 + e^{pd_{t+1}} + e^{bd_{t+1}} - e^{id_{t+1}}\right) \\
&\approx -r_{t+1} + \Delta d_{t+1} + \frac{1}{1 + e^{pd} + e^{bd} - e^{id}}\left[e^{pd}(pd_{t+1} - pd) + e^{bd}(bd_{t+1} - bd) - e^{id}(id_{t+1} - id)\right] \\
\delta_t &\approx r_{t+1} - \Delta d_{t+1} + \rho_\delta \delta_{t+1} - \rho_b bd_{t+1} + \rho_i id_{t+1} + \kappa, \qquad (6)
\end{aligned}
$$

where

$$
\begin{aligned}
&pd_t \equiv \log\left(\frac{P_t}{D_t}\right), \quad bd_t \equiv \log\left(BD_{t+1}\right), \quad id_t \equiv \log\left(ID_{t+1}\right), \quad \delta_t \equiv -pd_t, \\
&pd \equiv \mathbb{E}(pd_t), \quad bd \equiv \mathbb{E}(bd_t), \quad id \equiv \mathbb{E}(id_t), \\
&\rho_\delta \equiv \frac{e^{pd}}{1 + e^{pd} + e^{bd} - e^{id}}, \quad \rho_b \equiv \frac{e^{bd}}{1 + e^{pd} + e^{bd} - e^{id}}, \quad \rho_i \equiv \frac{e^{id}}{1 + e^{pd} + e^{bd} - e^{id}}, \\
&\kappa \equiv \rho_\delta pd + \rho_b bd - \rho_i id.
\end{aligned}
$$

[17] Of course, so long as D_t, P_t, buybacks, and issuance are all nonzero, which I show is true in yearly aggregates.

Taking logs of both sides leads from the third line to the fourth; to go to the fifth line, take a Taylor approximation using $(pd_{t+1}, bd_{t+1}, id_{t+1})$ around (pd, bd, id). Equation 6 is the novel present-value relationship I study, and I refer to bd_{t+1} and id_{t+1} as *buybacks* and *issuance* for simplicity.

The key point is that the very dividend-price ratio used by many preceding papers is, by definition, related to future issuance and buybacks. If we are investigating variation in the dividend-price ratio δ_t but ignoring buybacks and issuance, we have essentially imposed a constraint that those cash-flow expectations don't vary—we will see that the data reject that constraint.

All of the ρ parameters are positive (because e^{id} is much smaller than $1+e^{pd}+e^{bd}$), so if we take time-t expectations of both sides then (6) makes the following statements. News that future returns will be higher *increases* the dividend-price ratio, news that future dividend growth will be higher *decreases* the dividend-price ratio, and the dividend-price ratio *positively* predicts its future value. The preceding are well known both theoretically and empirically. The following present-value statements have not been analyzed, to the best of my knowledge. News that future buybacks will be higher, being (like dividends) cash paid to the household sector, *decreases* the dividend-price ratio. News that future issuance will higher, being cash paid *to the firm sector*, increases the dividend-price ratio.

Finally, we can easily derive a present-value restriction that should approximately hold. Following Cochrane (2008), project both sides of (6) onto δ_t, obtaining

$$1 = \phi_r - \phi_d + \rho_\delta \phi_\delta - \rho_b \phi_b + \rho_i \phi_i \tag{7}$$

for the projection coefficients $(\phi_r, \phi_d, \phi_\delta, \phi_b, \phi_i)'$.

2.3 Per-share and total-equity-payout ratios

Equations 4 and 6 do not say that other dividend-price-ratio decompositions are incorrect. Where are the net repurchases in the familiar dividend-price-ratio decomposition? Rework the original return identity to deliver the familiar

$$\frac{D_{n,t}}{P_{n,t}} \frac{D_{n,t+1}}{D_{n,t}} \left(1 + \frac{P_{n,t+1}}{D_{n,t+1}}\right).$$

There is no S in sight—what happened? The answer is that are deriving a present-value relationship for the *per-share* dividend-price ratio $D_{n,t}/P_{n,t}$.

Notice that the numerator share and denominator share could be subtly different economic objects. $D_{n,t}$ is the dividend paid for a share holding $1/S_{n,t-1}$ ownership of the firm, and $P_{n,t}$ is the price for a share holding $1/S_{n,t}$ ownership of the firm—when the number of adjusted shares $S_{n,t}$ varies, these objects differ from the ownership-stake share of standard macroeconomic models (e.g. Ljungqvist and Sargent, 2018, chap. 13). In this sense, the per-share dividend-price ratio hides the cash flows which occur when the number of adjusted shares changes: the ownership stake of each adjusted share responds to the firm's net repurchase decision, because the household sector always ultimately owns 100% of the firm's equity.[18]

My derivation takes a whole-firm approach to the dividend-price ratio as discussed in Campbell (2018).[19] His chapter 5.3 notes "the Campbell-Shiller formula can be applied in two different ways[, one of which where] the dividend can be interpreted as the total cash paid by the firm to investors, and the price can be interpreted as the total market value of the firm." This observation relates to the statement in Larrain and Yogo (2008) that dividend and payout yields "represent a subtle but important difference between a microeconomic and a macroeconomic view of investment... the portfolio strategy implicit in dividend-price ratio is feasible only at the microeconomic level, whereas the portfolio strategy implicit in equity payout yield is also feasible at the macroeconomic level." I add to this point that there is a subtle but important difference between per-share and aggregate dividend-price ratios: per-share dividend-price ratios involve a ratio between shares whose economic meaning can differ. Relatedly, I am taking a stand on the participation "question at hand" that Koijen and Van Nieuwerburgh (2011) posit: I am analyzing "an investor who participates in every stock repurchase." If the question is how news about future cash flows to the household sector affect aggregate prices, this is a sensible perspective. In aggregate, the household sector always participates.

I emphasize that this paper uses the aggregate dividend-price ratio of Campbell and Shiller (1988), Welch and Goyal (2007), Koijen and Van Nieuwerburgh (2011), and many others—there is nothing new about the predictor variable I study. It is the collection of forecast

[18]Here is another way to think about it. Basic textbook theory says that stock buybacks should affect the price per share even when they do nothing to affect total equity value (e.g. Brealey et al., 2003, chapter 16). Net repurchases, if they have no effect on future firm activities, can be used by firm managers to change per-share earnings and dividends—which they might desire to do if their compensation contracts target those quantities. But such changes in per-share ratios aren't necessarily the same macroeconomic forces upon which we'd like to focus. Moreover, the market's per-share ratio can appear a little unintuitive when shares are bought and issued—see the appendix.

[19]An alternative approach to the one I've taken could be to use a simpler Gordon growth model saying $(D+NETR)/P = R-G$ where $(D+NETR)$ is total cash to shareholders, P is total market value of equity, R is the return, and G is the growth rate of total cash to shareholders. Further break net repurchases $NETR$ into issuance and buybacks, and a standard loglinearization could follow. I am grateful to John Campbell for pointing this out.

targets that is novel—because the decompositions above show that dividends *and* net repurchases should be forecast at the same time. For the question "What moves aggregate stock prices?" with the two possibilities being discount-rate news or cash-flow news, expectations of future dividends *and* future net repurchases both reflect cash-flow news.[20] That is this paper's main idea.

Why not simply rewrite everything in terms of total equity payout, as used in Boudoukh et al. (2007), Larrain and Yogo (2008), Eaton and Paye (2017), and others? There are at least four reasons. First, I want to show that net repurchases drive the aggregate *dividend-price ratio* by definition. Second and related, I aim to stay close to previous literature studying predictability, and thereby stock price volatility. Retaining focus on the dividend-price ratio allows me to directly connect to numerous seminal percursors following Campbell and Shiller (1988). Third, we will be log-linearizing the expression and so require that the cash yield stay positive—but aggregate, total equity payout (even summing the cash flow over twelve months) is not always positive, even in recent data (see Section 3). Fourth, although they are all cash flows between firms and the household (Allen and Michaely, 2003), the economic mechanisms underlying each may be different. Therefore, it is possible that the household sector forms different expectations about future dividend growth, buybacks, and issuance (i.e. they have different degrees of predictability), and we can explore this empirically.

3 Data

For stock returns and cash flows, this paper uses data from CRSP, as have been used in Stephens and Weisbach (1998), Fama and French (2001), Bansal et al. (2005), Dichev (2007), Welch and Goyal (2007), Boudoukh et al. (2007), Larrain and Yogo (2008),Grullon et al. (2011), and Bessembinder (2018), amongst others, to extract distributions from firms. Some of those papers vary on how other data, for instance from Compustat or the Flow of Funds, are used to derive equity payouts other than dividends. This paper uses only CRSP data for its main results.

[20]In broad spirit, this is reminiscent of Aharoni et al. (2013)'s point that per-share empirical analysis did not accurately measure Miller-Modigliani valuation theory. Here I am saying that the per-share theory does not accurately reflect the driving forces of the well-known aggregate dividend-price ratio.

3.1 Basics

The basic idea (for instance in Dichev, 2007), is:

$$\text{(Distribution now)} = \text{(Mkt. cap. past)}[1 + \text{(Return now)}] - \text{(Mkt. cap. now)} \quad (8)$$

This expression hinges on the accuracy of CRSP data in identifying what are *stock* distributions using its cumulative factor to adjust shares, `CFACSHR`, which Campbell and Shiller (1988) argued is carefully constructed. Thereby, CRSP is identifying distributions that are *non-stock, cash* distributions between the firm and household sectors. When I refer to "cash flow" in this paper, I mean distributions calculated via (8). If the "Return now" in (8) is the cum-dividend return `RET`, then the distribution is "total equity payout". If the "Return now" is the ex-dividend return `RETX`, then the distribution is a "net repurchase": a "buyback" occurs in stock-months where the net repurchase is positive, and "issuance" occurs in stock-months where the net repurchase is positive. The difference between total equity payout and net repurchases is ordinary dividends, which is always nonnegative (because `RET` − `RETX` is always nonnegative). These are cash flows that tie together (i.e. are exactly implied by) CRSP's data for `RET`, `RETX`, `PRC`, and `SHROUT` variables.

Via Wharton Research Data Services (WRDS), I download the CRSP Monthly Stock File for the months December 1925 to March 2023.[21] Since I work at the `PERMNO` level, my definition of "firm" could be broad depending on the set of `PERMNO`s chosen—perhaps "stock" would be a better term, but I continue to use the word "firm" as well to highlight the macroeconomic perspective on firm versus household sectors. As I detail below, my benchmark sample uses common stocks and therefore the "firm" label is apt.

3.2 Cash flows

I calculate market capitalization $MKTCAP_{n,t} = \texttt{PRC}_{n,t}\texttt{SHROUT}_{n,t}$. Dividends are computed

$$DIV_{n,t} = (\texttt{RET}_{n,t} - \texttt{RETX}_{n,t})MKTCAP_{n,t-1}. \quad (9)$$

That is, I rely on CRSP's decision on what are ordinary dividends that should be excluded from `RETX` to calculate the total dividends paid out by the firm.

[21] I am using sans serif font to denote internet links that other researchers with WRDS access can select. Once in WRDS, select the CRSP vendor, then Quarterly Update, then Stock / Security Files, then Monthly Stock File.

Determining net repurchases is more involved. An aforementioned reason is SHROUT can change for reasons that do not involve a cash flow between firm and household: preeminent examples would be a stock split or stock dividend. These are the types of events that CRSP captures by its cumulative adjustment factor for shares CFACSHR. Therefore, identifying a net repurchase depends on seeing a change in the *adjusted* shares outstanding.[22]

This means I calculate a nonzero net repurchase for firm-month $(n, t+1)$ only when three conditions are met:

$$\begin{aligned} \texttt{CFACSHR}_{n,t} &= \texttt{CFACSHR}_{n,t+1}, \\ \texttt{SHROUT}_{n,t+1} &\neq \texttt{SHROUT}_{n,t}, \text{ and} \\ (1+\texttt{RETX}_{n,t+1})MKTCAP_{n,t} &\neq MKTCAP_{n,t+1}. \end{aligned} \tag{10}$$

Equation 10 is the one used to measure the net repurchase amount.[23] The first two conditions say that I only calculate (10) when the number of adjusted shares changed, but the cumulative adjustment factor was constant.

This $NETREP_{n,t+1}$ is *the* value of cash distributions that ties together the $\texttt{RETX}_{n,t+1}$, $\texttt{SHROUT}_{n,t}$, $\texttt{SHROUT}_{n,t+1}$, $\texttt{PRC}_{n,t}$ and $\texttt{PRC}_{n,t+1}$ variables in monthly data. This is a point worth emphasizing. If one argues that $NETREP_{n,t}$ is the wrong measure of net distributions (other than ordinary dividends) to the household, then an implication is that $\texttt{RETX}_{n,t}$ and $\texttt{RET}_{n,t}$ do not directly reveal the return received by the household sector on its ownership of all of the firm's equity. One-hundred percent ownership of the firm is claimed by shares outstanding that are ultimately owned by the household, so RETX differs from the change in firm market capitalization because of non-stock (what I am calling cash) distributions between firm and household.

[22] CRSP is careful in their construction of the adjustment factors. Using the example of AT&T's break-up, CRSP data say no net repurchase occurred because the adjusted shares outstanding do not change. In daily data, on February 16 1984, PERMNO 10401 loses 73.42% of its market capitalization measured as $\texttt{SHROUT} \times \texttt{PRC}$, but its $\texttt{RET} = \texttt{RETX} = 1.63\%$, which also equals the change in *adjusted* market capitalization $\texttt{SHROUT} \times \texttt{CFACSHR} \times \texttt{PRC}/\texttt{CFACPR}$ using also the cumulative adjustment factor for price. What has happened is that holders on PERMNO 10401 shares receive new shares in the baby Bells, PERMNOs 66122, 66093, 66026, 66018, 65883, 65875, and 65859. The important thing to note is that this is a *stock* distribution to shareholders, and so *should not* appear in our net repurchase measure because it was not a cash flow between households and firm—which is exactly what the adjusted number of shares tells us.

[23] Note I calculate the value of net repurchases using only RETX, PRC, and SHROUT (subject to the nonzero condition just described). This is done to avoid potentially tricky issues with the cumulative adjustment factors for shares and price, CFACSHR and CFACPR, that may arise when the two are not equal. Furthermore, there are further technical details with using monthly CRSP data, which I discuss further in the appendix.

Consistent with (5), I split $NETREP_{n,t}$ into its positive and negative parts:

$$[NETREP_{n,t}]^{+} \equiv BUY_{n,t} \text{ and } [NETREP_{n,t}]^{-} \equiv ISS_{n,t}. \tag{11}$$

Therefore for each stock-month at least one of these two variables is zero.

My definition of $NETREP$, and therefore ISS and BUY, is broad. ISS is nonzero any month in which the (cumulative factor adjusted) number of outstanding shares increases, and BUY is nonzero any month in which the number decreases. Therefore, ISS not only captures secondary equity offerings, but also equity-based employee pay, which Eisfeldt et al. (2022) and others note has grown in aggregate importance. For this paper's objective, we want such a broad measure of ISS because these are the implied cash flows making sense of `RETX` and the change in $MKTCAP$. Both equity-pay and seasoned equity offerings, ideally, involve cash flow from the household to firm—both of these are *de facto* stock issuance between the firm and household sectors.[24] The situation with BUY is more straightforward: the (cumulative factor adjusted) number of outstanding shares decreases when the firm buys back its own shares, transferring cash to the household sector.

3.3 Samples

I use CRSP monthly data from December 1925 to March 2023, using various sample periods. My benchmark sample includes all US-domiciled common stocks, identified where $10 \leq$ `SHRCD` ≤ 11. Additionally, I split the benchmark *common* stock sample into two subsamples: *nonfinancials* (where `SICCD` < 6000 or `SICCD` ≥ 7000) and *financials* (where $6000 \leq$ `SICCD` $<$ 7000) because previous work on payouts (e.g. Allen and Michaely, 2003) split these apart, implying that financials and nonfinancials may have broadly different payout policies.[25] I also consider a foreign-firm sample comprised of stocks where `SHRCD` $= 12$, but only in Section 5.3.

Aggregation proceeds in the usual way. For any month t I set

$$X_t = \sum_{n \in \mathcal{N}} X_{n,t}, \text{ for } X \in \{DIV, ISS, BUY, PAY\},$$

[24] This point could be sticky. For example, see Dechow et al. (1996) for a discussion of the history and disagreements with expensing equity pay.

[25] Furthermore, in an appendix robustness check I also consider a larger set of all `PERMNO`s that are not American Depositary Receipts (ADRs) (*Non-ADR*) because this is the set described as that used to calculate the CRSP index (see Center for Research in Security Prices, 2021, page 104): it is identified where `SHRCD` < 30 or `SHRCD` ≥ 40.

for a choice of sample $\mathcal{N}$, and total equity payout is calculated $PAY = DIV + BUY - ISS$ which I refer to simply as total payout hereafter. The value-weighted return is

$$R_{t+1} = \frac{\sum_{n\in\mathcal{N}} MKTCAP_{n,t} \times \mathtt{RET}_{n,t+1}}{\sum_{n\in\mathcal{N}} MKTCAP_{n,t}}.$$

Following Kelly and Pruitt (2013) I use observations after the Great Depression as my benchmark sample, because the pre-1940 period has extremely volatile cash-flow observations. Another popular choice is to use post-War data or to instead use all data beginning in 1926, and in robustness checks I consider those samples too.

3.4 Annual variable construction

Dividends display well-known seasonality, and for this reason empirical analysis typically uses yearly variables. The variables I have thus far described are monthly, in so far as their values needed only data from month t and $t-1$ for their construction. Now I construct annual variables that explicitly sum/compound over twelve consecutive months. What I call annual variables will have monthly observations, just as a twelve-month moving average has a new observation each month—my main results use these overlapping observations, but in robustness checks I use nonoverlapping observations instead.

Having extracted monthly observations, it is straightforward to take on board the idea that cash flows should be summed up without imparting return features, something discussed in detail by Binsbergen and Koijen (2010) and Koijen and Van Nieuwerburgh (2011). In particular, to construct a yearly cash-flow variable Koijen and Van Nieuwerburgh (2011) recommend summing up the cash flows using either a zero-rate (the simple sum of each month's cash, as in Campbell and Shiller (1988)) or the risk-free rate (i.e. each month's cash is compounded for each remaining month using the risk-free rate, and then these are summed). For simplicity my benchmark results employ zero-rate summing following Campbell and Shiller (1988), but for robustness I show they are unaffected by instead using the risk-free-rate sums.

The variables used in the empirical analysis are constructed as follows. Log market returns are calculated as

$$r_t \equiv \log\left(\prod_{j=0}^{11} R_{t-j}\right).$$

Dividend growth is calculated

$$\Delta d_t \equiv \log\left(\frac{\sum_{j=0}^{11} DIV_{t-j}}{\sum_{j=0}^{11} DIV_{t-12-j}}\right)$$

as the benchmark results employ zero-rate summing.[26] The (log) dividend-price ratio is calculated

$$\delta_t \equiv \log\left(\frac{\sum_{j=0}^{11} DIV_{t-j}}{MKTCAP_t}\right).$$

The (log) buyback and issuance variables are

$$bd_t \equiv \log\left(\frac{\sum_{j=0}^{11} BUYBACK_{t-j}}{\sum_{j=0}^{11} DIV_{t-j}}\right),$$

$$id_t \equiv \log\left(\frac{\sum_{j=0}^{11} ISSUE_{t-j}}{\sum_{j=0}^{11} DIV_{t-j}}\right).$$

Table I reports summary statistics for monthly and annual variables, for the latter calculating autocorrelations using a 12-month lag. Starting with monthly cash-flow variables in Panel A, we see that DIV, BUY, and ISS have very similar characteristics. They have similar means, with monthly dividends averaging \$10.5B, monthly buybacks higher at \$12.5B, and monthly issuance higher at \$17.3B; cash flows other than dividends have become more important in recent decades (see Grullon and Michaely, 2002) and it is reflected in the means.

Panel B reports annual variables' features. Annual dividends, buybacks, and issuance are always positive, as their minimum values are approximately \$1,900M, \$20M, and \$40M, respectively. On average \$62.3B is paid out to the household over a 12-month period, but total equity payout ($\sum_{j=0}^{11} PAY_{t-j}$) is not always positive. Its minimum value of $-$\$1056B occurs in September 2000 at the end of the tech bubble as aggregate annual issuance (achieving its maximum \$1490B) swamped aggregate dividends and buybacks (\$15B and \$28B, respectively). In fact, total equity payout is negative in about 15% of the months between January 1940 and March 2023. The earliest negative observation is December 1968; for the rest of the 1970s and 1980s payout stays positive. Negative aggregate payout occurs many times during the 1990s, starting in March 1992: in fact, in the years 1994 and 1996–2001 there are *only* negative values. More recently, all but one month of 2021 reported negative aggregate

[26] A simple adjustment of these variable definitions define risk-free-rate sums employed as robustness checks. Replace any monthly cash-flow variable x_{t-j} in the main text with $Z_{t,j}x_{t-j}$ for $Z_{t,j} \equiv \prod_{k=t-j+1}^{t}(1+r_k^f)$ for $j>0$ and $Z_{t,0}=1$, where r_t^f is the risk-free rate for month t.

Table I
Summary statistics

Notes – For the benchmark sample of monthly observations of common stocks over December 1940 to March 2023 (987 observations). p denotes a percentile. AC denotes autocorrelation, using a 1-month lag for Panel A and a 12-month lag for Panel B. Monthly variables are constructed using a single month's data; annual variables are constructed using twelve months in their construction.

	Mean	p_0	p_{10}	p_{50}	p_{90}	p_{100}	AC
			Panel A: Monthly				
DIV_t (\$B)	10.51	0.05	0.25	4.30	31.35	75.44	0.81
BUY_t (\$B)	12.46	0.00	0.00	0.32	43.95	196.22	0.76
ISS_t (\$B)	17.29	0.00	0.01	1.63	51.89	243.96	0.71
			Panel B: Annual				
$\sum_{j=0}^{11} DIV_{t-j}$ (\$B)	122.57	1.90	3.77	59.69	390.14	618.71	0.99
$\sum_{j=0}^{11} BUY_{t-j}$ (\$B)	145.25	0.02	0.08	10.87	535.71	1004.81	0.94
$\sum_{j=0}^{11} ISS_{t-j}$ (\$B)	205.51	0.04	1.15	42.59	582.38	1489.69	0.87
$\sum_{j=0}^{11} PAY_{t-j}$ (\$B)	62.31	−1055.87	−54.27	8.90	453.46	1083.19	0.73
δ_t	−3.56	−4.69	−4.20	−3.51	−2.94	−2.43	0.93
r_t (%)	10.57	−55.36	−11.30	12.80	29.15	51.14	−0.10
Δd_t (%)	7.02	−21.65	−0.42	6.75	15.10	30.30	0.23
bd_t (%)	−171.94	−639.31	−396.57	−164.20	41.38	77.72	0.90
id_t (%)	−27.31	−388.18	−126.83	−28.63	86.17	226.77	0.86

payout. Therefore, a total-equity-payout-to-price ratio takes negative values throughout the post-Great-Depression sample, which prohibits a log-linear decomposition and supports my focus on the dividend-price ratio.

The last five variables appear in the present-value relationship (6) and are used in the empirical analysis. It jumps out that of the three cash flow variables, Δd_t is clearly the least volatile, as it ranges from −22% to 30%. Buybacks and issuance range −172% to 78% and −388% to 226% respectively, which is a first statistical indication that their news might meaningfully drive dividend-price variation. Their negative means (divide the table values by 100) will lead the present-value constants ρ_b and ρ_i to be much less than 1, working against a contribution to dividend-price ratio variation. Considering the present-value constraint (7), buybacks and issuance will only really matter if their projection coefficients ϕ_b and ϕ_i are quite large relative to ϕ_r and ϕ_d.

Another observation is that δ_t, bd_t, and id_t are all rather persistent variables. However, the augmented Dickey-Fuller test rejects at the 5% level for all of them, implying we can reasonably assume covariance stationarity and spurious-regression concerns are abated. Nevertheless, the reader might suspect that predictability of bd_t and id_t by δ_{t-12} might unduly stem from the former's persistence and not from significant predictive information in the latter; however Section 5.2 will report that δ_{t-12} significantly predicts bd_t and id_t even while controlling for the latters' lagged values. This opens up an interesting possibility that the notorious persistence of the dividend-price ratio could result from the persistence of near-term expected cash flows, perhaps in addition to longer-term expected returns, should ϕ_b and ϕ_i be strongly significant.

4 Empirical results

This section begins by discussing the empirical framework. The following subsections present the benchmark results, results for nonfinancial and financial firm samples, robustness analysis, and explores issuance in more detail.

4.1 Set-up

Define $\boldsymbol{x}_t = (\delta_t, r_t, \Delta d_t, bd_t, id_t)'$, $\boldsymbol{\phi}_1 = (\phi_\delta, \phi_r, \phi_d, \phi_b, \phi_i)'$, and $\boldsymbol{\phi}_0$ the intercepts. A restricted VAR is estimated by the following moments

$$\mathbb{E}\left[\begin{array}{c} \boldsymbol{x}_{t+12} - \boldsymbol{\phi}_0 - \boldsymbol{\phi}_1 \delta_t \\ (\boldsymbol{x}_{t+12} - \boldsymbol{\phi}_0 - \boldsymbol{\phi}_1 \delta_t)\, \delta_t \end{array}\right] = 0. \tag{12}$$

To these restricted-VAR equations, let us add the present-value constraint (7). Hence, we have an over-identified GMM system, with eleven moment conditions for the ten elements in $(\boldsymbol{\phi}_0', \boldsymbol{\phi}_1')'$. We can directly test (7) or else use it in estimation (as in Larrain and Yogo, 2008). When I impose that the ten VAR moments hold, I am estimating $(\boldsymbol{\phi}_0', \boldsymbol{\phi}_1')'$ via OLS and hence I label this specification *OLS*. I also consider three alternative empirical specifications: equally-weighting all moments (*identity*), and two-step versions of the preceding (*two-step OLS* and *two-step identity*). In all cases, an overidentification test of the model is available via Hansen (1982)'s J test. Heteroskedasticity- and autocorrelation robust (HAR) spectral density estimates follow Newey and West (1987) with a bandwidth one-year *more* than the degree of overlap in the observations: in the baseline specification, this means 24 Newey-West lags.[27]

What do we expect the signs of of the projection coefficients $\boldsymbol{\phi}_1$ to be? Let's start with

[27] The main results are robust to instead using Hodrick (1992) standard errors.

the cash-flow variables of main interest. News that future buybacks will be higher (being cash paid *to the household*, like dividends) should raise prices today, thereby decreasing the dividend-price ratio: we expect a negative ϕ_b. News that future issuance will be higher (being cash paid *to the firm, from the household*) should lower prices today, thereby increasing the dividend-price ratio: we expect a positive ϕ_i. Meanwhile the standard intuition continues to apply to returns, dividend growth, and the dividend-price ratio itself: we expect a positive ϕ_r, a negative ϕ_d, and a positive ϕ_δ.

Via GMM we can test cross-equation null hypotheses. An interesting one is whether buybacks and issuance are forecasted differently; that is, $H_0 : \phi_b = \phi_i$. Of course, the present-value logic says this hypothesis should be rejected, as theoretically those projection coefficients have opposite signs. On notation: I will sometimes refer to the t-statistic for some parameter x as $t(x)$, and the p-value for some statistic x as $p(x)$. For ease I do not put a hat "ˆ" on estimates.

4.2 Benchmark

Table II presents benchmark results using all common stocks, on the sample December 1940 to March 2023, using 976 overlapping monthly observations. Panel A reports estimates when the OLS moments hold exactly. Just as in previous literature, the results suggest that the dividend-price ratio strongly and significantly forecasts future returns ($\phi_r = 0.09$, $t = 3.0$) to an economically-significant degree ($R^2 = 8.2\%$), but does not forecast future dividend growth ($\phi_d = 0.00$, $t = 0.1$).

The novel results are that the dividend-price ratio *also* strongly and significantly forecasts future buybacks ($\phi_b = -2.83$, $t = -12.8$), and issuance ($\phi_i = -1.51$, $t = -8.2$). The level of predictability is very economically significant as the variables' R^2s are around 62%. Hence, we easily reject the constraint that net repurchases do not drive the dividend-price ratio. The dividend-price ratio negatively predicts future buybacks, as that positive expected cash flow to the household (like dividends) raises the price. The dividend-price ratio also negatively predicts future *id*—this is a surprising result. We had expected a positive coefficient, as news about increased future issuance would lower today's prices and thereby increase the dividend-price ratio. Instead, we see a robustly significant negative prediction. Because $p(W) = 0.0\%$ we can reject that $\phi_b = \phi_i$, but not like how we expected to *a priori*.

All of the main conclusions hold in the remaining panels of Table II.[28] Panel B weights

[28]The predictability of buybacks and issuance may appear to contrast with what Eaton and Paye (2017)

Table II
Dividend-price-ratio forecasts

Notes – Benchmark results for the common stock sample. Each row presents estimates forecasting the variable listed in the first column, from the restricted VAR. Variables names are: δ is the log dividend-price ratio, r is the log value-weighted return, Δd is the log dividend growth rate, bd is log scaled buybacks, and id is log scaled issuance. There are 976 monthly observations over December 1940 to March 2023. In parentheses underneath parameter estimates are t-statistics from HAR (Newey-West) standard errors using 24 lags. The overidentification-test statistic is labeled J and reported alongside its p-value. The Wald-statistic of a null hypothesis that $\phi_i = \phi_b$ is labeled W and reported alongside its p-value. In the same column is reported R^2 for each forecast variable, and p-values for the J and W statistics, all in percentage. The row PV calculates the right-hand-side of (7). Each panel is a different GMM estimation choice: Panel A is one-step imposing the VAR parameters are the OLS estimate, and Panel B is one-step equally-weighting all moments; Panels C and D are two-step where the first-step estimates come from the preceding, respectively.

	Panel A: OLS				*Panel B: Identity*		
	δ_{t-12}	cons	$R^2/p(\%)$		δ_{t-12}	cons	$R^2/p(\%)$
δ_t	0.924	−0.287	86.6	δ_t	0.928	−0.272	86.6
	(27.17)	(−2.39)			(27.52)	(−2.28)	
r_t	0.094	0.442	8.2	r_t	0.098	0.457	8.2
	(2.98)	(4.13)			(3.11)	(4.26)	
Δd_t	0.002	0.078	0.0	Δd_t	−0.002	0.062	−0.1
	(0.12)	(1.14)			(−0.10)	(0.90)	
bd_t	−2.833	−11.748	61.5	bd_t	−2.833	−11.748	61.5
	(−12.76)	(−13.83)			(−12.76)	(−13.83)	
id_t	−1.508	−5.601	62.5	id_t	−1.508	−5.601	62.5
	(−8.19)	(−8.22)			(−8.19)	(−8.22)	
J	8.617		0.3	J	7.044		0.8
W	21.791		0.0	W	21.794		0.0
PV	0.987			PV	1.000		

	Panel C: Two-step OLS				*Panel D: Two-step Identity*		
	δ_{t-12}	cons	$R^2/p(\%)$		δ_{t-12}	cons	$R^2/(\%)$
δ_t	0.946	−0.207	86.5	δ_t	0.943	−0.217	86.5
	(28.53)	(−1.77)			(28.33)	(−1.85)	
r_t	0.090	0.425	8.2	r_t	0.089	0.422	8.1
	(2.84)	(3.99)			(2.81)	(3.95)	
Δd_t	0.011	0.113	−0.7	Δd_t	0.006	0.097	−0.4
	(0.60)	(1.68)			(0.34)	(1.43)	
bd_t	−2.520	−10.620	60.7	bd_t	−2.646	−11.077	61.2
	(−13.55)	(−14.49)			(−13.37)	(−14.37)	
id_t	−1.361	−4.999	61.1	id_t	−1.311	−4.840	60.9
	(−7.53)	(−7.58)			(−7.31)	(−7.37)	
J	7.213		0.7	J	6.276		1.2
W	17.877		0.0	W	21.867		0.0
PV	0.997			PV	1.000		

all moments equally. Coefficient estimates are only modestly affected, in fact with ϕ_b, ϕ_i unchanged.[29] GMM moves ϕ_r modestly to improve the fit of the present-value moment, which we can see as follows. Panel A's $p(J) = 0.3\%$, so we strongly reject the present-value restriction. Consistent with this, when we calculate the right-hand side of (7) using Panel A's estimates, which is called PV in the table, we get 0.987 instead of 1. Panel B instead weights this moment in estimation and GMM makes sure it is exactly satisfied. Now the remaining moments do not hold exactly, but are closer so that the $p(J)$ rises mildly to 0.8%.

Panels C and D present two-step GMM estimates (using the Panel A or Panel B estimates as their first step, respectively) that do very little to change the story. Common across them is that ϕ_r, ϕ_b, and ϕ_i move a bit closer to zero, while the ϕ_δ parameter moves a bit higher though remaining far from significant. In both cases $p(J)$ rises, in Panel C to 0.7% and in Panel D to 1.2% so that we fail to reject the overidentifying restriction at the 1% level. Consistent with this, Panel C's PV rises to 0.997 (of course Panel D's stays exactly at 1). Hence, return *and* cash-flow predictability looks strong across all GMM specifications.

4.3 Nonfinancials and financials

It is interesting to look across the nonfinancial and financial firm subsamples as the latter are often broken out in studies of corporate payout (e.g. Allen and Michaely, 2003). There are some noteworthy distinctions, as well as similarities. As in the benchmark sample, I find the differences across GMM specifications are not large. Hence, I opt to report just the OLS and two-step-identity estimates.

Between the nonfinancial and financial stocks, we firstly note that the degree of return predictability is similar. Secondly, and more to this paper's point, the degree of dividend-growth predictability is very different. In nonfinancials we see an absence of predictability, as in the benchmark sample. However, in financial stocks we see a $\phi_d = -0.1$ whose magnitude and statistical significance ($t = -3.7$) is nearly as large as the return coefficient's. This means that financial stocks, even when focusing on ordinary dividends, demonstrate a significant degree of cash-flow predictability, which to the best of my knowledge is a novel result.

Similar across the firm subsamples, buybacks and issuance continue to be statistically and

find (their Table 4), but note that their predictor (YLD in their notation) varies as the forecasted cash flow varies and so they don't use the dividend-price ratio to forecast net repurchases.

[29]The R^2 for Δd_t turns negative, and in fact is so for the remaining panels. This happens because $\boldsymbol{\phi}$ is no longer the OLS estimate, and when GMM is not solving least squares we are not assured the in-sample R^2 is non-negative.

Table III
Dividend-price-ratio forecasts, nonfinancials and financials

Notes – Results separately for nonfinancial and financial stocks, for the GMM specification given by the panel heading. Other than these differences, further details are as in Table II.

	Panel A: Nonfinancial, OLS				*Panel B: Nonfinancial, two-step identity*		
	δ_{t-12}	cons	$R^2/p(\%)$		δ_{t-12}	cons	$R^2/p(\%)$
δ_t	0.931	−0.263	87.6	δ_t	0.955	−0.179	87.5
	(25.73)	(−2.08)			(27.29)	(−1.46)	
r_t	0.092	0.435	8.5	r_t	0.085	0.411	8.5
	(2.79)	(3.89)			(2.58)	(3.66)	
Δd_t	0.007	0.093	0.3	Δd_t	0.014	0.118	0.0
	(0.38)	(1.38)			(0.76)	(1.78)	
bd_t	−2.786	−11.630	62.0	bd_t	−2.550	−10.769	61.6
	(−12.38)	(−13.52)			(−13.62)	(−14.64)	
id_t	−1.459	−5.485	63.7	id_t	−1.354	−5.030	62.6
	(−8.40)	(−8.50)			(−7.76)	(−7.86)	
J	7.204		0.7	J	6.525		1.1
W	22.108		0.0	W	20.550		0.0
PV	0.989			PV	1.000		

	Panel C: Financial, OLS				*Panel D: Financial, two-step identity*		
	δ_{t-12}	cons	$R^2/p(\%)$		δ_{t-12}	cons	$R^2/p(\%)$
δ_t	0.824	−0.635	71.2	δ_t	0.817	−0.656	71.2
	(22.24)	(−4.66)			(22.66)	(−4.90)	
r_t	0.117	0.520	8.2	r_t	0.126	0.548	8.1
	(4.37)	(5.90)			(5.15)	(6.71)	
Δd_t	−0.093	−0.232	6.5	Δd_t	−0.100	−0.249	6.3
	(−3.30)	(−2.15)			(−3.69)	(−2.37)	
bd_t	−2.068	−8.961	25.9	bd_t	−1.982	−8.767	25.5
	(−5.87)	(−7.42)			(−5.84)	(−7.31)	
id_t	−2.192	−8.020	40.1	id_t	−2.023	−7.379	39.8
	(−3.63)	(−3.60)			(−3.38)	(−3.34)	
J	0.752		38.6	J	0.615		43.3
W	0.024		87.8	W	0.002		96.1
PV	0.988			PV	1.000		

economically significant. Therefore the key conclusion we obtained from the benchmark sample is true in both nonfinancials and financials, separately. Yet some interesting differences emerge, particularly by focusing on financials. In them, the degree of net-repurchase predictability is much reduced as the R^2 almost halves, almost as if the increased dividend predictability comes at the cost of reduced net-repurchase predictability. Moreover, the ϕ_b and ϕ_i coefficient estimates are much closer, leading the Wald test of $\phi_b = \phi_i$ to accept.

Finally, we see differences in the model's overidentification test. In nonfinancials we see a failure to reject at the 1% level, just as for the benchmark. But in financials we see $p(J)$ is bigger than 10% and the overidentification condition is not rejected. One could say that the present-value model is most consistent with data displaying *both* return and cash-flow predictability.

Overall, splitting the sample into nonfinancial and financial firms does nothing to change the bottom line that buybacks and issuance are robustly forecasted by the dividend-price ratio.

4.4 Robustness

To persuade the reader that the previous subsections' main conclusions are robust, in Table IV are reported key estimates across a variety of different specifications. I report projection coefficients $(\phi_r, \phi_d, \phi_b, \phi_i)$ for the in-sample rows, and for the out-of-sample rows I report R^2 in percentage.

To begin with we consider different sample periods, as Koijen and Van Nieuwerburgh (2011) note they vary the qualitative and quantitative conclusions of dividend-price-ratio regressions forecasting future returns and dividend growth. Is this also true for issuance and buybacks? In the full 1926–2023 sample we continue to see buybacks and issuance robustly forecasted. Dividend growth broadly becomes predictable, in line with Koijen and Van Nieuwerburgh (2011)'s results, driven by extremely variable Great Depression observations (Kelly and Pruitt, 2013, suggested these be dropped from aggregate cash-flow predictions for this reason). In novel results, in the full sample financial-firm return predictability remains robust and the magnitude of dividend-growth predictability increases. Turning to the 1946–2009 sample featured in Cochrane (2011) and Koijen and Van Nieuwerburgh (2011), once again we see significant buyback and issuance predictability and the main takeaways are not much different than we have seen in the 1940–2022 sample used in previous tables. In the third row of each panel, restricting the sample to the most recent half (roughly) of data in 1970–2022 also changes little.

Table IV
Dividend-price-ratio forecasts, robustness

Notes – Estimates across a variety of specifications for all common, nonfinancial, or financial stocks: in-sample using the two-step identity GMM estimator (unless otherwise stated) and t-statistic; out-of-sample using OLS and the Clark and McCracken (2005) encompassing test. The column "Values" tells us what are the numbers in the row: ϕ estimates for the in-sample rows, and out-of-sample R^2 for the out-of-sample rows. Statistical significance denoted: * for 10% level, ** for 5% level, and *** for 1% level. Row names label the deviations from Tables II and III as follows. 1926–2022: using the full sample period. 1946–2009: using a familiar post-war sample period. 1970–2023: using the latter half of the data. Non-overlapping obs.: using non-overlapping annual observations, and 2 Newey-West lags (December to December, 81 observations). Risk-free-rate sum: sum up annual cash flows using the compounded risk-free rate. Two-year ahead, OLS: using δ_{t-24} to forecast the time-t target, and 36 Newey-West lags (955 observations). Out-of-sample, OLS, R^2: reports the out-of-sample R^2 (%) if positive, with statistical significance from the Clark and McCracken (2001) ENC-NEW test using a Newey-West adjustment as suggested by Clark and McCracken (2005) (using 24 lags); the out-of-sample forecasts are generated on a recursive window, starting with January 1981, making sure to leave a twelve-month gap between the training sample's end and the target-variable realization being forecast.

	Values	r	Δd	bd	id
			Panel A: Common		
1926–2023	ϕ	0.05	−0.07*	−2.46***	−1.55***
1946–2009	ϕ	0.12***	0.02	−2.34***	−1.36***
1970–2023	ϕ	0.11***	0.00	−2.05***	−1.19***
Non-overlapping obs.	ϕ	0.10***	0.00	−2.58***	−1.32***
Real	ϕ	0.07**	−0.01	−2.61***	−1.30***
Risk-free-rate sum	ϕ	0.09***	0.01	−2.59***	−1.30***
Two-year ahead, OLS	ϕ	0.08**	0.02	−2.82***	−1.42***
Out-of-sample, OLS	R^2	< 0	< 0	68.00***	75.89***
			Panel B: Nonfinancial		
1926–2023	ϕ	0.04	−0.06	−2.41***	−1.55***
1946–2009	ϕ	0.11***	0.02	−2.29***	−1.35***
1970–2023	ϕ	0.10**	0.01	−2.10***	−1.17***
Non-overlapping obs.	ϕ	0.09***	0.02	−2.43***	−1.27***
Real	ϕ	0.07*	−0.01	−2.52***	−1.34***
Risk-free-rate sum	ϕ	0.08**	0.01	−2.51***	−1.34***
Two-year ahead, OLS	ϕ	0.08**	0.02	−2.78***	−1.40***
Out-of-sample, OLS	R^2	< 0	< 0	68.4***	78.9***
			Panel C: Financial		
1926–2023	ϕ	0.08**	−0.15***	−1.76***	−2.45***
1946–2009	ϕ	0.15***	−0.06**	−1.92***	−1.63***
1970–2023	ϕ	0.10***	−0.12**	−1.61***	−0.74***
Non-overlapping obs.	ϕ	0.12***	−0.11***	−1.74***	−1.87***
Real	ϕ	0.11***	−0.12***	−1.95***	−2.00***
Risk-free-rate sum	ϕ	0.12***	−0.10***	−1.95***	−1.97***
Two-year ahead, OLS	ϕ	0.14***	−0.05	−2.30***	−1.43***
Out-of-sample, OLS	R^2	< 0	< 0	33.1***	29.6***

With respect to cash flows, little is changed by the next three rows. In the rows labeled "Nonoverlapping obs." I instead use non-overlapping December-to-December observations in our estimation, as do Cochrane (2008, 2011) and Koijen and Van Nieuwerburgh (2011), and find very similar results.[30] Converting everything to real values, as does the row labeled "Real", also does very little to change the previous cash-flow results. Meanwhile, using real returns diminishes some return predictability in the benchmark and nonfinancial firm samples. Using the compounded risk-free rate to sum up cash flows, as does row "Risk-free-rate sum", does very little to change matters.

The next row of each panel considers a longer-horizon forecast: instead of using δ_{t-12} we use δ_{t-24}, thus producing a two-year ahead forecast.[31] This follows from iterating the present-value restriction (see Cochrane, 2008, or below in Section 5) showing us that news about all future periods could in principle matter to prices now. Returns, buybacks, and issuance remain significantly predictable.[32] Meanwhile, the dividend-growth predictability in panel C vanishes, which could indicate that financial firms' dividend-growth news is primarily near-term in nature, possibly a manifestation of the importance of short-term dividend-growth expectations for S&P 500 firms highlighted by De La O and Myers (2021).

Finally, the last row of each panel considers out-of-sample evidence. From Welch and Goyal (2007) and Kelly and Pruitt (2013) we would suspect that the aggregate dividend-price ratio will not significantly forecast returns and dividend growth out-of-sample. Indeed the evidence here confirms those results: their out-of-sample R^2 is negative, meaning that the recursively-estimated dividend-price-ratio forecast does *worse* than simply using the recursively-estimated mean.[33]

But what about the net repurchases upon which we've focused? In stark contrast, the results for buybacks and issuance are very much in line with the in-sample evidence we've so far

[30]Since the annual observations do not overlap, I am conservative and employ 2 Newey-West lags; little is changed by using only 1 or modestly more.

[31]We do not further accumulate any variable: the target remains an annual variable, albeit one realized an additional year after the dividend-price ratio is realized.

[32]Consistent with the idea that objects farther in the future are harder to forecast, the R^2s uniformly drop (not reported). Also, I increase the number of Newey-West lags to 36 to make sure we account for residual autocorrelation.

[33]Following Kelly and Pruitt (2013), I emphasize that these out-of-sample results are not somehow better *per se* than the in-sample results we have seen so far. Instead, the point is to see if the out-of-sample and in-sample conclusions agree, because the former tell us something about the small-sample bias (overfit) which may have been present in the latter. In addition, using slope- and fitted-value restrictions as suggested in Campbell and Thompson (2008) does improve the out-of-sample R^2 of return predictions as in that paper, but not enough for the statistics to turn positive. This difference can be attributed to my sample beginning in 1940 (theirs starts in 1872), out-of-sample forecasts beginning in 1981 (theirs start in 1927), and my sample ending in 2023 (theirs ends in 2005).

seen. Across the panels, we see robust predictability for all three variables, significant at the 1% level using the Clark and McCracken (2001) ENC-NEW test statistic of forecast encompassing.[34]

In summary, Table IV shows us that the buyback and issuance results in Tables II and III are robust.[35] Moreover, that puzzling negative sign for ϕ_i is a widespread feature of the data. In the benchmark and nonfinancial firm samples, return predictability can sometimes be weak and dividend growth is nonexistent, consistent with the well-known existing results (e.g. Koijen and Van Nieuwerburgh, 2011). More novel (to the best of my knowledge) is the fact that return and dividend-growth in-sample predictability is strong in financial firms.

4.5 Stock-sale issuance

In Tables II-IV we have seen robust evidence of a surprising result: the dividend-price ratio *negatively* predicts future issuance. By equation 6's present-value logic this shouldn't be the case. Issuance is a negative cash flow from the household's perspective. Yet the evidence that $\phi_i < 0$ is strong and robust. What is happening?

Now I show that the issuance measured by (11) *does contain* these negative cash flows, but not exclusively. The negative estimates for ϕ_i come from the part of issuance that is not related to companies' reported sale of stock. To proceed, via WRDS I access Compustat Daily Updates - Fundamentals Quarterly and extract the Sale of Common and Preferred Stock (`SSTKY`) for every observation in the database; `SSTKY` is available (i.e. has a nonzero value for some firm) starting in 1970Q3—therefore I use the 1970–2022 sample.[36] I then project issuance onto the aggregated series (the appendix contains further detail) and proceed with construction of log scaled issuance variables as in Section 3.4. I refer to the eventual variable

[34] The nested model to which my ENC-NEW statistic refers is the recursively-estimated mean, so that R^2 and ENC-NEW measure a comparison against the same reference model. To account for overlapping forecasts, I use Newey-West with 24 lags to estimate the statistic's denominator, as suggested in Clark and McCracken (2005) and following Kelly and Pruitt (2013). When I say the forecast is "significant", it is with respect to this reference model.

[35] In the appendix I show that ϕ_b and ϕ_i continue to be statistically significant for further specifications. Using the break-adjusted dividend-price-ratio predictor suggested in Lettau and Van Nieuwerburgh (2007) continues to show significant buyback and issuance predictability. Furthermore, in the appendix I report results using foreign-incorporated firms, which I employ in the variance decompositions presented in Section 5.3 but do not report as rows in Table IV to save space. Finally, I show the main conclusions are robust to using the sample of all non-ADR stocks, which Center for Research in Security Prices (2021) describes as the universe used for the CRSP value-weighted return.

[36] So long as this series (when aggregated) is highly correlated with the sale of common stock only, it will work for my purpose. Compustat also includes the Sale of Common Stock (`SCSTKCY`) variable, but it is only available starting 1999Q4.

coming from the projection as *stock-sale issuance* denoted sid_t.

By its projection on stock sales, issuance's correlation with buybacks switches signs. Over the 1970–2023 sample, the correlation between bd and id is 0.68. But the correlation between bd and sid is -0.61: about as strong in the opposite direction.[37] Meanwhile, the projection little changes issuance's negative correlation with dividends (from -0.28 to -0.14) and does not boost the magnitude of issuance's correlation with returns (from 0.09 to -0.10).

Table V reports that stock-sale issuance is *positively* predicted by the dividend-price ratio, exactly as present-value intuition would tell us. In the benchmark sample, we see that $\phi_i = 0.39$ is significantly positive ($t = 2.9$). Similar stock-sale issuance results obtain for in the nonfinancial and financial subsamples, with the former a little smaller ($\phi_i = 0.28$, $t = 2.0$) and the latter a little larger ($\phi_i = 0.48$, $t = 3.1$). The return coefficient ϕ_r loses 5%-significance in the nonfinancial samples.

We conclude from Table V that issuance linked to firms' reported stock sales behaves just as issuance in the present-value relationship (6) says it should.[38] The projection coefficient of stock-sale issuance on dividend-price is positive and significant. It is noteworthy that buyback and stock-sale issuance expectations do not cancel each other out within the dividend-price ratio. Compare this result to Lettau and Ludvigson (2005) who argue that expected returns and dividend growth have offsetting effects in the dividend-price ratio. It appears that buyback and issuance expectations are like return expectations—robust drivers of aggregate stock prices—in contrast to dividend-growth expectations. Moreover both returns and stock-sale issuance are positively predicted by the dividend-price ratio, though their contemporaneous correlation is slightly negative, suggesting that their predictable parts are distinct.

What is the remaining issuance and why is it negatively predicted? This question deserves further study, but an answer is not required to accomplish this paper's goals. We can remain agnostic about what the predictive coefficient signs *should* be. As stated above, the net repurchases I measure in CRSP are the implied cash flows exactly implied by returns, prices, and outstanding shares—they are worthwhile to analyze regardless. Nevertheless, to make sure my conclusions are robust to a present-value condition that $\phi_i > 0$, I will also calculate variance decompositions using this stock-sale issuance variable sid and report results.

These dividend-price ratio variance decompositions are what we have been building up to,

[37] A full correlation table is in the appendix, Table A.5.

[38] One should note, these observations about issuance cash flows would not have been noticeable had I focused simply on total equity payout—another argument in favor of focusing on the dividend-price ratio.

Table V
Dividend-price-ratio forecasts, stock-sale issuance

Notes – Results using stock-sale issuance, the two-step identity estimator, on the sample 1970–2022. $PV = 1$ in all panels. Other than these differences, further details are as in Table II.

	Panel A: Benchmark				*Panel B: Nonfinancial*		
	δ_{t-12}	cons	$R^2/p(\%)$		δ_{t-12}	cons	$R^2/p(\%)$
δ_t	0.860	−0.549	83.5	δ_t	0.879	−0.475	85.2
	(21.32)	(−3.62)			(20.18)	(−2.94)	
r_t	0.088	0.430	7.6	r_t	0.084	0.421	8.3
	(2.20)	(3.03)			(1.91)	(2.67)	
Δd_t	0.003	0.074	−0.1	Δd_t	0.010	0.102	0.4
	(0.14)	(1.00)			(0.49)	(1.36)	
bd_t	−2.108	−8.515	53.8	bd_t	−2.123	−8.583	54.5
	(−8.22)	(−8.15)			(−8.01)	(−7.94)	
sid_t	0.393	2.210	11.0	sid_t	0.282	1.731	5.2
	(2.92)	(4.55)			(2.03)	(3.48)	
J	12.681		0.0	J	12.601		0.0
W	54.268		0.0	W	47.192		0.0

	Panel C: Financial		
	δ_{t-12}	cons	$R^2/p(\%)$
δ_t	0.701	−1.146	60.7
	(12.84)	(−5.33)	
r_t	0.084	0.406	4.3
	(2.79)	(3.57)	
Δd_t	−0.142	−0.451	3.3
	(−2.37)	(−1.85)	
bd_t	−1.596	−6.558	31.7
	(−7.28)	(−7.51)	
sid_t	0.480	2.755	18.3
	(3.07)	(5.03)	
J	6.689		1.0
W	39.546		0.0

and are discussed next.

5 Variance decomposition of the dividend-price ratio

The preceding section presents robust evidence that the aggregate dividend-price ratio forecasts future buybacks and issuance, with more modest evidence of return predictability. Now I build on those results to measure the drivers of aggregate stock prices.

5.1 Long-run coefficients

One way to attribute dividend-price-ratio movement to discount-rate versus cash-flow forces is to follow Cochrane (2008). Iterate forward our present-value identity (7)

$$\delta_t = \mathbb{E}_t \sum_{j=1}^{\infty} \rho_\delta^{j-1} \left(r_{t+j} - \Delta d_{t+j} - \rho_b b d_{t+j} + \rho_i i d_{t+j} \right).$$

After some algebra,[39] we have the relationship

$$1 = \frac{\phi_r - \phi_d - \rho_b \phi_b + \rho_i \phi_i}{1 - \rho_\delta \phi_\delta} \equiv \phi_r^{lr} - \phi_d^{lr} - \phi_b^{lr} + \phi_i^{lr}$$

which delivers *long-run* coefficients implied by our restricted VAR.

When these long-run coefficients are obtained from a GMM estimate that weights the present-value restriction, we are assured that they add up to 1. However, they do not measure the effect of orthogonal forces. Therefore, to use the long-run coefficients in a variance decomposition, I take $|\phi_r^{lr}|$ to measure the discount-rate news portion, $|\phi_d^{lr}| + |\phi_b^{lr}| + |\phi_i^{lr}|$ to measure the cash-flow news portion, and express each in proportion of their sum. I use estimates from the two-step identity estimates that impose the present-value restriction, and call this decomposition the *LRC* approach in Table VII.

[39]That is: multiply both sides by $\delta_t - \mathbb{E}(\delta_t)$; take the unconditional expectation of both sides, yielding the variance of δ_t on the left; divide both sides by the variance of δ_t; recognize slope coefficients as covariances divided by the predictor variances; and finally impose our restricted VAR(1) specification.

Table VI
Unrestricted VAR estimates

Notes – OLS estimates of an unrestricted VAR on the benchmark sample. In parentheses are t-statistics using Newey-West standard errors with 24 lags.

	Δd_{t-12}	bd_{t-12}	id_{t-12}	r_{t-12}	δ_{t-12}	cons	R^2
Δd_t	0.254	−0.006	0.003	0.132	−0.013	−0.018	16.0
	(2.48)	(−1.01)	(0.35)	(2.90)	(−0.58)	(−0.20)	
bd_t	−0.259	0.756	−0.077	0.993	−0.825	−3.409	83.5
	(−0.34)	(11.51)	(−0.66)	(2.52)	(−2.90)	(−3.00)	
id_t	−0.574	0.045	0.631	0.528	−0.324	−1.158	77.1
	(−1.21)	(1.12)	(7.19)	(1.52)	(−2.20)	(−1.97)	
r_t	−0.286	0.032	−0.012	−0.043	0.170	0.787	14.2
	(−1.90)	(2.14)	(−0.49)	(−0.53)	(3.53)	(4.37)	
δ_t	0.601	−0.025	0.005	0.172	0.852	−0.644	87.9
	(2.62)	(−1.59)	(0.20)	(1.81)	(15.15)	(−2.98)	

5.2 Structural VAR

Another way to attribute dividend-price-ratio variation is a structural VAR approach. I estimate an unrestricted version of the VAR we have heretofore analyzed. Then I adopt assumptions that partially identify structural shocks from the reduced-form estimates. At the cost of an identifying assumption, we have discount-rate and cash-flow shocks separated into orthogonal pieces of news, and we can then take a long-run variance decomposition—this is called the *SVAR* approach in Table VII. Hence, this structural-VAR decomposition is distinct from the long-run-coefficient decomposition in two main ways: using unrestricted VAR (therefore multivariately estimated) coefficients, and the identification assumption leading to orthogonal shocks.

To begin with, Table VI shows estimates of an unrestricted VAR on the benchmark sample. I have reordered the vector to $\tilde{\boldsymbol{x}}_t \equiv (\Delta d_t, bd_t, id_t, r_t, \delta_t)'$ to facilitate the identification discussion below. The dividend-price ratio, now in the last equation, is only significantly predicted by itself. The dividend-price ratio continues to significantly predict returns ($t = 3.5$), but not dividend growth. Interestingly, we see that returns are also positively predicted by buybacks ($t = 2.1$) while dividend growth is positively predicted by returns ($t = 2.9$), so that their R^2s are both economically significant (at 14.2% and 16%, respectively).

Turning to the predictors of buybacks and issuance, because of their persistence it is unsurprising that their own lags are significant. Buybacks predicts itself with a coefficient of 0.76

($t = 11.5$), and issuance itself with a coefficient of 0.63 ($t = 7.2$). More noteworthy is the fact that the dividend-price ratio *continues to* significantly predict them, with $t(\phi_b) = -2.9$ and $t(\phi_i) = -2.2$, despite now controlling for the persistence in each. Section 3.4 reported Dickey-Fuller results showing that *bd* and *id* are covariance stationary, in theory abating spurious-regression concerns. Table VI provides strong evidence that the dividend-price ratio forecasts buybacks and issuance due to significant predictive information, not merely as a statistical artifact of the latter's persistence.

To provide a structural variance decomposition between discount-rate and cash-flow news, I make a short-run assumption that partially identifies the structural shocks, amounting to two statements. One, the structural dividend-price-ratio shock is simply the present-value approximation error ensuring (6) holds with equality: therefore, it is a shock that has no effect on any of the other variables. Two, by virtue of the return definition (e.g. equation 4) all the cash-flow shocks must affect the return, so I define the structural return shock as one having no effect on the cash-flow variables.[40] Having ordered the variables as $\tilde{\boldsymbol{x}}_t$, we can use the Cholesky decomposition of the reduced-form residuals' covariance matrix to identify the structural shocks as in Sims (1980). Then we calculate the infinite-step-ahead forecast error variance decomposition and report the part of long-run dividend-price-ratio variation attributed to the structural return shock versus the structural cash-flow shocks (see the appendix for more detail).

5.3 Variance decompositions

Table VII reports the dividend-price ratio's long-run variance decomposition between discount-rate and cash-flow news, with Panel A using the sample of all common stocks. In the benchmark results, the LRC approach estimates 65% while the SVAR approach estimates 40%. Averaging those together, I get the paper's topline summary that expected discount rates and expected cash flows contribute equally to dividend-price-ratio volatility. Panel B shows that the main conclusion is robust for either nonfinancials or financials, separately, where once again there is evidence of greater cash-flow impact amongst financials. The last row

[40]Here is another way to think about this assumption. Suppose we are told that cash being paid in dividends or via buybacks will be higher than we previously thought: by the definition of a return, we know that the return will be higher next period. Now, instead suppose we are told the return next period will be higher than we thought: we do not necessarily know that dividends or buybacks will be higher. Hence, the structural return shock is the force not attributable to any surprise to cash flows, whereas cash-flow shocks necessarily affect returns.

Table VII
Long-run variance decomposition

Notes – Reports the dividend-price ratio's long-run variance decomposition, separating between discount-rate and cash-flow news, for different specifications. Columns labeled "LRC" use the long-run coefficient approach from two-step identity estimates of the restricted VAR as described in Section 5.1. Columns labeled "SVAR" use the structural VAR approach as described in Section 5.2. Rows with *sid* use stock-sale issuance. The sample period is 1940–2022 unless otherwise stated.

	Discount Rate		Cash Flow	
	LRC	SVAR	LRC	SVAR
Panel A: Common				
Benchmark	0.65	0.40	0.35	0.60
Nonoverlapping Obs.	0.68	0.17	0.32	0.83
Real	0.55	0.56	0.45	0.44
Real, Nonoverlapping Obs.	0.37	0.27	0.63	0.73
Risk-free-rate sum	0.64	0.38	0.36	0.62
Risk-free-rate sum, Nonoverlapping Obs.	0.69	0.16	0.31	0.84
1926–2022	0.30	0.30	0.70	0.70
1926–2022, Nonoverlapping Obs.	0.32	0.16	0.68	0.84
1970–2022	0.64	0.20	0.36	0.80
1970–2022, *sid*	0.65	0.29	0.35	0.71
1970–2022, Nonoverlapping Obs.	0.67	0.09	0.33	0.91
1970–2022, Nonoverlapping Obs., *sid*	0.69	0.30	0.31	0.70
Panel B: Firm subsamples				
Nonfinancial	0.61	0.46	0.39	0.54
Financial	0.45	0.11	0.55	0.89
Foreign firm, 1962–2023	0.39	0.49	0.61	0.51

shows that foreign-incorporated firms also support the importance of cash-flow news.[41]

It is noteworthy to connect these dividend-price results to a related decomposition. The return is of obvious interest to investors and doesn't depend on payout or fundamental valuation issues.[42] Campbell (1991) prominently advocates the decomposition of the unexpected return into discount-rate and cash-flow news, and many others (e.g. Campbell and Vuolteenaho, 2004; Campbell et al., 2018) have followed suit. Estimates in Campbell et al. (2018) imply that cash-flow news is responsible for about 24% of excess-return variance,

[41] I begin the foreign-firm sample in 1962 to avoid non-positive annual aggregate cash flow values prior to March 1962. The two-step-identity forecast estimates are reported in the appendix—returns and all cash-flow variables are significantly forecasted by the foreign-firm dividend-price ratio, which has much less persistence.

[42] I thank John Campbell for this point.

which is a non-negligible amount.[43] Using the reduced-form VAR estimates,[44] I find that cash-flow news is responsible for 27-28% of unexpected return variation in the benchmark sample, 52-59% in the financial sample, and 48-60% in the foreign-firm sample. Therefore, unexpected-return variance-decomposition results are more aligned to my cash-flow results than to results that cash-flow news insignificantly contributes to dividend-price variation.

Going back to Panel A and scanning across different specifications, one could say a value of 50% is moderately generous to discount rates' contribution. Across all specifications and approaches, the discount-rate contribution goes only as high as 69%—but the cash-flow contribution goes as high as 91% even in the post-Great-Depression period. We see that adjusting for inflation, using nonoverlapping observations, using risk-free summing of the cash flows, using stock-sale issuance, and changing the sample period don't significantly change the qualitative conclusion.[45] The long-run variance decomposition evidence strongly suggests that expected cash flows are *about equally important* as discount rates in driving the aggregate dividend-price ratio.

6 Conclusion

The aggregate dividend-price ratio fluctuates due to varying expectations of future discount rates or cash flows. This paper argues that dividends, buybacks, and issuance all represent cash flows between the firm and household sectors—and all, by definition, drive the dividend-price ratio. Across a variety of periods, firm samples, and empirical specifications we see that future buybacks and issuance are robustly predictable. Decomposing dividend-price-ratio variation in-sample, we find that about half comes from discount-rate expectations and half comes from cash-flow expectations; out of sample, only the cash-flow expectations are significant. Therefore, aggregate stock prices respond to cash-flow news.

Looking only to dividends, we find little predictability—perhaps this is unsurprising. Dividends smoothly vary and reportedly are related to expected long-term earnings.[46] Consider macroeconomic news about the next few years ahead of us. If we get news of good economic times ahead but we don't know how long it will last, that firms will be making more

[43]I use Campbell et al. (2018) table 2 and impose zero correlation between the excess-return shock and volatility news (because the correlation there is statistically insignificant).

[44]Concise details are in the appendix.

[45]I also checked if large-in-magnitude but statistically-insignificant coefficients drive this, and find they don't. If I zero-out any coefficient insignificant at the 10% level and then calculate the long-run variance decomposition, I get very similar results.

[46]See Brav et al. (2005) for survey evidence from corporate CFOs.

profits than we had previously anticipated but no one knows for how long—through what channels does any of this cash flow to the household? One might guess it would largely flow through buybacks that managers view as more flexible than dividends. Hence, the news driving prices today would show up in future net repurchases, not dividends. Essentially, this paper's results support such a narrative.

Consistent with the above intuition, using dividends to scale prices is a good way to impart stationarity—since we impart a minimum of backward-looking cash-flow variation to the forward-looking price whose fluctuations we ultimately care about—in addition to linking to seminal work like Campbell and Shiller (1988). Future research could combine corporate-payout theory and other data into more refined measures. Related, I found a surprising predictive link between future aggregate issuance and the dividend-price ratio, and that corporate accounting data from Compustat changes the macroeconomic story—future work should investigate further.

References

Aharoni, G., B. Grundy, and Q. Zeng (2013): "Stock returns and the Miller Modigliani valuation formula: Revisiting the Fama French analysis," *Journal of Financial Economics*, 110, 347–357.

Allen, F. and R. Michaely (2003): "Payout policy," in *Handbook of the Economics of Finance*, ed. by G. Constantinides, M. Harris, and R. M. Stulz, Elsevier, vol. 1, Part 1, chap. 07, 337–429, 1 ed.

Bansal, R., R. F. Dittmar, and C. T. Lundblad (2005): "Consumption, Dividends, and the Cross Section of Equity Returns," *Journal of Finance*, 60, 1639–1672.

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

Bessembinder, H. (2018): "Do stocks outperform Treasury bills?" *Journal of Financial Economics*, 129, 440–457.

Binsbergen, J. H. V. and R. S. J. Koijen (2010): "Predictive Regressions: A Present-Value Approach," *Journal of Finance*, 65, 1439–1471.

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

Brav, A., J. R. Graham, C. R. Harvey, and R. Michaely (2005): "Payout policy in the 21st century," *Journal of Financial Economics*, 77, 483–527.

Brealey, R. A., S. C. Myers, and F. Allen (2003): *Principles of Corporate Finance*, McGraw-Hill, 7th ed.

Brogaard, J., H. Nguyen, T. J. Putnins, and Y. Zhang (2022): "Are stock prices driven more by cash flow or discount rate news?" Tech. rep., U of Utah, U of Technology Sydney, Stockholm School of Economics in Riga.

Caballero, R. J. and A. Simsek (2020): "A Risk-Centric Model of Demand Recessions and Speculation," *Quarterly Journal of Economics*, 135, 1493–1566.

Campbell, J. and J. Cochrane (1999): "Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

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

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

Campbell, J. Y., S. Giglio, C. Polk, and R. Turley (2018): "An intertemporal CAPM with stochastic volatility," *Journal of Financial Economics*, 128, 207–233.

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

Campbell, J. Y. and S. Thompson (2008): "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *Review of Financial Studies*, 21, 1509–1531.

Campbell, J. Y. and T. Vuolteenaho (2004): "Bad Beta, Good Beta," *The American Economic Review*, 94, 1249–1275.

Center for Research in Security Prices, L. (2021): "Data Descriptions Guide: CRSP US Stock & US Index Databases," Tech. rep.

Chen, L., Z. Da, and X. Zhao (2013): "What Drives Stock Price Movements?" *Review of Financial Studies*, 26, 841–876.

Chen, L. and X. Zhao (2009): "Return Decomposition," *Review of Financial Studies*, 22, 5213–5249.

Clark, T. and M. McCracken (2005): "Evaluating Direct Multistep Forecasts," *Econometric Reviews*, 24, 369–404.

Clark, T. E. and M. W. McCracken (2001): "Test of equal forecast accuracy and encompassing for nest models," *Journal of Econometrics*, 105, 85–110.

Cochrane, J. H. (2005): *Asset Pricing*, Princeton University Press, 2 ed.

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

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

De La O, R. and S. Myers (2021): "Subjective Cash Flow and Discount Rate Expectations," *Journal of Finance*, 76, 1339–1387.

Dechow, P. M., A. P. Hutton, and R. G. Sloan (1996): "Economic Consequences of Accounting for Stock-Based Compensation," *Journal of Accounting Research*, 34, 1–20.

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

Dou, W. W., Y. Ji, and W. Wu (2021): "Competition, profitability, and discount rates," *Journal of Financial Economics*, 140, 582–620.

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

Eisfeldt, A. L., A. Falato, and M. Xiaolan (2022): "Human Capitalists," Tech. rep.

Fama, E. F. and K. R. French (2001): "Disappearing dividends: chaning firm characteristics or lower propensity to pay?" *Journal of Financial Economics*, 60, 3–43.

Golez, B. (2014): "Expected Returns and Dividend Growth Rates Implied by Derivative Markets," *Review of Financial Studies*, 27, 790–822.

Grullon, G. and R. Michaely (2002): "Dividends, Share Repurchases, and the Substitution Hypothesis," *Journal of Finance*, 57, 1649–1684.

Grullon, G., B. Paye, S. Underwood, and J. P. Weston (2011): *Journal of Financial and Quantitative Analysis*, 46, 1–24.

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.

Hodrick, R. J. (1992): "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies*, 5, 357–386.

Kelly, B. T. and S. Pruitt (2013): "Market Expectations in the Cross Section of Present Values," *Journal of Finance*, 68, 1721–1756.

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

Koudijs, P. and H.-J. Voth (2016): "Leverage and Beliefs: Personal Experience and Risk-Taking in Margin Lending," *American Economic Review*, 106, 3367–3400.

Lamont, O. (1998): "Earnings and Expected Returns," *Journal of Finance*, 53, 1563–1587.

Larrain, B. and M. Yogo (2008): "Does firm value move too much to be justified by subsequent changes in cash flow?" *Journal of Financial Economics*, 87, 200–226.

Lettau, M. and S. Ludvigson (2005): "Expected returns and expected dividend growth," *Journal of Financial Economics*, 76, 583–626.

Lettau, M. and S. Van Nieuwerburgh (2007): "Reconciling the Return Predictability Evidence: The Review of Financial Studies: Reconciling the Return Predictability Evidence ," *Review of Financial Studies*, 21, 1607–1652.

Ljungqvist, L. and T. J. Sargent (2018): *Recursive Macroeconomic Theory, Fourth Edition*, MIT Press Books, The MIT Press.

Newey, W. and K. West (1987): "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–08.

Pettenuzzo, D., R. Sabbatucci, and A. Timmermann (2020): "Cash Flow News and Stock Price Dynamics," *Journal of Finance*, 75, 2221–2270.

Pettenuzzo, D., A. Timmermann, and R. Sabbatucci (2022): "Firm Value and Payout Suspensions During Financial Market Distress," Tech. rep., Brandeis University, UCSD, and Stockhold School of Economics.

Sabbatucci, R. (2022): "Are Dividends and Stock Returns Predictable? New Evidence Using M&A Cash Flows," Tech. rep., Stockholm School of Economics.

Sims, C. A. (1980): "Macroeconomics and Reality," *Econometrica*, 48, 1–48.

Stephens, C. P. and M. S. Weisbach (1998): "Actual Share Reacquisitions in Open-Market Repurchase Programs," *Journal of Finance*, 53, 313–333.

Welch, I. and A. Goyal (2007): "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 21, 1455–1508.

A Appendix

A.1 Aggregate per-share aggregate dividend-price ratio

Following Cochrane (2008) derive an aggregate per-share aggregate price-dividend ratio as $\frac{1+\texttt{vwretd}_{t+1}}{1+\texttt{vwretx}_{t+1}} - 1$. From (3)

$$\begin{aligned} 1 + \texttt{vwretd}_{t+1} &= \frac{\sum_n S_{n,t} P_{n,t} \left(\frac{P_{n,t+1} + D_{n,t+1}}{P_{n,t}} \right)}{\sum_n S_{n,t} P_{n,t}} \\ &= \frac{\sum_n S_{n,t} P_{n,t+1} + S_{n,t} D_{n,t+1}}{\sum_n S_{n,t} P_{n,t}} \end{aligned}$$

and similarly

$$1 + \texttt{vwretx}_{t+1} = \frac{\sum_n S_{n,t} P_{n,t+1}}{\sum_n S_{n,t} P_{n,t}}.$$

Then

$$\begin{aligned} \frac{1 + \texttt{vwretd}_{t+1}}{1 + \texttt{vwretx}_{t+1}} - 1 &= \frac{\sum_n S_{n,t} P_{n,t+1} + S_{n,t} D_{n,t+1}}{\sum_n S_{n,t} P_{n,t+1}} - 1 \\ &= \frac{\sum_n S_{n,t} P_{n,t+1}}{\sum_n S_{n,t} P_{n,t+1}} + \frac{\sum_n S_{n,t} D_{n,t+1}}{\sum_n S_{n,t} P_{n,t+1}} - 1 \\ &= \frac{\sum_n S_{n,t} D_{n,t+1}}{\sum_n S_{n,t} P_{n,t+1}}. \end{aligned}$$

Note that the actual aggregate dividends paid are divided by $\sum_n S_{n,t} P_{n,t+1}$: the latter is the price per share at time $t+1$ times the number of shares outstanding at time t. If any firm's number of outstanding shares varies, this is in general not an observed price of the aggregate portfolio. Given that Campbell and Shiller (1988) start with the "price of a stock or stock portfolio", one might be surprised.

A.2 Data

I first drop duplicate observations for (`date`,`PERMNO`) pairs, retaining the first one. These exist for complex corporate actions, such as the break-up of AT&T in February 1984. For the variables I require every observation thereafter is identical (variables like `DIVAMT` and `DISTCD` are those that vary), so choosing the first is without loss of generality. I set to NaN any `RET` observation that is equal to 'C', 'B', -66, -77, -88, or -99. I verify that the resulting `RET` and `RETX` are NaN for the same observations. I convert any negative `PRC` observation to positive—the negative sign denotes a bid-ask average price, but CRSP uses such prices to calculate RET, so I use it too. Any PRC observation equal to 0 I set to NaN. I work with (date,PERMNO)-observations where $\texttt{PRC}_{n,t}$ and $\texttt{SHROUT}_{n,t}$ are nonmissing, of which there are 4,884,020. To get this observation count, I do forward-fill *gaps* of missing PRC and SHROUT values (an example is Berkshire-Hathaway within its first year of existence). I do this so I can use, since I need one-month-lagged values to calculate the variables of interest, the first nonmissing PRC and SHROUT observations after a gap. This contributes 111,275 observations: but note that, by definition, all the return and cash-flow values for those filled-in months are zero, and so these do not affect at the aggregate variables during those months. What rows with missing PRC and SHROUT values *are* dropped are those appearing at the end or beginning (the first observation for Berkshire-Hathaway is an example) of a PERMNO's history.

Net repurchases and monthly CRSP Technically, the market capitalization needed to calculate the net repurchases exists only in the Daily Stock File but not necessarily in the Monthly CRSP file. This is because CRSP ascribes a net repurchase to occur on a certain day of the month: therefore its value can be calculated from market capitalizations on that day and on the day before, exactly as (2) showed and other papers have calculated, but those market capitalizations are not necessarily visible in the monthly data. Therefore net repurchases from monthly data are technically a little different than what one extracts from the daily data, which ostensibly is the most precise. The reason is that RETX reflects the value of the distributions, each reinvested in the security until the end of the month (see Center for Research in Security Prices, 2021, page 101).

Start with the definition of the gross ex-dividend return and assume that $\texttt{CFACSHR}_t = \texttt{CFACSHR}_{t+1}$

$$1 + \texttt{RETX}_{n,t+1} \equiv \frac{\texttt{PRC}_{n,t+1}}{\texttt{PRC}_{n,t}} = \frac{\texttt{SHROUT}_{n,t+1}\texttt{PRC}_{n,t+1}}{\texttt{SHROUT}_{n,t}\texttt{PRC}_{n,t}} + \frac{(\texttt{SHROUT}_{n,t} - \texttt{SHROUT}_{n,t+1})\texttt{PRC}_{n,t+1}}{\texttt{SHROUT}_{n,t}\texttt{PRC}_{n,t}},$$
$$1 + \texttt{RETX}_{n,t+1} = \frac{MKTCAP_{n,t+1}}{MKTCAP_{n,t}} + q_{n,t+1}. \tag{A.1}$$

I have broken the ex-dividend gross return into pieces: the gross "return" in market capitalization, and the net-repurchase return $q_{n,t+1}$. This $q_{n,t+1}$ is the part of $\texttt{RETX}_{n,t+1}$ implied by the amount of capitalized net repurchases occurring during month $t+1$. That is, repurchase events happen on a particular day of the month that is not necessarily the last: this return reflects the value of the distributions, each reinvested in the security until the end of the month . Rearrange (A.1) and the data therefore tell us $q_{n,t+1}$ which we can use to calculate net repurchases as

$$NETREP_{n,t+1} = MKTCAP_{n,t} q_{n,t+1}. \tag{A.2}$$

Because the monthly data do not reveal it, this means I net out buyback and issuance cash flows *within* a month to arrive at one monthly value. In looking at the daily CRSP data, I have found tens of thousands of stock-months where this occurs, aggregating up to well over half a trillion dollars for buybacks and issuance. I leave analysis of these facts for future research to explore.

Stock-sale issuance In Section 4.5 I use stock-sale issuance, sid_t, which is constructed as follows. From WRDS I go to Compustat - Capital IQ>Compustat>North American>Fundamentals Quarterly, downloading `SSTKY` (Sale of Common and Preferred Stock) for all firms since 1961. Using `DATACQTR` (Calendar Data Year and Quarter) I aggregate up and assign the value to the last month of the quarter and linearly interpolate over the remaining months. On months where both values are nonzero, I project the CRSP issuance on `SSTKY`. I use the projected value instead of CRSP's aggregate issuance to thereafter construct sid_t as I describe the construction of id_t in the main text.

A.3 SVAR error decomposition

Write the structural representations of the the unrestricted systems using $\tilde{\boldsymbol{x}}_t$ and $\boldsymbol{v}_t \equiv (w_{d,t}, w_{b,t}, w_{i,t}, u_{r,t}, u_{\delta,t})'$ as

$$\boldsymbol{A}\tilde{\boldsymbol{x}}_t = \boldsymbol{b}_0 + \boldsymbol{B}\tilde{\boldsymbol{x}}_{t-12} + \boldsymbol{v}_t$$

where $\boldsymbol{v}_t$ is mean zero with an identity covariance matrix. The two identifying assumptions imply that the contemporaneous impact matrix looks like

$$\boldsymbol{A} = \begin{bmatrix} * & * & * & 0 & 0 \\ * & * & * & 0 & 0 \\ * & * & * & 0 & 0 \\ * & * & * & * & 0 \\ * & * & * & * & * \end{bmatrix}$$

where $*$ denotes an unknown element.[47] These assumptions are sufficient to identify $u_{\delta,t+1}$ and $u_{r,t+1}$ from each other and the space of cash-flow shocks named w: hence, the u shocks are identified while the w shocks are only partially identified (i.e. their space is identified). But I will not need to separately identify the w, so for my purposes this is sufficient. I can take any rotation of the w shocks as my representation of cash-flow shocks; a simple choice is to take $\boldsymbol{A}$ as lower triangular. This chooses a particular rotation of the cash-flow shocks. However, all I need is for the space of cash-flow shocks to be separated from the space of return and dividend-price-ratio structural shocks, so any rotation of them, including the convenient Cholesky-implied one, delivers identical results.

Therefore the reduced-form residuals' covariance matrices, which are estimable, are $\boldsymbol{A}^{-1}\boldsymbol{A}^{-1\prime}$, and therefore the lower-triangular Cholesky factor of those covariance matrices estimate $\boldsymbol{A}^{-1}$ for the system. Therefore, letting $\boldsymbol{\Phi}$ denote the reduced-form VAR slope estimates (i.e. the matrx reported in Table VI), the long-run effects of the structural shocks can be found $(\boldsymbol{I}-\boldsymbol{\Phi})^{-1}\boldsymbol{A}^{-1}$. Let the last row excluding the last column (the column for δ) of these matrices be $\boldsymbol{c}$: denote the last element (pertaining to r) as c_r, and the remaining subvectors (pertaining to w shocks) as $\boldsymbol{c}_{dbi}$. Then the discount-rate variance is given by $c_r^2/(\boldsymbol{c}'\boldsymbol{c})$, while the cash-flow variance is given by $(\boldsymbol{c}_{dbi}'\boldsymbol{c}_{dbi})/(\boldsymbol{c}'\boldsymbol{c})$.

A.4 Unexpected return decomposition

Write the reduced-form VAR $\tilde{\boldsymbol{x}}_{t+1} = \boldsymbol{a} + \boldsymbol{\Gamma}\tilde{\boldsymbol{x}}_t + \boldsymbol{u}_{t+1}$. Let $\boldsymbol{e}_z$ be a Euclidean basis vector with 1 located where z is located in $\tilde{\boldsymbol{x}}$. Define $\boldsymbol{e}_{CF} \equiv \boldsymbol{e}_d + \rho_b \boldsymbol{e}_b - \rho_i \boldsymbol{e}_i$. Then recursively applying present-value (6) and applying the operator $(\mathbb{E}_{t+1} - \mathbb{E}_t)$ to both sides, we get

$$
\begin{aligned}
0 &= (\mathbb{E}_{t+1} - \mathbb{E}_t) \sum_{j=0} \infty \rho_\delta^j (\boldsymbol{e}_r - \boldsymbol{e}_{CF})' \tilde{\boldsymbol{x}}_{t+j+1} \\
&= \sum_{j=0}^{\infty} \rho_\delta^j \boldsymbol{\Gamma}^j (\boldsymbol{e}_r - \boldsymbol{e}_{CF})' \boldsymbol{u}_{t+1} \\
\boldsymbol{e}_r' \boldsymbol{u}_{t+1} &= \boldsymbol{e}_{CF}' \sum_{j=0}^{\infty} \rho_\delta^j \boldsymbol{\Gamma}^j \boldsymbol{u}_{t+1} - \boldsymbol{e}_r' \sum_{j=1}^{\infty} \rho_\delta^j \boldsymbol{\Gamma}^j \boldsymbol{u}_{t+1}.
\end{aligned}
$$

[47] I am normalizing the structural shocks to have unit volatility, to ease algebra later on. This is WLOG as we could instead define the structural shocks' covariance to be diagonal with non-identical diagonal elements, in which case there would be 1s in the $\boldsymbol{A}$ I've assumed: but in the end this would be a choice of normalization that has no effect on the variance decomposition I'm after.

The left-hand side of the last line is the unexpected return realized at time $t+1$, and it is composed of two news terms on the right-hand side as noted by Campbell (1991)—the first *cash-flow news* and the second *discount-rate news*.

The approach taken by Campbell (1991), Campbell and Vuolteenaho (2004), and many others is to start with the discount-rate news calculated as

$$N_{DR,DR} = \boldsymbol{e}_r' \rho_\delta \boldsymbol{\Gamma} (\boldsymbol{I} - \rho_\delta \boldsymbol{\Gamma})^{-1} \boldsymbol{u}_{t+1}.$$

This is the discount-rate news calculated by focusing first on the discount-rate news, so I subscript it DR, DR. The corresponding cash-flow news is then

$$N_{CF,DR} = \boldsymbol{e}_r' (\boldsymbol{I} + \rho_\delta \boldsymbol{\Gamma} (\boldsymbol{I} - \rho_\delta \boldsymbol{\Gamma})^{-1}) \boldsymbol{u}_{t+1}.$$

An alternative approach is to start with the cash-flow news calculated as

$$N_{CF,CF} = \boldsymbol{e}_{CF}' (\boldsymbol{I} - \rho_\delta \boldsymbol{\Gamma})^{-1} \boldsymbol{u}_{t+1}$$

and then use this to calculate the corresponding discount-rate news

$$N_{DR,CF} = \left[\boldsymbol{e}_{CF}' (\boldsymbol{I} - \rho_\delta \boldsymbol{\Gamma})^{-1} - \boldsymbol{e}_r' \right] \boldsymbol{u}_{t+1}.$$

In the main text I report cash-flow-news contributions as the range between what is estimated by these two approaches to calculating cash-flow (and discount-rate) news.

A.5 Further results

Table A.1
Dividend-price-ratio forecasts, break-adjusted

Notes – Using the full sample, break adjusting the predictor $\tilde{\delta}_{t-12}$ at 1955 and 1995 as in Lettau and Van Nieuwerburgh (2007), estimated via OLS. Newey-West t-stats using 24 lags in parentheses.

	$\tilde{\delta}_{t-12}$	$R^2(\%)$
δ_t	0.723	17.1
	(4.91)	
r_t	0.156	5.4
	(2.16)	
Δd_t	-0.175	18.0
	(-2.69)	
bd_t	-1.093	3.5
	(-2.07)	
id_t	-1.552	18.3
	(-3.76)	

Table A.1 estimates, for each target variable x_t,

$$x_t - \bar{x} = \phi_x \tilde{\delta}_{t-12} + error$$

where $\tilde{\delta}_t = \delta_t - \bar{\delta}_t$ and $\bar{\delta}_t$ is equal to the mean dividend-price ratio: during 1926–1954 if t is before January 1955; during 1955–1994 if t is after December 1954 and before January 1995; during 1995–2023 if t is after December 1994.

Table A.2
Dividend-price-ratio forecasts, risk-free-rate summing

Notes – All common stocks, 1940–2023, computing yearly cash flows using risk-free-rate summing, for the two-step-identity estimator. Other than these differences, further details are as in Table II.

	δ_{t-12}	cons	$R^2/p(\%)$
δ_t	0.945	−0.209	86.4
	(28.68)	(−1.81)	
r_t	0.087	0.413	7.8
	(2.76)	(3.91)	
Δd_t	0.006	0.095	−0.4
	(0.31)	(1.38)	
bd_t	−2.595	−10.848	60.2
	(−12.91)	(−13.81)	
id_t	−1.296	−4.760	60.1
	(−7.19)	(−7.23)	

Table A.3
Dividend-price-ratio forecasts, foreign-firm

Notes – Results separately for foreign (SHRCD = 12) stocks, for the two-step-identity estimator on the sample period 1962–2023 to avoid nonpositive annual aggregate cash flow values prior to March 1962. Other than these differences, further details are as in Table II.

	δ_{t-12}	cons	$R^2/p(\%)$
δ_t	0.647	−1.316	46.1
	(10.34)	(−5.70)	
r_t	0.173	0.729	7.6
	(2.61)	(3.04)	
Δd_t	−0.216	−0.702	14.1
	(−3.99)	(−3.53)	
bd_t	−2.136	−9.541	13.9
	(−4.24)	(−4.90)	
id_t	−1.977	−7.493	19.4
	(−4.62)	(−4.65)	

Table A.4
Dividend-price-ratio forecasts, non-ADR

Notes – Results separately for non-ADR stocks, for the two-step-identity estimator. Further details are as in Table II.

	δ_{t-12}	cons	$R^2/p(\%)$
δ_t	0.924	−0.280	83.9
	(25.54)	(−2.22)	
r_t	0.109	0.485	9.6
	(3.17)	(4.21)	
Δd_t	−0.003	0.070	−0.2
	(−0.13)	(0.93)	
bd_t	−2.870	−11.778	57.3
	(−12.21)	(−13.56)	
id_t	−1.667	−5.959	68.0
	(−11.11)	(−10.55)	
J	4.693		3.0
W	18.130		0.0

Table A.5
Correlation of predicted variables

Notes – Correlation of predicted variables, 1970–2023

	r_t	Δd_t	bd_t	id_t
r_t	1.00			
Δ_t	−0.00	1.00		
bd_t	0.02	−0.17	1.00	
id_t	0.09	−0.28	0.68	1.00
sid_t	−0.10	−0.14	−0.61	