The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXV, NO. 4 • AUGUST 2020

Cash Flow News and Stock Price Dynamics

DAVIDE PETTENUZZO, RICCARDO SABBATUCCI, and ALLAN TIMMERMANN*

ABSTRACT

We develop a new approach to modeling dynamics in cash flows extracted from daily firm-level dividend announcements. We decompose daily cash flow news into a persistent component, jumps, and temporary shocks. Empirically, we find that the persistent cash flow component is a highly significant predictor of future growth in dividends and consumption. Using a log-linearized present value model, we show that news about the persistent dividend growth component predicts stock returns consistent with asset pricing constraints implied by this model. News about the daily dividend growth process also helps explain concurrent return volatility and the probability of jumps in stock returns.

ON MOST DAYS, A MULTITUDE of firms announces cash flow news, but the number of firms, as well as the industries they belong to, can vary greatly over time. Such variation gives rise to a highly irregular cash flow news process, which complicates investors' attempt to infer the underlying growth rate of cash flows for individual firms, industries, and the economy as a whole. This is important because cash flow news is a key determinant of investors' forecasts of future cash flows and of their risks, and hence ultimately of how stocks are priced.[1]

*Davide Pettenuzzo is at International Business School, Brandeis University. Riccardo Sabbatucci is at Department of Finance, Stockholm School of Economics. Allan Timmermann is at Rady School of Management, UC San Diego. We are grateful to the editor, Stefan Nagel, an Associate Editor, and two anonymous referees for many valuable suggestions on our paper. We also received helpful comments from Jules van Binsbergen, Chris Polk, Federico Bandi, and participants at the Utah Winter Finance Conference 2019, CEPR Gerzensee 2018, BI-SHoF 2018, and conferences in Toronto (Western University, April 2018), Toulouse (Toulouse School of Economics, May 2018), and University of Chicago (May 2019). We are grateful to INQUIRE Europe for financial support. We have read *The Journal of Finance* disclosure policy and have no conflicts of interest to disclose.

Correspondence: Allan Timmermann, University of California, San Diego, Rady School of Management, 4S146 Otterson Hall, MC 0553, 9500 Gilman Drive, La Jolla, CA 92093-0553; e-mail: atimmerm@ucsd.edu

[1] Patton and Verardo (2012) develop a rational learning model to explain the observed patterns in firms' betas around earnings announcements. Their model contains unobserved firm-specific and common earnings innovation terms, and investors' extraction of these components is modeled as a Kalman filtering problem. Savor and Wilson (2016) develop a learning model in which investors decompose cash flow news into firm-specific and market-wide components. Positive covariances between the cash flow processes of individual firms and the broader market imply that bad (good) news on individual firms' cash flows result in decreased (increased) forecasts of aggregate cash

DOI: 10.1111/jofi.12901

While information extracted from firms' cash flow announcements is critical to understanding investors' cash flow expectations and, in turn, movements in stock prices, relatively few studies analyze the predictability of cash flows, in most cases focusing instead on quarterly or annual changes in aggregate dividends or earnings.[2] However, data aggregated in this manner may conceal rich dynamic patterns in cash flows recorded at a higher frequency that reduces our ability to study important questions such as how fast can cash flow growth respond to changes in the underlying state of the economy.

Several challenges complicate attempts to measure daily cash flow dynamics. First, the cash flows of most firms have a pronounced seasonal component related to weather patterns and holiday sales. Second, the number of firms announcing cash flow news on any given day can fluctuate between as little as 0 to more than 100 firms. Third, the particular date on which a firm announces dividends can vary widely from year to year, requiring that care be taken in constructing daily proxies that account for firm-specific effects. Fourth, there is considerable heterogeneity across individual firms' cash flow processes. The combined effect of these various factors is that daily cash flow news tends to be very lumpy.

To address these challenges, we develop a new approach that allows us to measure and model dynamics in high-frequency (daily) cash flows. Specifically, we account for firm-level heterogeneity and seasonality effects by taking a bottom-up approach that starts from changes in individual firms' dividends on a given day relative to their value over the same quarter during the previous year. In contrast to conventional smoothed estimates, only data on those firms that announce dividend news on a given day are used to update the dividend growth estimate, which ensure that our measure picks up changes in the cash flow process in a timely fashion.[3] Moreover, by computing a dollar-weighted growth estimate, we account for variation in the size of the firms that pay dividends on any given day.[4] Our dividend growth measure uses dividend announcements by publicly traded firms as opposed to dividend payments, which form the basis for

flows. In turn, this cash flow learning channel implies that the stock returns of announcing firms and of the aggregate market are positively correlated, justifying an announcement risk premium for exposure to individual firms' cash flows. These models do not allow for jumps in cash flows, although, in practice, this is an important feature of earnings and dividend data.

[2] Cochrane (2008) finds little evidence of dividend growth predictability, while van Binsbergen and Koijen (2010), Kelly and Pruitt (2013), and Jagannathan and Liu (2019) find that growth in dividends is predictable.

[3] To illustrate the loss of information from the common practice of aggregating cash flow news over the most recent 12-month period and updating this on, say, a monthly basis, suppose that firms' announcement dates are uniformly distributed across calendar dates. Each month when the cash flow estimate is updated, the same weight is assigned to firms announcing cash flows close to the cutoff date and firms whose announcement occurred almost one year previously. This weighting automatically makes the resulting growth estimate stale and also introduces serial correlation in the estimate.

[4] Cash flow news is often announced after the regular trading sessions in the stock market have closed and thus aggregating across firms that announced cash flows within a 24-hour interval—as opposed to modeling, say, hourly cash flow news—seems appropriate.

the Center for Research in Security Prices (CRSP) measure of dividend growth conventionally used in the finance literature. This distinction is important because dividend announcements precede dividend payments by several weeks, which gives our dividend measure a significant timing advantage and more closely aligns our measure with movements in market prices following dividend news.[5]

To account for lumpiness in daily values of year-over-year changes in firm-matched cash flows, we decompose cash flow news into a slowly evolving component, which picks up time variation (predictability) in the mean of the cash flow process, a transitory component, with volatility that changes over time, and large jumps, the probability of which depends on the number of firms that announce cash flow news on a given day. All three components turn out to be important in capturing predictability in the dividend growth process and evolution in the uncertainty surrounding growth in cash flows.

An important test of our approach is whether it can be used to generate more accurate forecasts of cash flows than existing methods. Empirically, we find that our estimate of the persistent dividend growth component is a strong predictor of future dividend growth. Moreover, the predictive power of our approach compares favorably to alternative predictors of dividend growth proposed by van Binsbergen and Koijen (2010) and Kelly and Pruitt (2013). We also find that our measure of the persistent dividend growth component is a positive and significant predictor of future GDP growth and aggregate consumption growth. In sharp contrast, "raw" dividend growth, as well as the individual jump or transitory shock components, are very noisy and turn out not to have any predictive power with respect to future dividend growth. Our results suggest that firms closely monitor the state of the economy and adjust their dividend policies in anticipation of changes in slow-moving economic indicators such as GDP and consumption growth.

A key limitation of empirical tests of asset pricing models is that while high-frequency data are available on movements in individual and aggregate stock prices, individual firms' cash flows are observed at much lower frequencies—typically quarterly. The absence of high-frequency cash flow data reduces researchers' ability to estimate and test asset pricing models that rely on the joint dynamics of stock prices, expected returns, and cash flow growth expectations.

In the second part of our paper, we address these challenges by using our new daily cash flow estimates to study predictability of stock returns in the context of the log-linearized present value model developed by Campbell and Shiller (1988a). The present value model offers several advantages. As shown by Cochrane (2008) and van Binsbergen and Koijen (2010), the present value model implies a set of cross-equation restrictions that tie the coefficients on the variables in the prediction model for dividend growth to the coefficients on the corresponding variables in the return prediction model. In particular, any variable that helps predict future dividend growth should also have the ability

[5] Announced dividends precede actual dividend payments by approximately 42 days, on average.

to forecast stock returns, after controlling for predictor variables such as the log dividend-price ratio. We find strong empirical evidence that variation in the persistent dividend growth component is a strong and significant predictor of daily stock returns and that the cross-equation restrictions implied by the present value model are supported by the data. Moreover, we show that our new dividend growth measure can be used to predict stock returns several days into the future. At the monthly horizon, our persistent dividend growth measure provides more accurate out-of-sample forecasts of stock market returns than all but one of the 14 predictor variables considered by Goyal and Welch (2008).

We also study whether the different components extracted from the dividend process help explain concurrent (same-day) dynamics in stock returns. We find strong evidence that positive shocks to the persistent dividend growth component are associated with higher same-day mean stock returns, lower stock market volatility, and a reduced likelihood of outliers (jumps) in stock returns, again in a manner that is consistent with the cross-equation restrictions implied by the present value relation. Greater uncertainty about dividend growth translates into higher stock market volatility as well as a larger chance of a jump (outlier) in stock returns, suggesting that daily dividend news is an important determinant of dynamics in aggregate stock market returns and is not just limited to affecting the mean of stock returns.

Our paper is related to a literature that estimates the effect of new information on stock prices on days with news releases using event-study methodology (see, e.g., Cutler, Poterba, and Summers (1989), McQueen and Roley (1993)). A limitation of this approach is that news stories are qualitative (e.g., "world events"), heterogeneous across news categories (e.g., news on housing starts versus monetary policy shocks), and, sometimes, rare (e.g., monetary policy shocks). In addition, the effect of such news may be state-dependent. For example, good news about the economy can be bad news for stock prices in a high-growth state (McQueen and Roley (1993)) and may depend on whether the economy is operating above or below trend. Macroeconomic news can also be difficult to decipher because it comprises a bundle of information about cash flows and expected returns (Boyd, Hu, and Jagannathan (2005)).

The methodology developed in our paper is also related to papers in the asset pricing literature such as Chib, Nardari, and Shephard (2002) and Eraker, Johannes, and Polson (2003) that estimate models of stock return dynamics with stochastic volatility and jumps. There are three key differences between our approach and these papers. First, to the best of our knowledge, no existing study attempts to model the high-frequency dynamics in dividends using such methods, let alone estimate and test a model as general as ours. Second, we test asset pricing implications implied by the present value model using a much higher data frequency (daily) than previously employed. Third, our paper models the transmission from daily cash flow news—in the form of mean, volatility, or jump components—to contemporaneous dynamics in stock market returns, including variation in the volatility and jump probability of returns.

The paper is organized as follows. Section I introduces our data and explains how we construct a daily cash flow index from dividend announcement data. Section II describes our econometric modeling approach and reports empirical model estimates. Section III analyzes dividend growth predictability. Section IV develops the present value framework to explore the implications of predictability of dividend growth for return predictability, while Section V analyzes the effect of news about the dividend growth process on concurrent stock returns. Section VI concludes. Additional technical details, extensions, and robustness tests are reported in the Internet Appendix.[6]

I. Data

We start our analysis by explaining how we construct our daily dividend growth series and describing the data sources that we use. Our analysis of daily cash flows focuses on growth in dividends, which as Kelly and Pruitt (2013) point out, is the focus of a large literature on asset pricing.[7] Because earnings can be negative, defining earnings growth poses challenges that are quite different from those that arise when studying dividends.

The greatest effect of dividend news on asset prices is likely to come through its information content, so we focus on dividends as initially *announced* as opposed to actual dividend payments. However, for robustness in the Internet Appendix, we analyze daily dividends from the perspective of the payment date, which allows us to compare the information effect to the direct cash flow effect from dividend payments.

A. Sample Construction

Our sample includes all ordinary cash dividends declared by firms with common stocks (share codes 10 and 11) listed on NYSE, NASDAQ, or AMEX from 1926 to 2016.[8] We do not exclude special dividends from our measure. For example, Costco (permno 87055), which usually pays dividends around 30 cents per quarter, declared two "special dividends" during our sample period: $7 per share on November 28, 2012, and $5 per share on January 30, 2015. Those two dividends, with CRSP codes 1232 and 1272, respectively, are included in our dividend measure. What we do *not* include in our measure is nonordinary dividends such as M&A cash flows, buybacks, and new issues. We also require that firms have valid stock prices and a number of shares outstanding when dividends are announced. Furthermore, we ensure that there are no duplicate observations in the data set and that each firm pays only one

[6] The Internet Appendix is available in the online version of this article on *The Journal of Finance* website.

[7] See, for example, Campbell and Shiller (1988a), Cochrane (1992), Lettau and Ludvigson (2005), Koijen and van Nieuwerburgh (2011), and Maio and Santa-Clara (2015).

[8] Ordinary cash dividends have CRSP distribution codes below 2000.

dividend at any point in time.[9] Overall, our sample consists of 503,591 declared dividends.[10]

Corporate dividends have a strong firm-specific component and often display pronounced seasonal variation. Our analysis therefore computes dividend growth by comparing same-firm, same-(fiscal) quarter, year-over-year changes in cash flows. To this end, let D_t^i be the total dividends declared by firm i on day t, calculated as the dividend per share times the number of shares outstanding. Moreover, let I_t^i be an indicator variable that equals 1 if company i announces quarterly dividends on day t, and 0 otherwise, and let $\tilde{t}$ be the same-quarter, prior-year dividend announcement date for firm i.[11] Consider, for example, a company that declared dividends on May 17, 2014, while it declared the corresponding quarter's prior-year dividends on May 9, 2013. In this case, t is May 17, 2014 and $\tilde{t}$ is May 9, 2013.

Aggregating across firms, the total dollar value of dividends paid out on day t is $\sum_{i=1}^{N_t} I_t^i D_t^i$, where N_t is the number of (publicly traded) firms in existence on day t. Similarly, the total value of dividends paid out by the *same* set of firms for the *same* fiscal quarter during the prior year is given by $\sum_{i=1}^{N_t} I_{\tilde{t}}^i D_{\tilde{t}}^i$. Taking the ratio of these two numbers, we obtain a measure of the aggregate, year-over-year (gross) growth in dividends:[12]

$$G_t = \frac{\sum_{i=1}^{N_t} I_t^i D_t^i}{\sum_{i=1}^{N_t} I_{\tilde{t}}^i D_{\tilde{t}}^i}. \tag{1}$$

Note that the number of firms used in this calculation—as well as their identity—changes on a daily basis and from year to year as firms shift their dividend announcement dates. Only firms for which $I_t^i = I_{\tilde{t}}^i = 1$ are included in this calculation, which ensures that the *same*firms are used in both the numerator and denominator of the ratio. Equation (1) uses the dollar amount paid in dividends by individual firms, implicitly applying value weights since large firms tend to have larger dividend payouts.[13]

[9] There are instances in CRSP in which a company declares or pays multiple dividends on the same day, using different distribution codes but still classified as ordinary dividends. We aggregate such dividends to convert them into a single dividend. As an example, on November 23, 1983, PPL Corporation (permno 22517) declared two ordinary dividends of 39 and 21 cents.

[10] Following a recent update, CRSP no longer provides the dividend declaration date prior to 1962 and data until 1964 appear to be incomplete. However, we also have an older version of the database in which the declared dividend dates start in 1926. As a result, we have 101,476 pre-1964 observations and 402,115 post-1964 observations.

[11] We note that $\tilde{t}$ depends on firm i, and the same-quarter, next-year date t, so a more precise notation would be $\tilde{t}(i, t)$.

[12] Only 7% of individual firms' year-over-year dividend growth observations in our sample are constant, suggesting that firms often change their dividends (even marginally) every year.

[13] An alternative approach that accounts for heterogeneity in firm size more explicitly is to first define individual firms' cash flow growth as

$$\tilde{G}_t^i = \begin{cases} \dfrac{D_t^i}{D_{\tilde{t}}^i} & \text{if } I_t^i = 1 \\ 0 & \textit{otherwise} \end{cases}.$$

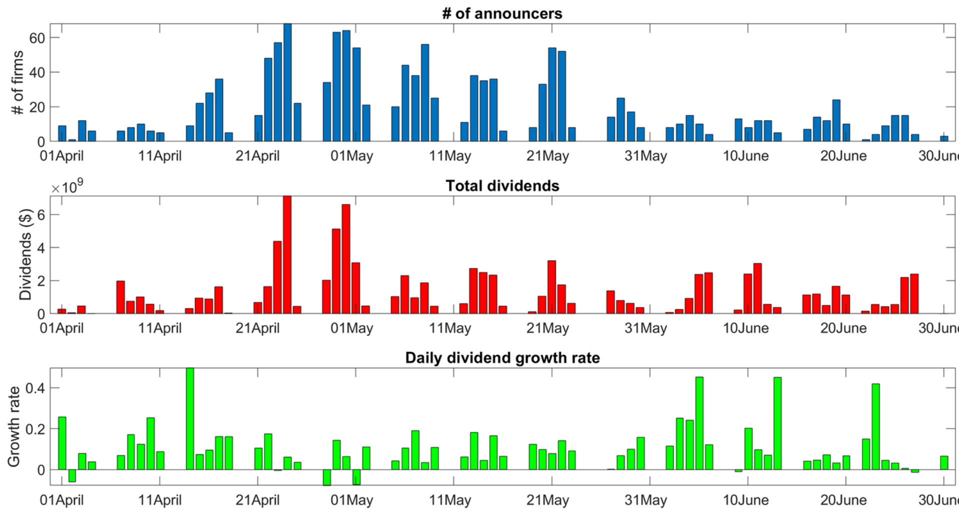

Figure 1. Distribution of dividend announcements within a quarter. This figure plots time series of dividend announcements for 2014 Q2. For each day within this quarter, the top panel shows the number of firms announcing dividends. The middle panel shows the overall nominal amount of dividends announced by those firms (in billion dollars), while the bottom panel shows the daily (net) dividend growth rate. (Color figure can be viewed at wileyonlinelibrary.com)

As a first illustration of our data, Figure 1 plots the number of firms, as well as the dollar dividend and the (net) growth rate from equation (1), during a single quarter (2014Q2). The top panel shows substantial intraquarter variation in the number of firms announcing dividends. During this particular quarter, the maximum number of firms announcing dividends on any one day was 68 (on April 24), while the minimum number was 1 (on June 22). On several days, more than 50 firms announced dividends.

The middle panel in Figure 1 shows the total value of dividends declared on individual days. As can be seen, this value depends on both the number and the size of firms announcing dividends, as large firms tend to announce bigger dividends.[14] The bottom panel in Figure 1 shows the daily net dividend growth

In a second step, we can use individual firms' market capitalization to aggregate cash flow growth rates across the firms that pay dividends on day t,

$$\tilde{G}_t = \sum_{i=1}^{N_t} I_t^i \omega_t^i \tilde{G}_t^i, \text{ where}$$

$$\omega_t^i = \frac{MktCap_t^i \times I_t^i}{\sum_{i=1}^{N_t} MktCap_t^i \times I_t^i}$$

is the weight on company i in the daily year-over-year value-weighted dividend growth calculation. Results based on this alternative measure are very similar to those based on the measure in equation (1).

[14] The largest amount of dividends declared during 2014Q2, \$7.12bn, happened on April 24, while only \$3.6m of dividends were announced on June 30.

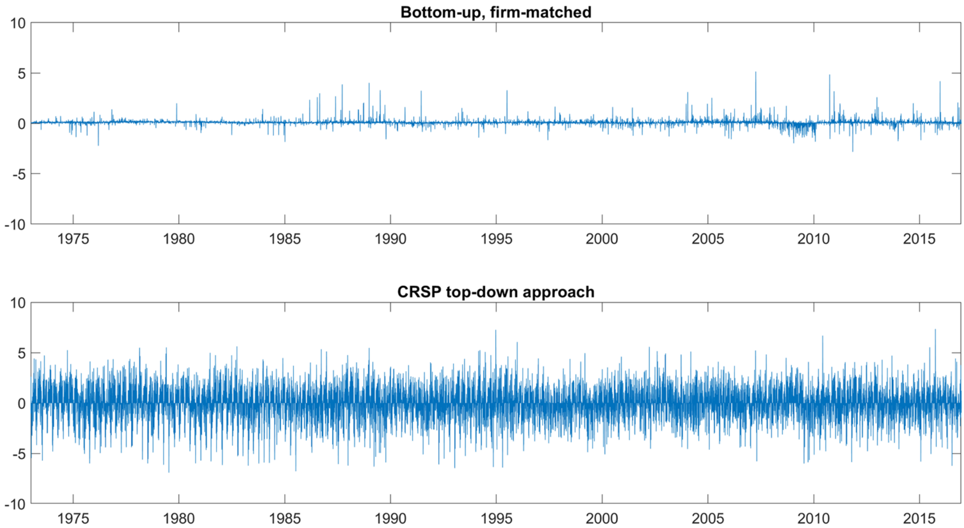

Figure 2. Comparison between our daily "bottom-up" dividend growth series versus a daily "top-down" dividend growth measure extracted from CRSP. The top panel plots the log of the daily dividend growth series G_t defined in equation (1). The bottom panel plots the CRSP-extracted daily dividend growth, calculated as dividends (paid out) on an given day divided by dividends distributed on the same day one year earlier. Both plots use daily data over the period 1973 to 2016. (Color figure can be viewed at wileyonlinelibrary.com)

during the quarter. Peaks in this measure need not coincide with days on which most firms announce dividends (top panel) or days on which the overall amount of announced dividends peaked (middle panel). This is because the dividend growth rate depends on dividends announced by the same group of firms during the prior year as reflected in the denominator of equation (1). For example, the gross dividend growth rate on June 22 (1.15) is generated by a single firm announcing dividends on that day—the firm announced \$155m in dividends in 2014Q2 and \$135m in 2013Q2. Thus, variation in daily dividend growth rates reflects both heterogeneity across firms' dividend behavior and changes in the number of firms announcing dividends on a given day.

B. Features of Daily Dividend Growth

Our data span the period 1927–2016, but the first part of the sample is dominated by the Great Depression. For robustness, we therefore split the sample into halves and study both the full sample and the recent subsample (i.e., 1973 to 2016). The top panel of Figure 2 plots $\Delta d_t = \ln(G_t)$ from 1973 to 2016. On days with no dividend announcements, we set the series to zero.[15] The daily dividend growth series is very spiky and dominated by days with

[15] This happens on 170 days, or 1.4% of the full sample.

unusually large positive or negative dividend growth. There is also evidence of a sustained decline in dividends during the financial crisis.

The features displayed by our daily series of year-over-year growth in dividends in Figure 2 can be summarized as follows. First, the daily dividend growth series is very lumpy, reflecting variation in individual firms' cash flow growth as well as variation in the composition of firms that, on any given day, announce dividends. Second, daily dividend news is driven by a persistent component that is particularly pronounced during the financial crisis of 2008 to 2009. And third, the volatility of daily cash flow news changes over time, between unusually calm periods on the one hand and more volatile periods on the other.

C. Comparison with the Standard Dividend Growth Measure

The conventional alternative to our bottom-up approach is to extract dividends top-down using market returns with and without dividends as published by the CRSP. Three limitations render this conventional approach unattractive.

First, the daily CRSP index accounts for dividends *distributed* on a particular day but does not indicate *when* those dividends were announced. This distinction is crucial as firms typically announce dividends several days prior to the payment date and it is the announcement effect that we would expect to be important for movements in stock prices.

Second, the set of firms announcing or paying dividends on any given day is generally different from the set of firms announcing dividends on the same day one year earlier. As a consequence, year-over-year estimates of dividend growth from daily values of the CRSP index are difficult to interpret as they do not control for firm fixed effects and this may be distorted due to changes in the composition of the set of firms paying dividends on a given calendar day, which would confound dividend growth information with composition shifts.[16]

Third, the CRSP index contains assets such as exchange-traded funds (ETFs) and mutual funds (Sabbatucci (2017)), whose dividends follow a pattern different from that of individual firms.

These limitations turn out to make a crucial difference in the calculation of daily dividend growth. To see this, the bottom panel of Figure 2 plots the daily dividend growth rate constructed using the conventional top-down CRSP approach over the post 1973 sample.[17] While the daily dividend growth series based on our bottom-up methodology (top panel) is affected by occasional jumps, it clearly contains a persistent component linked to the state of the economy. In contrast, the daily dividend growth series constructed from the CRSP return indexes is very noisy throughout the sample considered.

[16] Changes in the composition of dividend-paying firms have a far smaller effect on dividend growth measured over longer intervals such as a quarter or a year.

[17] To replicate our methodology as closely as possible, the daily dividend measure from CRSP shown here computes the growth rate as the change in aggregate dividends paid on a given day relative to aggregate dividends paid on the same day (or whichever day is closest) during the previous year.

II. Econometric Model

To match the features of the daily dividend growth data noted above, in this section, we develop an econometric model that incorporates multiple components to capture different parts of the dividend dynamics. To do so, we first account for lumpiness by allowing for a jump component in daily cash flow growth. Since the lumpiness introduced by firm-level heterogeneity in dividend payments is more likely to be diversified away if a large number of firms announce dividends on the same day, we allow the jump intensity to depend on the number of firms announcing dividends. Next, we incorporate a persistent component in the mean growth equation. Finally, we account for time-varying volatility by modeling ordinary shocks to daily dividend growth as a stochastic volatility process.

A. A Components Model for Daily Dividend Growth

Our econometric model decomposes the daily dividend growth process, Δd_{t+1}, into three parts, namely (i) a persistent term, μ_{dt+1}, which captures a smoothly evolving mean component, (ii) a jump component, $\xi_{dt+1}J_{dt+1}$, where $J_{dt+1} \in \{0, 1\}$ is a jump indicator that equals 1 in the case of a jump in dividends and 0 otherwise, while ξ_{dt+1} measures the magnitude of the jump, and (iii) a temporary cash flow shock, ε_{dt+1}, whose volatility is allowed to be time-varying. Adding these terms, we have

$$\Delta d_{t+1} = \mu_{dt+1} + \xi_{dt+1}J_{dt+1} + \varepsilon_{dt+1}. \tag{2}$$

We next motivate this decomposition and introduce our assumptions on the components.

As we reduce the time interval over which dividends are measured, the data become more scarce, with sizeable day-over-day variation in both the number and the type of firms that announce dividends. To minimize the effect of such composition shifts, our main empirical analysis focuses on an aggregate, market-wide measure of dividend growth, which offers the broadest coverage of dividend events and minimizes the number of days with very few firms announcing dividends.

To understand how daily changes in the composition of firms announcing dividends can affect our dividend growth measure, suppose the data-generating process for aggregate dividend growth, Δd_{At+1}, takes the form

$$\Delta d_{At+1} = \mu_{dt+1} + \sigma_A \varepsilon_{At+1}, \quad \varepsilon_{At+1} \sim \mathcal{N}(0, 1), \tag{3}$$

where μ_{dt+1} is the conditional mean of the aggregate dividend growth process at time $t+1$ and σ_A is the volatility of the daily shocks, ε_{At+1}.

We capture any persistence that may be present in the mean of the dividend growth process by assuming that μ_{dt+1} follows a mean-reverting first-order autoregressive process,

$$\mu_{dt+1} = \mu_d + \phi_\mu(\mu_{dt} - \mu_d) + \sigma_\mu \varepsilon_{\mu t+1}, \quad \varepsilon_{\mu t+1} \sim \mathcal{N}(0, 1), \tag{4}$$

where $|\phi_\mu| < 1$ and $\varepsilon_{\mu t+1}$ is assumed to be uncorrelated with ε_{At+1}. In the special case in which $\phi_\mu = 0$, d_{At+1} follows a random walk process whose changes, Δd_{At+1}, are unpredictable.

In practice, we do not observe the aggregate dividend process at a daily frequency. Rather, on any given day, we observe the dividends announced by a subset of firms. Shifts in the composition of firms announcing dividends on a given day can introduce measurement issues and must be carefully accounted for when modeling the dynamics in the observed dividend process. To see how changes in the composition of dividend-announcing firms affect our dividend growth measure, let Δd_{it+1} be the observed value of firm i's year-over-year dividend growth rate on day $t+1$. We decompose this growth rate into a systematic component that is driven by firm i's cash flow beta (β_i) times the growth in aggregate cash flows, $\beta_i \Delta d_{At+1}$, and an idiosyncratic component, ε_{it+1}, which is assumed to be uncorrelated across firms,

$$\Delta d_{it+1} = \beta_i \Delta d_{At+1} + \sigma_i \varepsilon_{it+1}, \quad \varepsilon_{it+1} \sim \mathcal{N}(0,1). \tag{5}$$

Let ω_{it} be the weight on the dividends announced by firm i on day t, measured relative to the total amount of dividends paid out by all firms on this day. Thus, $\omega_{it} = 0$ if firm i does not announce dividends on day t, and $\sum_{i=1}^{N_{dt}} \omega_{it} = 1$ each day, where N_{dt} is the number of firms announcing dividends on day t. Our measure of (announced) dividend growth on day $t+1$, Δd_{t+1}, aggregated across all dividend-announcing firms, can then be computed as

$$\begin{aligned}
\Delta d_{t+1} &= \sum_{i=1}^{N_{dt+1}} \omega_{it+1} \Delta d_{it+1} = \sum_{i=1}^{N_{dt+1}} \omega_{it+1} \left[\beta_i \Delta d_{At+1} + \sigma_i \varepsilon_{it+1}\right] \\
&= \left[\sum_{i=1}^{N_{dt+1}} \omega_{it+1} \beta_i\right] \Delta d_{At+1} + \left[\sum_{i=1}^{N_{dt+1}} \omega_{it+1} \sigma_i \varepsilon_{it+1}\right] \\
&\equiv \beta_{t+1} \Delta d_{At+1} + \sigma_{dt+1} \varepsilon_{t+1}, \quad \varepsilon_{t+1} \sim \mathcal{N}(0,1),
\end{aligned} \tag{6}$$

where $\beta_{t+1} \equiv \sum_{i=1}^{N_{dt+1}} \omega_{it+1} \beta_i$ is the weighted average of the cash flow betas of firms announcing dividends on day $t+1$, and $\sigma_{dt+1} = [\sum_{i=1}^{N_{dt+1}} \sigma_i^2 \omega_{it+1}^2]^{1/2}$. If all stocks have identical cash flow betas, then β_{t+1} in (6) will be constant over time and equal to the average beta, $\bar{\beta}$. Even if there is heterogeneity in firms' cash flow betas, however, provided that the number of firms announcing dividends each day is large, the weights on individual firms (ω_{it+1}) are small, and the subset of firms announcing dividends on a given day is reasonably random over time, we would expect little time variation in β_{t+1}. This makes it easy to extract a daily estimate of the aggregate dividend growth rate from equation (6) by means of filtering methods.

The number of firms announcing dividends can be low on some days and so the announced dividends can be dominated by a few companies or by companies within the same industry. If cash flow betas vary systematically across industries, this can lead to time variation in the cash flow beta of the subset of dividend-announcing firms.

While these composition effects can lead to time variation in the cash flow beta of the observed dividend growth process, Δd_{t+1}, such effects are temporary and so $\beta_{t+1} - \bar{\beta}$ will be mean-reverting toward zero. Using (6), we can therefore write

$$\begin{aligned} \Delta d_{t+1} &= \bar{\beta}\Delta d_{At+1} + (\beta_{t+1} - \bar{\beta})\Delta d_{At+1} + \sigma_{dt+1}\varepsilon_{t+1} \\ &= \bar{\beta}\big[\mu_{dt+1} + \sigma_A \varepsilon_{At+1}\big] + (\beta_{t+1} - \bar{\beta})\big[\mu_{dt+1} + \sigma_A\varepsilon_{At+1}\big] + \sigma_{dt+1}\varepsilon_{t+1} \\ &\propto \mu_{dt+1} + \frac{(\beta_{t+1} - \bar{\beta})}{\bar{\beta}}\mu_{dt+1} + \left[1 + \frac{\beta_{t+1} - \bar{\beta}}{\bar{\beta}}\right]\sigma_A\varepsilon_{At+1} + \frac{\sigma_{dt+1}}{\bar{\beta}}\varepsilon_{t+1}. \end{aligned} \tag{7}$$

Compared to the aggregate dividend process in (3), there are some additional time-varying terms in the observed dividend growth rate: (i) $\frac{(\beta_{t+1} - \bar{\beta})}{\bar{\beta}}\mu_{dt+1}$, which can be quite volatile because of random variation in the first part, $(\beta_{t+1} - \bar{\beta})/\bar{\beta}$, driven by the process determining the number and types of firms announcing dividends on any given day (this term tends to be particularly volatile on days where few firms announce dividends, increasing the probability that $|\beta_{t+1} - \bar{\beta}|/\bar{\beta}$ will be large), (ii) $\frac{\beta_{t+1} - \bar{\beta}}{\bar{\beta}}\sigma_A\varepsilon_{At+1}$; this is the product of a random component $(\beta_{t+1} - \bar{\beta})/\bar{\beta}$ and an uncorrelated temporary shock, ε_{At+1}, and (iii) $(\sigma_{dt+1}/\bar{\beta})\varepsilon_{t+1}$. The last two components can introduce time-varying and persistent volatility in Δd_{t+1} if firms in certain industries tend to cluster their dividend announcement dates around certain dates and if the volatility of the idiosyncratic dividend growth component differs across industries, as we might expect.

Our model in (2) captures the effect of these components in Δd_{t+1} in two ways. First, the jump component, $\xi_{dt+1}J_{dt+1}$, can account for temporary composition shifts that lead to large variation in the daily cash flow beta of our composite dividend growth measure. It can also be used to absorb outliers in company-specific dividend shocks, ε_{it+1}, on days when individual firms dominate dividend announcements. Because the magnitude of this component is likely to be larger on days with few firms announcing dividends, we allow the jump probability to depend on the number of announcers, N_{dt+1},

$$\Pr(J_{dt+1} = 1) = \Phi(\lambda_1 + \lambda_2 N_{dt+1}), \tag{8}$$

where Φ is the cumulative distribution function (CDF) of a standard normal distribution. The magnitude of the jumps is modeled as $\xi_{dt+1} \sim \mathcal{N}(0, \sigma_\xi^2)$.

Second, our model captures time-varying heteroskedasticity in Δd_{t+1} by modeling the variance of the residuals in the dividend growth process as a stochastic volatility process,

$$\varepsilon_{dt+1} \sim \mathcal{N}(0, e^{h_{dt+1}}), \tag{9}$$

where h_{dt+1} is the log-variance of ε_{dt+1}, which is assumed to follow a mean-reverting process

$$h_{dt+1} = \mu_h + \phi_h(h_{dt} - \mu_h) + \sigma_h \varepsilon_{ht+1}, \qquad \varepsilon_{ht+1} \sim \mathcal{N}(0, 1), \tag{10}$$

where ε_{ht+1} is uncorrelated with both ε_{dt+1} and $\varepsilon_{\mu t+1}$. Stochastic volatility in the dividend growth process (7) may arise due to the clustering of days with firms from similar industries (e.g., financial institutions) announcing dividends.

To summarize, our dividend growth model in (2) accounts for a persistent mean-reverting component, time-varying volatility, and jumps. We evaluate the importance of these features by comparing results from the general model in (2) to a simpler (no-jump) model that ignores jump dynamics and stochastic volatility and so takes the form

$$\Delta d_{t+1} = \mu_{dt+1}^{NJ} + \varepsilon_{dt+1}, \qquad \varepsilon_{dt+1} \sim \mathcal{N}(0, \sigma_d^2), \tag{11}$$

where μ_{dt+1}^{NJ} follows the process in (4). This comparison allows us to gauge the importance of incorporating jump dynamics and stochastic volatility for our estimate of μ_{dt+1}.

B. Estimation

We adopt a Bayesian approach that uses Gibbs sampling to estimate the model parameters. Internet Appendix Section I provides details on our estimation procedure, while Internet Appendix Section II documents the convergence properties of our estimation algorithm.

It is worth briefly describing the priors that underlie our model. We choose standard normal-gamma conjugate priors, which simplify the process of drawing from the conditional distributions of the model parameters in the Gibbs samplers. Moreover, we specify independent priors for the parameters of the mean, variance, and jump processes. For almost all of the prior hyperparameters, we use loose and mildly uninformative priors. The main exceptions are the persistence parameters, ϕ_μ and ϕ_h, whose priors we center on 0.98. Further details are provided in Internet Appendix Section I.

C. Empirical Estimates

We next present estimates of the parameters of our econometric model and evaluate empirically the behavior of the three components of the dividend growth process.

Table I presents parameter estimates for our general dividend growth model in equations (2) to (10) for both the recent subsample (1973 to 2016) and the full sample (1927 to 2016). We focus our discussion on the parameter estimates for the post-1973 sample, but note that the estimates for the full sample are similar.

First, consider the parameters determining the mean of the dividend growth process in (4). The long-run mean estimate, $\mu_d = 0.080$, corresponds to an 8% annualized nominal dividend growth rate, which is close to the mean of

Table I

Parameter Estimates for the Dividend Growth Rate Model

This table shows parameter estimates for a model fitted to the daily dividend growth series. The equations for the components model, described further in Section II.A, take the form

$$
\begin{aligned}
\Delta d_{t+1} &= \mu_{dt+1} + \xi_{dt+1} J_{dt+1} + \varepsilon_{dt+1},\\
\mu_{dt+1} &= \mu_d + \phi_\mu(\mu_{dt} - \mu_d) + \sigma_\mu \varepsilon_{\mu t+1},\\
\varepsilon_{dt+1} &\sim \mathcal{N}(0, e^{h_{dt+1}}),\\
h_{dt+1} &= \mu_h + \phi_h(h_{dt} - \mu_h) + \sigma_h \varepsilon_{ht+1},\\
\Pr(J_{dt+1} = 1) &= \Phi(\lambda_1 + \lambda_2 N_{dt+1}),\\
\xi_{dt+1} &\sim \mathcal{N}(0, \sigma_\xi^2),
\end{aligned}
$$

where μ_{dt+1} captures the mean of the smooth component of the underlying dividend process, $J_{dt+1} \in \{0, 1\}$ is a jump indicator that equals 1 in the case of a jump in dividends and 0 otherwise, ξ_{dt+1} measures the jump size, ε_{dt+1} is a temporary cash flow shock, $\varepsilon_{\mu t+1} \sim \mathcal{N}(0, 1)$ is assumed to be uncorrelated at all times with the innovation in the temporary dividend growth component, ε_{dt+1}, $|\phi_\mu| < 1$, h_{dt+1} denotes the log-variance of ε_{dt+1}, $\varepsilon_{ht+1} \sim \mathcal{N}(0, 1)$ is uncorrelated at all times with both ε_{dt+1} and $\varepsilon_{\mu t+1}$, N_{dt+1} denotes the number of firms announcing dividends on day $t+1$, Φ is the CDF of a standard normal distribution, and $\xi_{dt+1} \sim \mathcal{N}(0, \sigma_\xi^2)$ captures the magnitude of the jumps. The columns report the posterior mean, standard deviation, and 90% credible sets for the parameter estimates.

	Parameter Estimates					
	1927 to 2016			1973 to 2016		
	Mean	Std	90% Credible Set	Mean	Std	90% Credible Set
μ_d	0.058	0.016	[0.035, 0.080]	0.080	0.012	[0.062, 0.097]
ϕ_μ	0.999	0.000	[0.999, 0.999]	0.998	0.001	[0.997, 0.999]
σ_μ	0.002	0.000	[0.002, 0.002]	0.002	0.000	[0.002, 0.002]
μ_h	−5.076	0.037	[−5.138, −5.014]	−5.337	0.045	[−5.415, −5.262]
ϕ_h	0.749	0.008	[0.735, 0.762]	0.833	0.008	[0.820, 0.846]
σ_h	1.320	0.031	[1.269, 1.369]	0.752	0.028	[0.706, 0.797]
σ_ξ	2.651	0.040	[2.584, 2.718]	2.761	0.040	[2.695, 2.829]
λ_1	−1.387	0.039	[−1.452, −1.323]	−1.354	0.045	[−1.428, −1.281]
λ_2	−0.054	0.006	[−0.064, −0.045]	−0.024	0.002	[−0.028, −0.021]

the standard dividend growth measure (extracted from CRSP data) of 7.8% over the same sample period—a figure well within the 90% credible set of [0.062, 0.097]. The persistence parameter for the mean-reverting component of the dividend growth process, ϕ_μ, is centered on 0.998, corresponding to a half-life of 350 trading days. While highly persistent, shocks to the mean-reverting component (4) are very small as shown by the estimate $\sigma_\mu = 0.002$. Our model thus identifies a small, highly persistent component in the dividend growth process.

The top and bottom panels in Figure 3 plot the persistent dividend growth component, μ_{dt}, extracted using either the no-jump model (equation (11), top panel) or the general jump model (equation (2), bottom panel). The μ_{dt}^{NJ} series

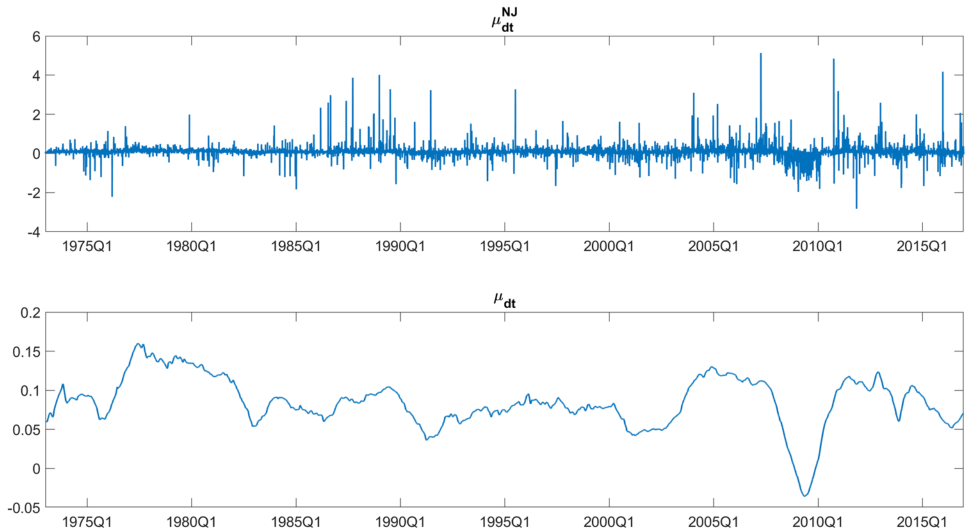

Figure 3. Time series of daily dividend growth and the persistent growth component. The top panel plots the persistent dividend growth component, μ_{dt}^{NJ}, extracted from a model without jumps and stochastic volatility. The bottom panel plots the persistent dividend growth component, μ_{dt}, extracted from the daily dividend series using a model that accounts for jumps and stochastic volatility. All plots use daily data over the period 1973 to 2016. (Color figure can be viewed at wileyonlinelibrary.com)

erroneously assigns large daily spikes in the observed series to the persistent component, μ_{dt}. In contrast, the jump model succeeds in separating temporary spikes (noise) in the daily dividend series from the persistent component, μ_{dt}, which, as a result, is far smoother. Indeed, values of the persistent dividend growth component extracted from the general model fall on a far narrower scale than the unfiltered dividend growth series, ranging from just below 0 to 0.15. The financial crisis of 2008 to 2009 is associated with a notable drop in the persistent dividend growth component, which, for the only time in our sample, turns negative, followed by a notable bounce-back in the second half of 2009 and 2010.

Our empirical analysis uses the dividend growth measure in (1), but one could alternatively use a market (cap-) weighted growth measure. To briefly examine whether our choice of dividend growth measure makes a difference for the estimates of our components model, we construct the daily dividend series as described in footnote 13 and plot the resulting estimate of μ_{dt} along with our original estimate in Figure 4. The two series are virtually indistinguishable with a correlation of 98.11%, which implies that the dividend weighting scheme makes little or no difference to our results.

To empirically quantify the importance to our estimates of μ_{dt} of including days with very few firms announcing dividends—presumably days with a higher likelihood of being dominated by "noise" due to the potential for large composition shifts—we reestimate our model using only observations for

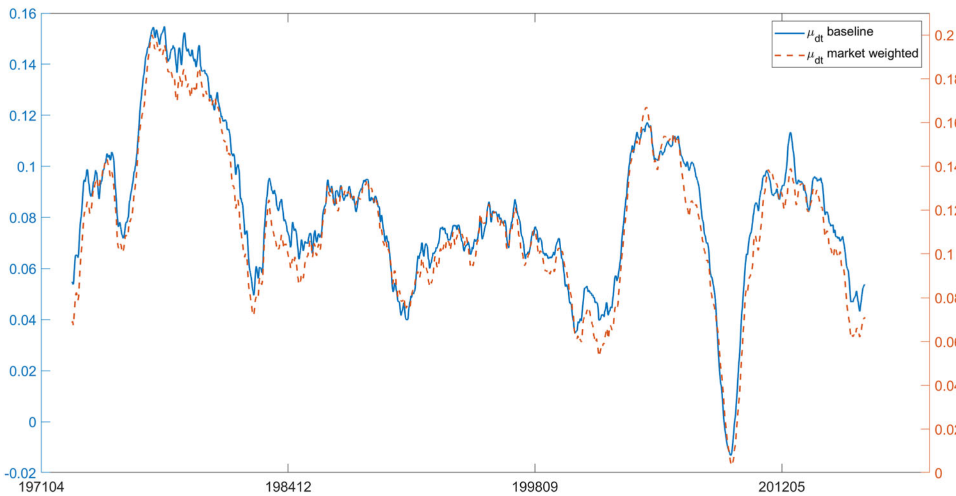

Figure 4. Comparison of alternative approaches for constructing μ_{dt}. This figure plots estimates of the persistent dividend growth measure μ_{dt} extracted from our model using dollar-weighted (baseline) and market-weighted (market weighted) dividend growth data. (Color figure can be viewed at wileyonlinelibrary.com)

dividend growth on days with at least five announcing firms. This requirement leads us to drop around 10% of the original sample. The correlation between our original μ_{dt} series (constructed without requiring at least five announcing firms) and our new estimate of μ_{dt} (constructed with at least five announcing firms) on days when they overlap is 99.64%, so our model seems to succeed in isolating the effect of outliers in the dividend growth data caused by a scarcity of firms announcing dividends. Allowing for jumps in the dividend growth process is key to ensuring such robustness. To see this, we extract μ_{dt} in a similar way, using only days with at least five announcing firms, but now in a model without jumps and stochastic volatility. When we compare this series to an estimate from the same model constructed without the requirement of at least five announcing firms, we find that the correlation is only 72%, suggesting that the no-jump model is far more strongly affected by days with few firms announcing dividends.

Our estimates of equation (10) show that the stochastic volatility process is moderately persistent with an autoregressive parameter, ϕ_h, whose estimated mean is 0.833.

Finally, consider the jump component extracted from our model. The negative estimates of the jump intensity parameters, λ_1 and λ_2, imply that a jump occurs every 60 days on average, with lower jump probabilities on days when a large number of firms announce dividends. On many days in the sample period, the jump probability indicator, J_{dt}, shown in the top panel of Figure 5, is close to 1. On such days, spikes in daily dividend growth are attributed mostly to jumps rather than to clusters with high volatility from the transitory component, ε_{dt}, in equation (2). Jumps can be very large in magnitude, as shown in the

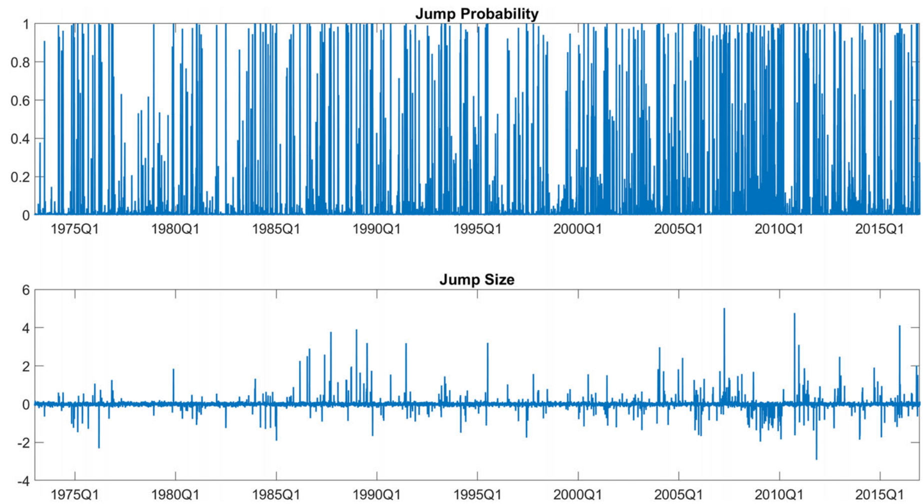

Figure 5. Jumps in the daily dividend growth series. The top panel plots the probability of a jump in the daily dividend growth series, while the bottom panel plots the magnitude of such jumps. Both plots use daily dividend data over the period 1973 to 2016. (Color figure can be viewed at wileyonlinelibrary.com)

bottom panel of Figure 5, which displays the estimated jump size, ξ_{dt}. Indeed, the estimated standard deviation of the jump size, σ_ξ, has a mean of 2.761, which is almost four times larger than the estimated mean of σ_h (0.75). Shocks to daily dividend growth originating from the jump component thus tend to be far larger than the regularly occurring ε_{dt} shocks.

To get a sense of how sensitive the dividend growth jump probability is to the number of firms announcing dividends on a given day, N_{dt}, Figure 6 plots the jump intensities for three values of N_{dt} chosen to match the 25th, median, and 75$^{\text{th}}$ percentiles of the distribution of the daily number of announcing firms. On days with a large number of announcing firms (75$^{\text{th}}$ percentile, or 36 firms on average), the jump intensity distribution is centered on a number a little over 0.005, which corresponds to a jump every 200 days on average. On days with a typical number of announcing firms (median, or 22 firms), the jump intensity is centered around its average value near 0.016, implying a jump every 60 days. Finally, on days with a small number of announcing firms (25$^{\text{th}}$ percentile, or 12 firms), the probability of a jump is centered just below 0.03, which corresponds to a jump in dividend growth every 35 days.

D. Firm Heterogeneity and Composition Effects

Earlier studies find evidence of heterogeneity in the cash flow process across firms with different characteristics. For example, Vuolteenaho (2002) and Cohen, Polk, and Vuolteenaho (2009) link firms' book-to-market ratios to cross-sectional variation in their cash flow betas, finding that value stocks have higher cash flow betas than growth stocks.

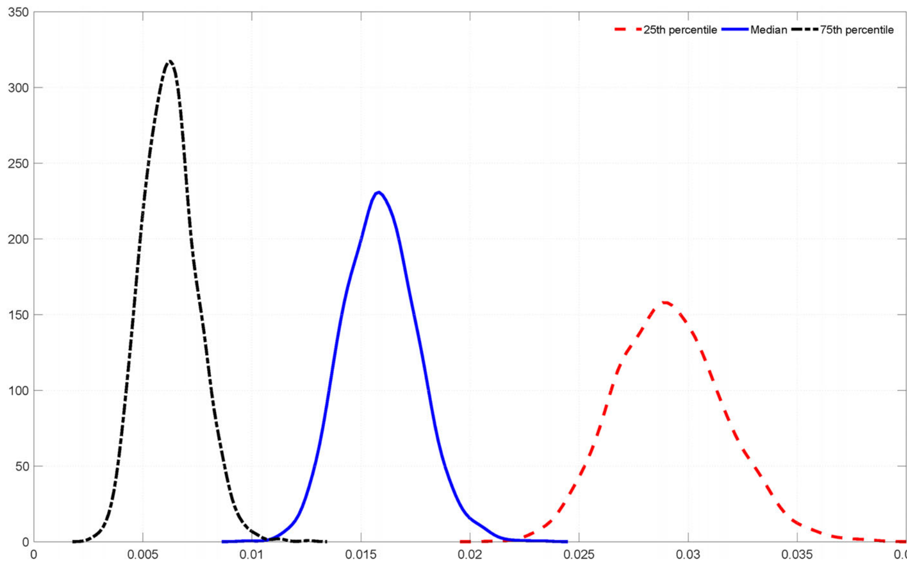

Figure 6. Jump intensities and the number of firms announcing dividend news. This figure shows the sensitivity of the dividend growth jump probability to the number of firms announcing dividends on a given day, N_{dt}, chosen to match the 25^{th}, median, and 75^{th} percentiles of the distribution of the daily number of firms announcing dividends. On days with a large number of announcing firms (dashed curve), the jump intensity distribution is centered around 0.005, corresponding to a jump every 200 days on average. On days with a typical (median) number of announcing firms (solid curve), the jump intensity is centered around 0.016, implying a jump roughly every 60 days. Finally, on days with a small number of announcing firms (dotted curve), the probability of a jump is 0.03, corresponding to a jump every 35 days. (Color figure can be viewed at wileyonlinelibrary.com)

To address concerns about heterogeneity in cash flow betas caused by daily shifts in the composition of the subset of dividend-announcing firms, we construct a number of daily dividend series that only include data on firms in specific sectors or firms with certain book-to-market ratio or size characteristics.[18] We then apply our methodology separately to each of these disaggregate series and extract a persistent dividend growth component. Provided that cash flow betas are more homogeneous within each of these groups, composition effects due to daily variation in the set of dividend-announcing firms matter less.

Figure 7 plots the μ_{dt} estimates obtained from firms sorted into high versus low book-to-market categories (top panel) or firms sorted into small-, medium-, or large-size categories (bottom panel). The figure confirms that there are cross-sectional differences in the behavior of the persistent dividend components (μ_{dt}) extracted from our model. Specifically, the persistent dividend growth component displays much stronger time-series variation for value firms

[18] For example, each month we sort all firms in our sample by their book-to-market (BM) ratios and form two portfolios consisting of firms whose BM ratios are above (“high”) or below (“low”) the median BM ratio. To be included in the sort in a given month, a stock needs to have reported book and market values and must be traded on the NYSE, AMEX, or NASDAQ.

Figure 7. Heterogeneity in μ_{dt} and firm characteristics. The top panel plots estimates of the persistent dividend growth component, μ_{dt}, computed separately for firms with high versus low book-to-market ratios. The bottom panel plots estimates of the persistent dividend growth component, μ_{dt}, computed separately for small, medium, and large firms. All estimations use daily data over the period 1973 to 2016. (Color figure can be viewed at wileyonlinelibrary.com)

than for growth firms (top panel) and for small firms compared with large or medium-sized firms (bottom panel). Comparing the individual μ_{dt} series with our aggregate, market-wide estimate of μ_{dt}, we find correlations of 97.80% for large firms, 88.38% for medium firms, 71.67% for small firms, 90.40% for firms with high book-to-market, and 92.89% for firms with low book-to-market.

In Internet Appendix Section III, we also estimate the persistent dividend growth component (μ_{dt}) for the five Fama-French industries and compute parameter estimates for our decomposition applied to these series. The parameter estimates and time-series plots show that the basic features of the dividend process remain the same across very different industries.

These results demonstrate that while the predictable dividend growth component varies across firms with different characteristics and in different industries, some features are common across firms: estimates of ϕ_μ close to unity (indicating a highly persistent component in dividend growth), similar jump sizes, and negative dependence between the jump probability and the number of firms announcing dividends on a given day.

III. Predictability of Dividend Growth

Predictability of dividend growth features prominently in discussions of asset pricing models. Cochrane (2008) finds little evidence of predictability of U.S. dividends, while van Binsbergen and Koijen (2010), Kelly and Pruitt (2013), and Jagannathan and Liu (2019) argue that dividend growth is, at least to some extent, predictable.[19] The parameter estimates from our dividend model show that the daily dividend growth process contains a small but highly persistent component. In this section, we explore the implications of our results for predictability in dividend growth.

A. Predictive Regressions

Existing studies on dividend growth predictability use time-aggregated dividends measured over horizons longer than our daily interval. To explore whether our estimate of the persistent dividend growth component can predict dividends—and to make our results directly comparable to existing results—we use the conventional top-down approach to construct monthly and annual measures of dividend growth from CRSP data, denoted by Δd_t^{CRSP}.[20]

Next, we estimate a predictive regression of future dividend growth, Δd_{t+1}^{CRSP}, on the persistent dividend component measured at the end of the previous period, μ_{dt}, the log dividend-price ratio extracted from CRSP, dp_t, and current and lagged dividend growth,

$$\Delta d_{t+1}^{CRSP} = \alpha + \beta \mu_{dt} + \gamma dp_t + \sum_{j=1}^{3} \rho_j \Delta d_{t+1-j}^{CRSP} + \varepsilon_{t+1}. \tag{12}$$

[19] Recent literature uses dividend futures to estimate the term structure of dividends. van Binsbergen, Brandt, and Koijen (2012) and van Binsbergen and Koijen (2017) recover prices of dividend strips and show that their expected returns are higher than those on the underlying index.

[20] Most researchers extract aggregate dividends, D_t, from CRSP as the difference between the cum-dividend return (VWRETD), R_t^{cum}, and the ex-dividend return (VWRETX), R_t^{ex}, multiplied by the previous ex-dividend index level, P_{t-1}^{ex}, that is, $D_t = (R_t^{cum} - R_t^{ex}) \times P_{t-1}^{ex}$. Using the resulting aggregate dividend series, the log dividend growth rate can be computed as $\Delta d_t^{CRSP} = \ln(\frac{D_t}{D_{t-1}})$.

We include the log dividend-price ratio in the regression because a variety of studies suggest that it helps predict cash flow growth (e.g., Cochrane (1992), Cochrane (2008), Lettau and van Nieuwerburgh (2008)).

Panel A of Table II shows that the persistent component for dividend growth, μ_{dt}, has strong predictive power over future dividend growth recorded at the quarterly horizon. In the post-1973 sample, the lagged persistent dividend growth component has a t-statistic of 4.6 after accounting for the effect of lagged dividend growth and the lagged dividend-price ratio. Moreover, the R^2 is quite high at 0.26. These findings are robust to the sample period: starting the sample in 1927, the coefficient on μ_{dt} obtains a t-statistic of 2.5, and the R^2 of the predictive regression rises to 0.37. The coefficient on the lagged dividend-price ratio is not significant in any of these regressions, while the first two lags of dividend growth are significant in the quarterly models.

The predictive power of μ_{dt} for future dividend growth is very similar at the annual and quarterly horizons in the post-1973 sample, though is somewhat weaker at the annual frequency using the full sample that begins in 1927. Still, μ_{dt} remains highly statistically significant at the annual horizon for both samples, and this result is again robust to the presence of lagged dividend growth and the dividend-price ratio in the predictive regression.

B. Comparison with Alternative Predictors of Dividend Growth

Previous studies report evidence of predictability in dividend growth. Campbell and Shiller (1988b) develop a vector autoregressive (VAR) model for the joint dynamics in dividend growth, the price-dividend ratio, and the cyclically adjusted price-earnings ratio (CAPE). Lags of all three variables obtain significant coefficients in the dividend growth equation, suggesting significant in-sample predictability of dividend growth at the annual frequency over the period 1871 to 1987. Similarly, Jagannathan and Liu (2019) propose a three-variable VAR that includes dividend growth, the dividend-earnings (payout) ratio, and a latent component that can capture the long-run risk surrounding future dividend growth. Although their analysis focuses on large stocks in the S&P500 index, their findings on predictability of annual dividend growth are consistent with ours. Finally, Donaldson and Kamstra (1996) propose a neural net model to capture nonlinear predictability in discounted future dividend growth. Computing stock prices as the present value of expected discounted future dividends, they can match the boom and crash in stock prices experienced during the Great Depression.

The two studies most closely related to ours are van Binsbergen and Koijen (2010) and Kelly and Pruitt (2013). For example, van Binsbergen and Koijen (2010) use a latent variable Kalman filtering approach to estimate a log-linearized present value model that contains expected returns and expected dividend growth rates for the aggregate stock market and identify a predictable component in dividend growth.

Kelly and Pruitt (2013) assume that individual firms' stock returns and log cash flow growth rates are a linear function of a set of unobserved common

Table II

Dividend Growth Regressions

Panel A reports results from predictive regressions of dividend growth extracted from CRSP data, Δd_{t+1}^{CRSP}, on the persistent component μ_{dt} estimated from our daily dividend growth model and the log dividend price ratio, dp_t, at quarterly and annual frequencies. Panel B compares the predictive power of our persistent dividend growth component to that of two alternative dividend growth variables. The first measure, g_t^{VBK}, is taken from van Binsbergen and Koijen (2010) and uses cash-reinvested dividends, measured annually. The second measure, g_t^{KP}, is taken from Kelly and Pruitt (2013) and uses monthly dividend data. Panel C reports results from forecast encompassing regressions that compare the predictive power of our μ_{dt} measure to the two alternative measures, both in-sample and out-of-sample (available only for Kelly and Pruitt (2013)). Square brackets report t-statistics computed using Newey-West standard errors with three lags. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Unless otherwise indicated, the sample period is 1973 to 2016.

Panel A: $\Delta d_{t+1}^{CRSP} = \alpha + \rho_i \sum_{i=1}^{3} \Delta d_{t+1-i}^{CRSP} + \beta \mu_{dt} + \gamma dp_t^{CRSP} + \varepsilon_{t+1}$

	Quarterly (1)	Annual (2)	Quarterly (1927-) (3)	Annual (1927-) (4)
μ_{dt}	0.36***	2.50***	0.14**	0.96***
	[4.58]	[3.95]	[2.49]	[3.20]
dp_t	−0.01	0.03	−0.00	0.04
	[−1.29]	[0.50]	[−0.73]	[1.11]
Δd_t^{CRSP}	0.16**	−0.64***	0.34***	−0.16
	[2.20]	[−4.58]	[5.58]	[−1.07]
Δd_{t-1}^{CRSP}	0.07	−0.53***	0.22***	−0.12
	[1.39]	[−3.87]	[4.58]	[−0.84]
Δd_{t-2}^{CRSP}	0.03	−0.12	0.01	0.01
	[0.45]	[−0.78]	[0.26]	[0.05]
R^2	26.50%	27.45%	37.42%	6.31%
Observations	169	40	353	86

Panel B: $\Delta d_{t+1}^{i} = \alpha + \beta \mu_{dt} + \gamma g_t^{i} + \varepsilon_{t+1}$

	Annual			Monthly		
	(1)	(2)	(3)	(4)	(5)	(6)
μ_{dt}	1.28***		1.10***	1.39***		1.16***
	[4.04]		[3.97]	[7.23]		[5.12]
g_t^{VBK}		0.90**	0.34			
		[2.16]	[0.91]			
g_t^{KP}					1.45***	0.43
					[6.41]	[1.58]
R^2	25.84%	13.14%	27.27%	31.65%	20.30%	32.67%
Observations	42	42	42	527	527	527

Panel C: $\Delta d_{t+1}^{i} = \beta_1 \mu_{dt} + (1-\beta_1)\gamma g_t^{i} + \varepsilon_{t+1}$

	(3)	(6)
β_1	0.87***	1.08***
p-Value	(0.00)	(0.00)
$1-\beta_1$	0.13	−0.08
p-Value	(0.66)	(0.53)
β_1 (OOS)		1.25***
p-Value (OOS)		(0.00)
$1-\beta_1$ (OOS)		−0.25***
p-Value (OOS)		(0.00)

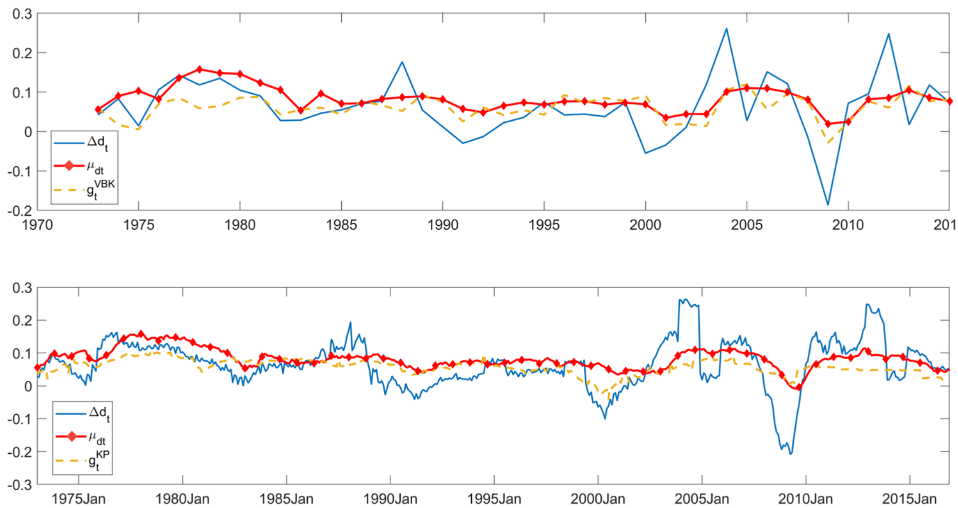

Figure 8. Actual versus predicted dividend growth under alternative modeling approaches. The top panel plots actual dividend growth, Δd_t, against the persistent dividend growth component extracted from our model, μ_{dt}, and the measure proposed by van Binsbergen and Koijen (2010), g_t^{VBK}. The latter assumes cash-reinvested dividends. The bottom panel plots actual dividend growth against our persistent dividend growth component and the measure of Kelly and Pruitt (2013), g_t^{KP}. In both cases, we extend the sample period originally used by the papers after replicating their results. (Color figure can be viewed at wileyonlinelibrary.com)

factors that can be estimated using a three-pass regression (partial least squares) methodology. In turn, cash flow growth can be projected on the common factors to generate a dividend growth forecast. Empirically, Kelly and Pruitt (2013) find strong in-sample evidence of predictability in annual cash flow growth, while their out-of-sample results are somewhat mixed. In the Depression era (1930 to 1940), dividend growth appears to be highly unstable and hard to predict, while out-of-sample predictability is stronger over the 1940 to 2010 period.

We next compare our dividend growth estimates to results based on these two approaches.[21] To this end, the top panel in Figure 8 plots realized values of annual dividend growth against the persistent growth component estimated from our model, μ_{dt} (sampled annually) and the van Binsbergen and Koijen (2010) measure, g_t^{VBK}. The bottom panel plots monthly dividend growth against our persistent dividend growth series, μ_{dt} (sampled monthly) and the Kelly and Pruitt (2013) estimate, g_t^{KP}. While these dividend growth estimates are clearly correlated, there are also some notable differences. Notably, our persistent dividend growth measure shows a sharper decline during the global financial crisis compared with the two alternative estimates.

[21] We are grateful to Seth Pruitt for sharing data and computer code that allowed us to replicate the results in Kelly and Pruitt (2013).

Panel B in Table II reports more formal results from regressions of observed future dividend growth on the growth estimate implied by the three approaches. Because van Binsbergen and Koijen (2010) study cash-reinvested, annual dividend growth, while Kelly and Pruitt (2013) use monthly dividend growth extracted from CRSP, their growth estimates are not directly comparable. We therefore report separate results for the annual and monthly series used in the two studies. To make comparisons between the three approaches easier to interpret and to relate the results to other parts of our paper, we use the post-1973 sample adopted throughout much of our analysis.

All three growth estimates clearly have predictive power for future dividends in the univariate regressions. For example, the growth estimate of van Binsbergen and Koijen (2010) obtains a t-statistic of 2.16 with an R^2 of 13% in the annual sample. In comparison, the t-statistic on our μ_{dt} estimate is 4.04 and the associated R^2 is 25%. Including both μ_{dt} and g_t^{VBK} in the regression, we continue to observe a large t-statistic on μ_{dt} (3.97), while the t-statistic on the estimate of van Binsbergen and Koijen (2010) drops to 0.91 with an R^2 of 27%. These results demonstrate the predictive power of our estimate of the persistent growth component.

In the monthly dividend growth regressions, the growth estimate of Kelly and Pruitt (2013) generates a t-statistic of 6.41 and an R^2 of 20%. In comparison, the t-statistic obtained when we instead use our μ_{dt} measure is 7.23 and the R^2 is 31%. Including both μ_{dt} and g_t^{KP} as predictors in the regression, μ_{dt} obtains a t -statistic of 5.12, while the t-statistic of the Kelly and Pruitt (2013) growth estimate declines to 1.58.

These results show that the persistent component in dividend growth extracted from daily dividend announcements possesses strong predictive power for actual dividend growth at both the monthly and the annual frequencies. Moreover, our μ_{dt} estimate adds substantial predictive power to existing dividend growth estimates.

To formally test and compare the predictive power of the three dividend growth estimates, we run the following pair of forecast encompassing regressions:

$$\begin{aligned}\Delta d_{t+1}^{CRSP} &= \alpha + \beta_1 \mu_{dt} + (1-\beta_1) g_t^{VBK} + \varepsilon_{t+1},\\ \Delta d_{t+1}^{CRSP} &= \alpha + \beta_1 \mu_{dt} + (1-\beta_1) g_t^{KP} + \varepsilon_{t+1}.\end{aligned} \tag{13}$$

The larger is β_1, the greater is the weight on our dividend growth estimate and the smaller is the weight on the competing model estimate. Specifically, a value of $\beta_1 = 1$ suggests that μ_{dt} dominates (encompasses) either g_t^{VBK} (top regression) or g_t^{KP} (bottom regression).[22]

The top row of Table II, Panel C, shows that the estimate of β_1 equals 0.87 in the encompassing regression that includes μ_{dt} and g_t^{VBK}. The persistent dividend growth estimate from our model thus obtains a weight of 87%, while the weight on the van Binsbergen and Koijen (2010) estimate

[22] Note that $g_t = E_t \Delta d_{t+1}$ is the forecast of dividend growth next period.

equals 13%. Moreover, the estimated weight on μ_{dt} is statistically significant at the 1% level. In the second regression, the weight on μ_{dt} is 108%, which is significantly different from zero, while the weight on the Kelly and Pruitt (2013) dividend growth estimate is -8%, which is not significantly different from zero.

We conclude from these regressions that our dividend growth estimate adds significant predictive value to state-of-the-art measures proposed in existing finance literature.

B.1. Out-of-Sample Forecasts

The analysis so far uses data up to 2016 and as a result does not address whether our approach could have been used in real time to generate accurate forecasts of dividend growth. To address this question, we conduct an out-of-sample forecasting experiment that only uses historically available data to estimate our model parameters and generate forecasts. Specifically, using an expanding estimation window and an initial warm-up period from 1973 to 1986, we reestimate our model each week and construct real-time, out-of-sample daily forecasts over the 1986 to 2016 period. We then take the last daily value of our forecast each month and compare this to the monthly out-of-sample forecasts generated using the approach of Kelly and Pruitt (2013). Finally, we evaluate the accuracy of these forecasts using either univariate regressions or the forecast encompassing regressions in (13).

A univariate regression of dividend growth on our out-of-sample, real-time estimate, μ_{dt}, generates an R^2 of 27.76% (t-statistics of 6.70), while the same regression using the out-of-sample forecast of Kelly and Pruitt (2013) produces an R^2 of 2.88% (t-statistics of 2.38).[23] Similarly, in the forecast encompassing regression, the weight on the μ_{dt} series increases from 1.08 to 1.25 with a t-statistic of 12, while the weight on the Kelly and Pruitt (2013) forecast is -0.25. Our dividend growth forecasts thus perform very well even out-of-sample.[24]

To better assess the periods in which our $\Delta\mu_{dt}$ measure produces more accurate forecasts than existing alternatives, we follow Goyal and Welch (2008) and inspect the out-of-sample cumulative sum of squared forecast error differentials (CSSED) from a forecasting model based on $\Delta\mu_{dt}$ versus the Kelly-Pruitt dividend growth measure. Using $\Delta\mu_{dt}$ leads to more accurate forecasts than the Kelly-Pruitt model between the start of the sample (1973) and 1980, while our model underperforms during a brief spell in the late 1990s. During the 16-year period after 2000, our $\Delta\mu_{dt}$ measure consistently produces more

[23] Results ending in 2007, excluding the recent financial crisis, are qualitative similar with an R^2 of 25.47% (t-statistics of 9.01) for our measure and an R^2 of 4.05% (t-statistics of 2.15) for Kelly and Pruitt (2013).

[24] Since our out-of-sample forecasts are generated recursively, our findings are not overly sensitive to the Great Recession that produced a considerable amount of cash flow news, see, for example, Campbell, Giglio, and Polk (2013).

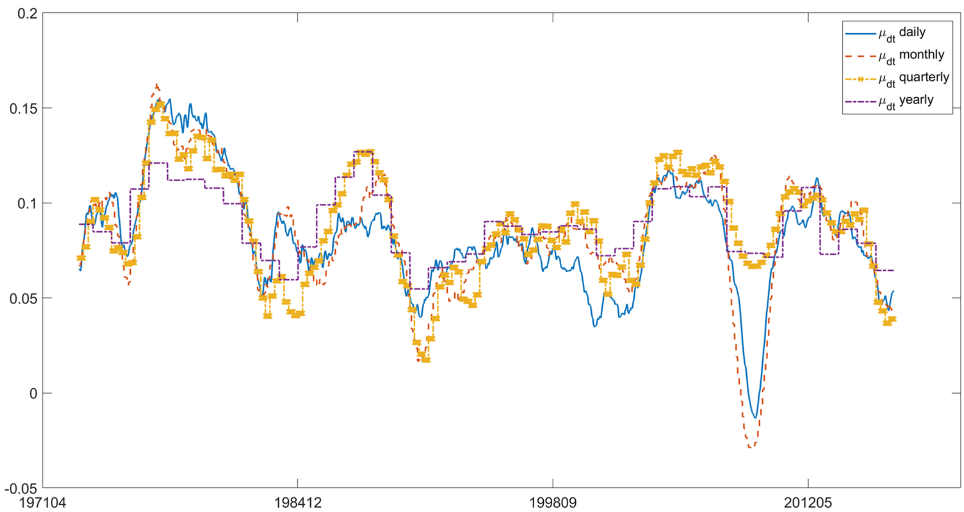

Figure 9. Comparison of estimates of the persistent dividend growth component, μ_{dt}, extracted from data at the daily, monthly, quarterly, and annual frequencies. This figure plots time series of μ_{dt} estimated at the daily, monthly, quarterly, and annual frequencies. (Color figure can be viewed at wileyonlinelibrary.com)

accurate forecasts of dividend growth than the forecasting model based on the Kelly-Pruitt measure.

C. Data Frequency and Dividend Dynamics

Prior work on predictability of dividend growth focuses on long horizons using monthly, quarterly, or annual data, while our results pertain to the daily frequency. To more formally evaluate the role of data frequency in our analysis, we estimate our components model using dividend data at the longer frequencies, extract the corresponding estimates of the persistent growth component, μ_{dt}, and replicate the analysis in Panel A of Table II. Figure 9 shows that the dividend growth estimates extracted from the lower frequency data possess many of the features of the daily estimates. However, they also tend to be very smooth compared to the estimates of μ_{dt} extracted from daily data. This means that the lower frequency estimates miss some of the important troughs and peaks in the dividend growth process. For example, the severity of the decline in dividend growth during the global financial crisis is completely missed by the quarterly and annual estimates, while the monthly estimate fares better in this regard.

Table III shows that our daily μ_{dt} estimate outperforms the monthly, quarterly, and yearly measures in terms of predictive accuracy, generating higher R^2s even when the basis for the comparison is the monthly, quarterly, and annual forecast horizons used to estimate the lower frequency measures.

Table III

Dividend Growth Regressions at Different Frequencies

This table reports results from predictive regressions of the conventional dividend growth measure extracted from CRSP data, Δd_{t+1}^{CRSP}, on the persistent component, μ_{dt}, extracted from daily, quarterly, or annual dividend growth series, and the log dividend-price ratio, dp_t. Columns (1), (4), and (7) (monthly, quarterly, and annual, respectively) regress dividend growth on a daily estimate of μ_{dt}, measured at the end of the corresponding period. Columns (3), (6), and (9) report forecast encompassing regressions that impose the restriction that the coefficients on the daily estimate of μ_{dt} and the monthly, quarterly, or annual estimates of μ_{dt} sum to unity. Square brackets report t-statistics computed using Newey-West standard errors with three lags. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. The sample period for these regressions is 1973 to 2016.

$$\Delta d_{t+1}^{CRSP} = \alpha + \beta \mu_{dt}^{i} + \gamma dp_t^{CRSP} + \sum_{i=1}^{3} \rho_i \Delta d_{t+1-i}^{CRSP} + \varepsilon_{t+1}$$

	Monthly			Quarterly			Annual		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
μ_{daily}	0.17***		1.14***	0.36***		0.73***	2.50***		1.72
	[5.77]		[15.01]	[4.58]		[7.64]	[3.95]		[1.44]
$\mu_{monthly}$		0.14***	−0.14*						
		[6.51]	[−1.90]						
$\mu_{quarterly}$					0.22***	0.27***			
					[3.19]	[2.87]			
μ_{yearly}								2.64**	−0.72
								[2.58]	[−0.60]
dp_t	−0.00**	−0.00	−0.03***	−0.01	0.00	−0.02***	0.03	0.07	0.05
	[−2.49]	[−0.97]	[−9.86]	[−1.29]	[0.20]	[−2.79]	[0.50]	[1.12]	[0.73]
Δd_t^{CRSP}	−0.05	−0.05	−0.58***	0.16**	0.22**	−0.05	−0.64***	−0.52***	−0.52***
	[−0.97]	[−0.94]	[−4.77]	[2.20]	[2.59]	[−0.72]	[−4.58]	[−3.07]	[−2.81]
Δd_{t-1}^{CRSP}	−0.07*	−0.07**	−0.58***	0.07	0.10	−0.10	−0.53***	−0.36***	−0.39*
	[−1.96]	[−2.01]	[−5.69]	[1.39]	[1.65]	[−1.50]	[−3.87]	[−3.01]	[−2.03]
Δd_{t-2}^{CRSP}	0.20***	0.20***	−0.31***	0.03	0.06	−0.13*	−0.12	−0.08	−0.05
	[3.68]	[3.59]	[−3.99]	[0.45]	[1.05]	[−1.84]	[−0.78]	[−0.52]	[−0.26]
R^2	18.29%	17.80%		26.50%	21.76%		27.45%	22.14%	
Observations	525	525	525	169	169	169	40	40	40

Next, we perform forecast encompassing regressions, regressing monthly, quarterly, or annual dividend growth on the daily growth estimate plus one of the monthly, quarterly, or annual growth estimates subject to the constraint that the coefficients ("weights") sum to unity. At all three horizons, we find that our daily dividend growth estimate gains substantially more weight than the corresponding low-frequency dividend growth estimates.

These results show that data frequency is important to our findings and that applying our decomposition on daily data yields stronger prediction results than applying it on data recorded at lower frequencies. The reason is that the daily dividend growth estimate better captures changes in the momentum of dividend growth, which tends to be smoothed out in time-series averages based on longer periods.

The twin effects of using daily data on growth in announced dividends *and* decomposing changes in this series into a smoothly evolving persistent component, temporary jumps, and (small) shocks with time-varying volatility combine to yield better forecasts of dividend growth. So, *both* the higher data frequency *and* our decomposition approach help produce more accurate forecasts of dividend growth.

Past studies suggest that dividends are sticky, as firms are cautious in raising dividends and even more hesitant in reducing them (e.g., Brav et al. (2005)). This behavior might imply that the predictability of dividend growth differs on the upside and the downside. To address this question, we extend the regression in the first column of Table III to include μ_{dt} interacted with top and bottom quintile dummies defined according to whether the prior-month dividend growth rate was in the highest or lowest 20% of historical dividend growth rates, respectively. We find that the slope coefficient on μ_{dt} does indeed increase significantly from 0.17 on average to 0.28 during downside months, indicating downside stickiness.

Our analysis of predictive dynamics in dividend growth is based on announced dividends rather than the commonly used paid-out dividends, which lag dividend announcements. To provide evidence on the importance of this distinction, we sample both announced and paid dividends at the monthly, quarterly, and annual horizons and estimate our components model separately on these series. Next, we produce forecasts of dividend growth for the ensuing period. We find that predictability of dividend growth is stronger if we use forecasts based on announced dividends than if we use forecasts based on paid dividends. Interestingly, some advantage from using announced rather than distributed dividends to generate forecasts remain even at the annual horizon. Further details are provided in the Internet Appendix.

D. Cash Flow News and Economic Activity

We next examine the relation between our estimate of the persistent component of dividend growth news and two measures of macroeconomic growth, namely, GDP growth and consumption growth, both of which have been examined in previous work such as Liew and Vassalou (2000) and Bansal and Yaron

(2004).[25] Specifically, to evaluate the statistical significance of these relations, we estimate the predictive regressions

$$\Delta y_{t+1} = \alpha + \beta_1 \mu_{dt} + \beta_2 \Delta y_t + \varepsilon_{t+1}, \tag{14}$$

where Δy_{t+1} is the future change in log GDP or log consumption. We include one lag of the dependent variable, Δy_t, to control for persistence in consumption or GDP growth. Table IV reports the results from the regression in (14). We first consider the results based on the estimate of μ_{dt} extracted from daily data (Panel A). In the univariate regressions, our persistent dividend growth measure, μ_{dt}, generates positive coefficients of 0.15 with t-statistics of 5.68 and 6.58 for GDP growth and consumption growth, respectively. Moreover, with R^2s of 23.5% and 31.1%, μ_{dt} has strong predictive power for future GDP and consumption growth. The coefficient on μ_{dt} remains statistically significant after we add a lag of the dependent variable, although the t -statistic decreases to around three.

To evaluate the importance of the data frequency, we also report results for cases in which μ_{dt} is extracted from either monthly dividend growth data (Panel B) or quarterly dividend growth data (Panel C). While the μ_{dt} estimates extracted from the lower frequency data remain statistically significant in helping predict both GDP and consumption growth, the R^2s decline notably—particularly in the quarterly data for which we obtain R^2s of 5.9% and 9.4%, compared with 23.5% and 31.1% for the daily data. High-frequency data can therefore significantly boost the predictive power of our dividend growth estimate.

We conclude from this evidence that our persistent cash flow measure μ_{dt} helps predict variation in macroeconomic growth. This is consistent with our earlier finding that μ_{dt} predicts future dividend growth and shows that this result carries over to broader measures of economic growth. Taken together, our results indicate that firms adjust their dividend payments in anticipation of changes in economic conditions measured by slow-moving economic indicators such as GDP and consumption growth.

E. Relation to Other Activity Measures

Our daily dividend growth measure reflects general macroeconomic conditions and thus can be viewed as an economic indicator similar to existing measures such as the macroeconomic uncertainty measure of Jurado, Ludvigson, and Ng (2015), the economic policy uncertainty measure of Baker, Bloom, and Davis (2016), the Aruoba, Diebold, and Scotti (2009) (ADS) business conditions index, the credit spread indicator of Gilchrist and Zakrajsek (2012),

[25] The GDP series is downloaded from the Federal Reserve Economic Data (FRED) and is seasonally adjusted. Consumption expenditures are the sum of nondurable consumption plus services from Table 2.3.5 of the National Income and Product Accounts (NIPAs) from the Bureau of Economic Analysis (BEA) website.

Table IV

Predictive Regressions of GDP and Consumption Growth on the Persistent Dividend Growth Component

This table reports estimates from quarterly predictive regressions of future GDP and consumption growth, Δy_{t+1}, on the persistent dividend growth component, μ_{dt}, estimated from our daily dividend growth model (Panel A) or from dividend data sampled at the monthly (Panel B) or quarterly (Panel C) frequencies. Square brackets report t-statistics computed using Newey-West standard errors with three lags. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. The sample period for these regressions is 1973 to 2015.

Panel A: $\Delta y_{t+1} = \alpha + \beta\mu_{daily} + \gamma\Delta y_t + \varepsilon_{t+1}$

	ΔGDP		$\Delta Consumption$	
	(1)	(2)	(3)	(4)
μ_{daily}	0.15***	0.09***	0.15***	0.04***
	[5.68]	[3.30]	[6.58]	[2.84]
Δy_t		0.37***		0.67***
		[3.87]		[12.16]
R^2	23.52%	33.56%	31.08%	59.97%
Observations	171	171	171	171

Panel B: $\Delta y_{t+1} = \alpha + \beta\mu_{monthly} + \gamma\Delta y_t + \varepsilon_{t+1}$

	ΔGDP		$\Delta Consumption$	
	(1)	(2)	(3)	(4)
$\mu_{monthly}$	0.13***	0.07***	0.13***	0.04*
	[4.28]	[2.53]	[4.80]	[1.90]
Δy_t		0.40***		0.68***
		[4.59]		[11.15]
R^2	20.11%	32.37%	27.36%	59.85%
Observations	171	171	171	171

Panel C: $\Delta y_{t+1} = \alpha + \beta\mu_{quarterly} + \gamma\Delta y_t + \varepsilon_{t+1}$

	ΔGDP		$\Delta Consumption$	
	(1)	(2)	(3)	(4)
$\mu_{quarterly}$	0.08**	0.04	0.09***	0.02
	[2.06]	[1.44]	[2.65]	[1.63]
Δy_t		0.49***		0.74***
		[5.77]		[17.22]
R^2	5.85%	28.68%	9.36%	58.64%
Observations	171	171	171	171

and "noise" in the Treasury market (Hu, Pan, and Wang (2013)).[26] Previous research addresses whether these measures can be used to predict the state

[26] Aruoba, Diebold, and Scotti (2009) measure economic activity at the daily frequency using a variety of stock and flow data observed at mixed frequencies. Their approach extracts the state of the business cycle from a latent factor that affects all observed variables. Jurado, Ludvigson,

Table V

Correlations between the Persistent Dividend Growth Component μ_{dt} and Macroeconomic and Financial Activity Measures

This table reports correlations between the persistent dividend growth component μ_{dt} extracted from our daily cash flow model and the following daily macroeconomic variables/indicators: the VIX index, the policy uncertainty index of Baker, Bloom, and Davis (2016), the ADS index of Aruoba, Diebold, and Scotti (2009), the liquidity noise index of Hu, Pan, and Wang (2013), and the daily inflation index of Cavallo and Rigobon (2016). Panel A correlates the levels of these variables, while Panel B correlates changes in the variables. *p*-Values are reported below the correlation estimates. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. Sample periods vary according to the length of the individual series.

Panel A: Persistent Dividend Growth and Indexes (Levels)

	VIX	PU	ADS	Liquidity	Inflation	μ_{dt}
VIX	1					
PU	0.34***	1				
	(0.00)					
ADS	−0.49***	−0.26***	1			
	(0.00)	(0.00)				
Liquidity	0.67***	0.19***	−0.54***	1		
	(0.00)	(0.00)	(0.00)			
Inflation	−0.66***	−0.36***	0.50***	−0.63***	1	
	(0.00)	(0.00)	(0.00)	(0.00)		
μ_{dt}	−0.54***	−0.20***	0.20***	−0.51***	0.76***	1
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	

Panel B: Persistent Dividend Growth and Indexes (Changes)

	ΔVIX	ΔPU	ΔADS Index	ΔLiquidity	ΔInflation	$\Delta\mu_{dt}$
ΔVIX	1					
ΔPU	−0.02	1				
	(0.10)					
ΔADS Index	0.01	0.00	1			
	(0.65)	(0.94)				
ΔLiquidity	−0.01	0.05***	−0.00	1		
	(0.53)	(0.00)	(0.95)			
ΔInflation	0.03	−0.03	−0.00	−0.00	1	
	(0.24)	(0.27)	(0.89)	(0.94)		
$\Delta\mu_{dt}$	−0.02	−0.00	0.04	−0.01	0.10***	1
	(0.16)	(0.77)	(0.66)	(0.58)	(0.00)	

of the economy, especially during recessions and financial crises, so we next explore the relation between μ_{dt} and these alternative measures.

Panel A of Table V reports estimates of the correlations between the persistent dividend component μ_{dt} and these daily measures of financial and

and Ng (2015) provide econometric estimates of time-varying macroeconomic uncertainty and show that important uncertainty episodes appear far more infrequently than indicated by popular uncertainty proxies. However, when such episodes do occur, they tend to be larger, more persistent, and more correlated with real economic activity.

macroeconomic conditions. We find that μ_{dt} has a highly significant negative correlation of −0.54 with the VIX, suggesting that dividend growth is lower in times with high uncertainty, which tends to coincide with economic recessions. Further, μ_{dt} has a significantly negative correlation of −0.20 with the policy uncertainty index of Baker, Bloom, and Davis (2016) and a negative correlation of −0.51 with the liquidity noise index of Hu, Pan, and Wang (2013), indicating that firm payouts are lower in times with greater uncertainty. Finally, μ_{dt} has a significantly positive correlation of 0.20 with the ADS index of Aruoba, Diebold, and Scotti (2009) and with the daily inflation index of Cavallo and Rigobon (2016) (correlation of 0.76). These findings show that our persistent dividend growth measure is significantly negatively correlated with risk proxies, for example, stock market volatility and policy uncertainty, but positively correlated with economic growth and inflation.

Panel A in Table V uses levels, and thus, the correlation estimates described above are driven by common, persistent factors reflecting the state of the economy. Panel B highlights short-run effects by reporting the correlations between daily changes in the underlying indexes. Changes in our daily dividend growth index are only significantly correlated with changes in daily inflation, suggesting that daily variation in our new dividend growth measure is not captured by daily variation in other financial and macroeconomic variables but rather constitutes a new source of (high-frequency) information.

IV. Implications for Return Predictability

A long debate in the asset pricing literature relates to the importance of cash flows expectations as a source of variation in stock returns. For example, Cochrane (2008) argues that dividend growth cannot be predicted for the most part by means of variables such as the dividend-price ratio. If true, this has important asset pricing implications because it implies that movements in the dividend-price ratio must instead reflect time-varying risk premia.

Finding that the dividend-price ratio fails to predict future dividend growth does not, of course, rule out the possibility that other variables—such as our persistent dividend growth measure—can predict dividend growth and, in turn, stock returns. In this section, we explore whether such return predictability from our persistent dividend growth measure can be established and whether it is consistent with basic principles of asset pricing.

A. A Predictive Present Value Model

We use a simple reduced-form present value model to explore the implications of dividend growth predictability for the predictability of stock returns. This model provides a structured framework that allows us to quantify and test the implications of dividend growth predictability for the predictability of stock returns.[27]

[27] Other, less parametric approaches to asset pricing are also available, for example, the pricing kernel bounds proposed by Hansen and Jagannathan (1997). An important advantage of the

Our analysis of return dynamics uses the following notations. Let p_{t+1} denote the logarithm of the stock price on day $t+1$, while d^{cf}_{t+1} is the logarithm of the dividends *paid out* (distributed) on day $t+1$ and, as before, d_{t+1} is the logarithm of dividends *announced* on day $t+1$. The present value model refers to paid-out cash flows (*cf*) and the two dividend measures are very different at the daily horizon—recall that, on average, dividend announcements precede distributions by 42 days. At longer horizons such as a quarter or a year, differences between announced and distributed dividends tend to be much smaller.[28] Finally, $\Delta d^{cf}_{t+1} = d^{cf}_{t+1} - d^{cf}_t$ denotes the growth rate in distributed dividends, while r_{t+1} is the log return on day $t+1$.

Following Campbell and Shiller (1988a), consider the approximate log-linearized present value model

$$r_{t+1} \approx k + \rho(p_{t+1} - d^{cf}_{t+1}) + \Delta d^{cf}_{t+1} + (d^{cf}_t - p_t), \tag{15}$$

where $\rho = 1/(1 + \exp(\overline{d - p})$ and $k = -\ln(\rho) - (1-\rho)\ln(1/\rho - 1)$ are log-linearization constants. At the daily frequency, the average dividend yield is very small and so $\rho = 0.99998$ is extremely close to unity. Using this value for ρ, the log-linearization in (15) is highly accurate with a correlation between exact and approximate returns of 0.9998.

Again, we emphasize the distinction between announced and distributed dividends for the daily present value model: at the daily horizon, the correlation between growth in announced and distributed dividends (Δd_{t+1} and Δd^{cf}_{t+1}) is essentially zero (-0.0072). Moreover, if we (incorrectly) use announced instead of distributed dividends in (15), the accuracy of the present value approximation deteriorates sharply.

Studying return predictability in the context of the present value model in (15) offers a number of important advantages. First, previous studies of return predictability, such as Campbell and Shiller (1988a), Cochrane (2008), and van Binsbergen and Koijen (2010), use this model.[29] Adopting a similar framework makes it easier to compare our findings to existing literature.

Second, the present value model (15) has strong economic implications that directly link predictability of dividend growth to predictability of stock returns not only from popular predictors such as the log dividend-price ratio but also from other potential predictors such as news about our persistent dividend growth component. In particular, (15) implies that if a variable helps predict dividend growth, Δd^{cf}_{t+1}, it should also predict returns—unless this variable has an equivalent and opposite effect on the price-dividend ratio, $p_{t+1} - d^{cf}_{t+1}$. This

present value model is that it lends itself to empirical tests through a set of linear restrictions on the model parameters.

[28] The correlations between log-dividends calculated using announced and paid-out dividends are 33.23%, 72.45%, and 83.61% at the monthly, quarterly, and annual frequencies, respectively.

[29] Jagannathan and Liu (2019) embed a dividend growth VAR specification in an asset pricing model with recursive Epstein-Zin preferences and find significant predictability of annual stock returns from a stock yield variable, particularly in the presence of investor learning about dividend growth dynamics.

latter scenario arises if a variable forecasts equal movements in expected cash flow growth and discount rates so that the net effect on return predictability is zero.

Third, the present value model has implications for return predictability at different horizons. Because (15) holds regardless of the horizon, if a variable helps forecast dividend growth at multiple horizons, we would, generally, expect the same variable to possess predictive power for stock returns across these horizons. This is relevant here, given our empirical evidence that the persistent dividend growth component helps predict dividend growth at short as well as long horizons.

At the long forecast horizons conventionally used in the finance literature, it is common to predict the dividend yield, $d^{cf}_{t+1} - p_{t+1}$, because it is easy to extract a highly persistent component from the dividend-price ratio that is very smooth; see, for example, Cochrane (2008). At the daily horizon, however, $d^{cf}_{t+1} - p_{t+1}$ tends to be dominated by large variations in daily dividend distributions although these variations are temporary and thus do not reflect predictability in long-term dividend prospects.

To address this difference and obtain a more stable forecasting model, we instead predict daily *changes* in the dividend-price ratio $\Delta(d^{cf}_{t+1} - p_{t+1}) = (d^{cf}_{t+1} - p_{t+1}) - (d^{cf}_{t} - p_{t})$. From an econometric perspective, the main effect of predicting $\Delta(d^{cf}_{t+1} - p_{t+1})$ rather than $d^{cf}_{t+1} - p_{t+1}$ is to anchor daily variation in the dividend-price ratio on a random walk, imposing a coefficient of unity on $d^{cf}_{t} - p_{t}$. Ferreira and Santa-Clara (2011) find that imposing a similar constraint on the dividend yield reduces estimation error and leads to more accurate forecasts of stock returns. The constraint can be viewed as a sensible economic prior since we would expect the dividend-price ratio to contain a highly persistent (near unit root) component at the daily frequency.

Following this discussion, we approximate (15) using

$$r_{t+1} \approx k - \Delta\left(d^{cf}_{t+1} - p_{t+1}\right) + \Delta d^{cf}_{t+1}. \tag{16}$$

The approximation in (16) is highly accurate as a result of ρ being close to unity at the daily horizon, the correlation between r_{t+1} and the approximation in (16) is 0.9997 in our sample.

Note that Δd^{cf}_{t+1} can be dropped from (16), so $r_{t+1} \approx k + (p_{t+1} - p_t)$, suggesting a predictive regression that simply projects capital gains, Δp_{t+1}, on candidate predictors. However, daily capital gains and stock returns are similar, so regressing capital gains rather than returns on the predictors does not provide additional economic insights. In contrast, stock prices can move because of changes in valuations (reflected in the dividend-price ratio) or because of cash flow news (dividend growth). Our decomposition in (16) allows us to identify predictability in these separate components, which offers the potential for new economic insights on return predictability at the daily horizon and for testing the validity of the cross-equation restrictions implied by the present value model (15).

Next, consider the list of predictors. Our first predictor is the lagged change in the persistent dividend growth component, $\Delta\mu_{dt}$. Using μ_{dt} as a predictor might seem the natural choice. However, μ_{dt} is constructed to capture the predictive component in *announced* dividends, whereas the present value model (16) is based on *distributed* dividends, which are uncorrelated with announced dividends at the daily frequency. Moreover, given its very high persistence, μ_{dt} is not well suited for predicting daily stock returns, which are not very persistent. In contrast, $\Delta\mu_{dt}$ effectively measures innovations to μ_{dt} and so captures news about changes in the state of the economy.

Our second predictor is the lagged dividend-price ratio, which is extensively used in the finance literature to forecast variation in the investment opportunity set and thus is important to include. To see why, note from (15) that variation in the log dividend-price ratio should reflect changes in expected future returns and/or changes in expected future dividend growth,

$$d_t^{cf} - p_t \propto E_t\left[\sum_{j=0}^{\infty} \rho^j \left(r_{t+j+1} - \Delta d_{t+j+1}^{cf}\right)\right]. \tag{17}$$

At the frequencies typically considered in empirical studies—monthly, quarterly, or annual—the dividend-price ratio tends to be very smooth and highly persistent. This makes it well suited as a predictor of slow movements in the investment opportunity set operating at longer horizons. In contrast, because of the lumpiness in daily dividends, the daily dividend-price ratio is very volatile. To obtain a predictor variable with persistence features more similar to the conventional dividend-price ratio obtained at lower frequencies, we measure the dividend-price ratio as a smoothed 30-day moving average of announced dividends divided by the current stock price, denoted by $\overline{d_t - p_t}$.[30]

Finally, we include the lagged value of daily dividend growth, Δd_t^{cf}, to account for short-term reversals in the daily dividend growth process, as days with very large dividend distributions tend to be followed by days with much smaller dividend payments.

Following this discussion, our daily prediction model for growth in distributed dividends takes the form

$$\Delta d_{t+1}^{cf} = \theta_0 + \theta_\mu \Delta\mu_{dt} + \theta_{dp}\left(\overline{d_t - p_t}\right) + \theta_d \Delta d_t^{cf} + \varepsilon_{dt+1}. \tag{18}$$

We predict changes in the daily dividend-price ratio using the same set of predictors,

$$\Delta\left(d_{t+1}^{cf} - p_{t+1}\right) = \gamma_0 + \gamma_\mu \Delta\mu_{dt} + \gamma_{dp}\left(\overline{d_t - p_t}\right) + \gamma_d \Delta d_t^{cf} + \varepsilon_{t+1}^{cf}. \tag{19}$$

[30] Our empirical results are robust to using other windows to smooth the dividend process.

Using equations (18) and (19) in the log-linearized return equation (16), we have

$$r_{t+1} \approx k + \theta_0 - \gamma_0 + (\theta_\mu - \gamma_\mu)\Delta\mu_{dt} + (\theta_{dp} - \gamma_{dp})(\overline{d_t - p_t}) + (\theta_d - \gamma_d)\Delta d_t^{cf} + \varepsilon_{dt+1} - \varepsilon_{t+1}^{cf}. \quad (20)$$

This equation can be rewritten as

$$r_{t+1} = \lambda_0 + \lambda_\mu \Delta\mu_{dt} + \lambda_{dp}(\overline{d_t - p_t}) + \lambda_d \Delta d_t^{cf} + \varepsilon_{rt+1}, \quad (21)$$

where $\varepsilon_{rt+1} = \varepsilon_{dt+1} - \varepsilon_{t+1}^{cf}$. Comparing (21) to (20), it follows that the present value model imposes the following cross-equation restrictions on the parameters of the return equation:

$$\begin{aligned} \lambda_0 &= k + \theta_0 - \gamma_0, \\ \lambda_\mu &= \theta_\mu - \gamma_\mu, \\ \lambda_{dp} &= \theta_{dp} - \gamma_{dp}, \\ \lambda_d &= \theta_d - \gamma_d. \end{aligned} \quad (22)$$

The restriction on λ_{dp} implies that if, at any horizon, $\overline{d_t - p_t}$ predicts future dividend growth ($\theta_{dp} \neq 0$ in (18)), it must also predict future returns ($\lambda_{dp} \neq 0$). Similarly, if, after controlling for the effect of the log dividend-price ratio, some other predictor such as $\Delta\mu_{dt}$ forecasts dividend growth ($\theta_\mu \neq 0$ in (18)), then the present value model requires that this variable also forecasts returns ($\lambda_\mu \neq 0$) unless $\theta_\mu = \gamma_\mu$, in which case the predictive effects from $\Delta\mu_{dt}$ on Δd_{t+1}^{cf} and $\Delta(d_{t+1}^{cf} - p_{t+1})$ cancel out. This case arises if the variable predicts both higher (lower) dividend growth and higher (lower) discount rates in equal amounts (see equation (17)).[31]

An important implication of (21) and (22) is that dividend news will generally also contain information about discount rates. Hence, care should be exercised when interpreting the effect of dividend news on variables such as stock returns and the dividend-price ratio since such news may also have sizeable discount rate effects. This interpretation is reinforced by our finding that the persistent dividend growth component helps predict consumption growth, which, of course, is an important component of the pricing kernel in standard consumption-based asset pricing models.

[31] The model in Cochrane (2008) is a special case of our specification that is obtained by omitting $\Delta\mu_{dt}$ and Δd_t^{cf} as predictor variables (i.e., setting $\lambda_\mu = \lambda_d = 0$), assuming that the distributed and announced dividends are the same, and setting $\rho \approx 1$, which, as we have shown, holds to a very close approximation for daily data.

B. Empirical Findings

Before proceeding with the empirical analysis of the present value model, it is important to note that the behavior of the components in (15) is quite different at the daily frequency than at the longer horizons conventionally used in empirical studies.

First, as noted previously, the distinction between announced and distributed dividends is not important at longer horizons but is crucial at the daily horizon for which the correlation is only 52.12% and 0.72% for the two dividend series in log-levels and log-changes, respectively.

Second, the contribution to return volatility of the individual return components in (15) is very different at the daily and longer frequencies. For example, in quarterly data, Δd^{cf}_{t+1} is quite smooth and its volatility is roughly six times smaller than the volatility of the dividend-price ratio, $d^{cf}_{t+1} - p_{t+1}$. Conversely, in daily data, the volatility of Δd^{cf}_{t+1} is almost 50% *higher* than the volatility of $d^{cf}_{t+1} - p_{t+1}$ due to shifts in the composition of firms paying dividends on any given day. Hence, the volatility of dividend growth is nearly an order of magnitude greater at the daily than at the quarterly horizon, measured relative to the importance of variation in the dividend-price ratio.

Third, days with unusually large dividend payments (large, positive dividend growth) are typically followed by days with smaller dividend payments (negative dividend growth). This produces negative first-order autocorrelation in daily dividend growth—the first-order autocorrelation in Δd^{cf}_t is -0.35—which does not carry over beyond one day, however, and so is not helpful for predicting long-term dividend growth.

These observations are important for interpreting our empirical estimates and suggest that we should only expect to find limited predictability of dividend growth and of the log dividend-price ratio at the daily frequency using our three predictors.

The top row in Table VI presents estimates from the predictive dividend growth regression (18). We show results for the full sample (i.e., 1927 to 2016), but the estimates are very similar in the recent sample (i.e., 1973 to 2016). Both the lagged dividend-price ratio and lagged daily dividend growth obtain significantly negative coefficients. The negative coefficients on $\overline{d_t - p_t}$ and Δd^{cf}_t are largely a consequence of the short-term negative serial correlation (reversal) in daily dividend growth, which also explains the large t -statistic on Δd^{cf}_t. In contrast, the coefficient on news about persistent growth in announced dividends ($\Delta \mu_{dt}$) fails to be significant. This is not surprising given the lead time of dividend announcements relative to dividend distributions.

Turning to the prediction equation for the change in the dividend-price ratio, (19), the second row in Table VI shows that the coefficient on $\Delta \mu_{dt}$ is positive but insignificant. The coefficient on the dividend-price ratio is also insignificant, whereas lagged growth in cash flows, Δd^{cf}_t, produces a negative and significant coefficient, again as a result of short-term reversals in daily dividend distributions.

Table VI

Parameter Estimates and Test Statistics for the Predictive Present Value Model Fitted to Daily Data

This table reports parameter values and test statistics for a predictive present value model estimated at the daily frequency over the period 1927 to 2016,

$$\Delta d^{cf}_{t+1} = \theta_0 + \theta_\mu \Delta\mu_{dt} + \theta_{dp}(\overline{d_t - p_t}) + \theta_d \Delta d^{cf}_t + \varepsilon_{dt+1},$$

$$\Delta(d^{cf}_{t+1} - p_{t+1}) = \gamma_0 + \gamma_\mu \Delta\mu_{dt} + \gamma_{dp}(\overline{d_t - p_t}) + \gamma_d \Delta d^{cf}_t + \varepsilon^{cf}_{t+1},$$

$$r_{t+1} = \lambda_0 + \lambda_\mu \Delta\mu_{dt} + \lambda_{dp}(\overline{d_t - p_t}) + \lambda_d \Delta d^{cf}_t + \varepsilon_{rt+1},$$

with the following restrictions implied by the model:

$$\lambda_0 = k + \theta_0 - \gamma_0,\ \lambda_\mu = \theta_\mu - \gamma_\mu,\ \lambda_{dp} = \theta_{dp} - \gamma_{dp},\ \lambda_d = \theta_d - \gamma_d.$$

***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

	Const.	$\Delta\mu_{dt}$	$\overline{d_t - p_t}$	Δd^{cf}_t	R^2
Δd^{cf}_{t+1}	−0.24**	13.93	−0.04**	−0.36***	12.87%
	(−2.31)	(0.30)	(−2.54)	(−58.32)	
$\Delta(d^{cf}_{t+1} - p_{t+1})$	−0.23**	12.72	−0.04	−0.36***	12.87%
	(−2.31)	(0.28)	(−2.53)	(−58.32)	
ret_{t+1}	0.00	1.28***	0.00	−0.00	0.08%
	(0.86)	(3.31)	(0.27)	(−0.14)	
p-Value of restriction	0.37	0.87	0.13	0.75	

Finally, in the prediction equation for daily stock returns, (21), only news about the persistent dividend growth component, $\Delta\mu_{dt}$, turns out to be significant. The estimated coefficient on $\Delta\mu_{dt}$ is 1.28 with a t-statistic of 3.31. This indicates that positive news about persistent dividend growth is associated with larger stock returns the following day.[32]

To see if these results are consistent with the asset pricing implications that follow from the present value model, the bottom row in Table VI reports the outcome of Wald tests of the cross-equation restrictions in (22). We fail to reject any of the coefficient restrictions, which suggests that the relatively precisely estimated coefficient on $\Delta\mu_{dt}$ in return equation (21) is consistent with the present value model. In fact, the estimated coefficient on $\Delta\mu_{dt}$ (1.28) is very close to the value implied by the present value model (1.21).

Similarly, the present value model implies small coefficients on both $\overline{d_t - p_t}$ and Δd^{cf}_t in the return equation, which again is consistent with our estimates of (21).[33] Although these variables are significant predictors of variation in both daily dividend growth and changes in the dividend-price ratio, the two effects

[32] To be consistent with standard approaches to testing the present value model, we use one-period-ahead returns, r_{t+1}, in these regressions. However, the results are robust if we instead skip one day and use r_{t+2}, as the results in Table VII show.

[33] Similar results hold in the recent sample (i.e., 1973 to 2016), where we again fail to reject the cross-equation restrictions in (22).

cancel out and so these variables do not have predictive power over returns. This case arises when the predictors correlate with future cash flow growth and discount rates in the same amount but in opposite directions.

In Internet Appendix Section III, we assess the economic value of return predictability from our μ_{dt} measure. We find that using information on μ_{dt} leads to a 2.32% increase in the annualized certainty-equivalent return for a mean-variance investor with moderate risk aversion.

C. Return Predictability at Longer Horizons

A large empirical literature in finance examines return predictability at longer horizons such as a month or a quarter. To make our results comparable to this literature and to explore whether our new dividend growth measure possesses predictive power for returns at longer horizons, we conduct two exercises.

First, we regress multiday stock returns $r_{t+2::t+h} = \sum_{\tau=t+2}^{t+h} r_\tau$ on $\Delta\mu_{dt}$ in addition to a range of alternative measures of dividend growth or dividend dynamics such as the "raw" daily dividend growth, Δd_t, news about persistent dividend growth extracted from a model without a jump component, $\Delta\mu_{dt}^{NJ}$, jumps in the dividend growth process, $J_{dt}\xi_t$, the stochastic volatility component extracted from the dividend growth process, $h_{dt}/2$, and temporary shocks to dividend growth, ε_{dt+1}. These regressions use a conservative approach and skip one day to compute stock returns, which therefore start on day $t+2$. We do this to show that our results are not due to price slippage effects.

Table VII reports results for different models and forecast horizons ranging from one through three, five, and 21 trading days. We find that $\Delta\mu_{dt}$ has significant predictive power for stock returns at one-, three-, and five-day horizons as well as one month (21 trading days) ahead. This finding strongly depends on allowing for jumps in the dividend growth process to extract $\Delta\mu_{dt}$. We find essentially no return predictability from the alternative measure that excludes jumps, $\Delta\mu_{dt}^{NJ}$. We also find little to no return predictability from raw daily dividend growth, Δd_t, or from any of the other predictors.

Second, we consider the more common monthly horizon for which return predictability from variables such as the dividend-price ratio has been established in the literature. In particular, we compare return predictability from $\Delta\mu_{dt}$ to predictability from a set of 14 standard predictors used by Goyal and Welch (2008). Following common practice, we focus on univariate predictive regressions conducted one predictor at a time. We use the 20-year window 1973:01 to 1992:12 as an initial training sample and compute out-of-sample return forecasts for the remaining 1993:01 to 2016:12 period, employing an expanding estimation window that adds new observations as they become available.

We consider out-of-sample R^2s and average utility gains. The out-of-sample R^2 measures the percent reduction in mean squared forecast error obtained when conditioning on a predictor relative to using the prevailing mean forecast. Average utility gains can be interpreted as the portfolio management fees that an investor with mean variance preferences and a risk aversion coefficient of

Table VII

Predictive Regressions for Stock Returns

This table reports estimates of predictive regressions using cumulative stock market returns as the dependent variable and the following predictor variables: (i) daily growth in announced dividends, Δd_t; (ii) changes in the persistent dividend growth component $\Delta\mu_t^{NJ}$ estimated from a model without jumps and stochastic volatility; and (iii) changes in the persistent dividend growth component $\Delta\mu_t$ estimated from a model that accounts for jumps and stochastic volatility. The dependent variables are the one-day-ahead return r_{t+2} (columns (1) to (3)), the three-day-ahead cumulative return $r_{t+2:t+4}$ (columns (4) to (6)), the five-day-ahead cumulative return $r_{t+2:t+6}$ (columns (7) to (9)), and the one-month (21-trading-day) ahead cumulative return $r_{t+2:t+22}$ (columns (10) to (12)). We skip one day in the predictive regressions to account for dividends announced after trading hours. t-Statistics, reported in square brackets, are computed using Newey-West standard errors with two lags for one-, three-, and five-day-ahead cumulative returns and 10 lags for one-month-ahead cumulative returns. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. The sample period is 1927 to 2016.

	Return Horizon											
	r_{t+2}			$r_{t+2:t+4}$			$r_{t+2:t+6}$			$r_{t+2:t+22}$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Δd_t	−0.00**			−0.00			−0.00			−0.00		
	[−2.19]			[−0.07]			[−0.17]			[−0.32]		
$\Delta\mu_{dt}^{NJ}$		−0.00**			−0.00			−0.00			−0.00	
		[−2.20]			[−0.00]			[−0.08]			[−0.26]	
$\Delta\mu_{dt}$			1.14***			3.17***			5.40***			21.60***
			[3.36]			[3.90]			[4.77]			[4.98]
$\xi_{dt} J_{dt}$			−0.00			0.00			0.00			0.00
			[−0.04]			[1.49]			[0.53]			[0.37]
$h_{dt}/2$			0.00			0.00			0.00			0.01
			[0.13]			[0.85]			[1.52]			[0.72]
ε_{dt+1}			−0.00			0.00			−0.00			−0.00
			[−1.16]			[0.37]			[−0.30]			[−0.77]
R^2	0.03%	0.03%	0.09%	0.00%	0.00%	0.22%	0.00%	0.00%	0.37%	0.00%	0.00%	1.27%
Observations	20,965	20,965	20,965	20,963	20,963	20,963	20,961	20,961	20,961	20,945	20,945	20,945

Table VIII

Out-of-Sample Return Predictability

This table compares the out-of-sample predictive performance of our $\Delta\mu_{dt}$ measure to that of 14 standard predictors used by Goyal and Welch (2008) to forecast stock returns for the period 1993:01 to 2016:12. The predictor variables are the log dividend-price ratio (log(DP)), the log dividend yield (log(DY)), the log earnings-price ratio (log(EP)), the log dividend payout ratio (log(DE)), the stock variance (SVAR), the book-to-market ratio (BM), the net equity expansion (NTIS), the Treasury bill rate (TBL), the long-term yield (LTY), the long-term rate of return (LTR), the term spread (TMS), the default yield spread (DFY), the default return spread (DFR), and the inflation rate (INFL). For all predictor variables, we use a 20-year initial estimation window from 1973:01 to 1992:12 followed by an expanding estimation window that adds new data points to the sample as they become available. We report both the out-of-sample R^2 and average utility gains for the full out-of-sample period as well as separately for recession and expansion months as defined by the NBER. For any given predictor, the out-of-sample R^2 measures the percentage reduction in the mean squared forecast error associated with a model that uses this predictor relative to the mean squared forecast error of forecasts based on the historical mean return. Positive values indicate better performance than this benchmark, while negative values suggest worse performance. The average utility gain can be interpreted as the portfolio management fee that an investor with mean variance preferences and a risk aversion coefficient of five would be willing to pay to have access to the forecasts generated by a particular predictor, again measured relative to the historical average forecasts.

	Full Sample			Expansions			Recessions		
	$R^2_{OS}(\%)$	$p-val$	$\Delta U(\%)$	$R^2_{OS}(\%)$	$p-val$	$\Delta U(\%)$	$R^2_{OS}(\%)$	$p-val$	$\Delta U(\%)$
$\Delta\mu_{dt}$	0.43	0.14	1.28	−0.02	0.35	0.11	1.68	0.03	11.40
log(DP)	−1.98	0.97	−2.44	−3.06	0.99	−3.63	1.01	0.09	7.66
log(DY)	−2.01	0.96	−2.28	−3.25	0.99	−3.73	1.41	0.04	10.18
log(EP)	−0.85	0.67	0.74	−0.94	0.86	−1.17	−0.61	0.48	17.45
log(DE)	−2.04	0.80	−0.77	−1.29	0.99	−1.24	−4.13	0.65	3.19
SVAR	1.49	0.12	1.30	0.08	0.34	0.10	5.38	0.13	11.80
BM	−0.45	0.80	−0.77	−0.51	0.75	−0.96	−0.28	0.82	1.00
NTIS	−2.57	0.97	−1.37	−1.31	0.72	−0.23	−6.04	0.99	−11.33
TBL	−0.55	0.43	−0.57	0.42	0.12	0.93	−3.22	0.98	−13.20
LTY	−0.33	0.65	−0.14	−0.17	0.50	−0.17	−0.76	0.82	0.37
LTR	−0.20	0.31	−0.31	−0.28	0.33	−0.38	0.01	0.41	−0.04
TMS	−1.20	0.65	−1.33	−1.11	0.51	−0.22	−1.46	0.80	−11.35
DFY	−3.24	0.99	−4.13	−3.02	0.96	−2.66	−3.84	0.96	−17.29
DFR	−2.13	0.41	1.08	−2.17	0.49	−0.25	−2.00	0.41	12.41
INFL	−0.75	0.79	−1.06	−0.33	0.53	−0.30	−1.89	0.89	−7.62

five is willing to pay to have access to the forecasts generated using a predictor, again computed relative to the prevailing mean.

The results, reported in Table VIII, show that $\Delta\mu_{dt}$ has significant out-of-sample predictive power for stock returns, coming in second among the 15 predictor variables in terms of both the strength of its out-of-sample predictive power and its utility gains. Moreover, consistent with findings in recent literature (e.g., Rapach, Strauss, and Zhou (2010), Dangl and Halling (2012)), the degree of return predictability, as well as the size of the utility gains, is larger during recessions than in expansions.

V. Dividend News and Return Dynamics

The present value model can also be used to quantify and assess the effect of dividend news on concurrent (same day) stock returns. In particular, the model has implications for which of the different dividend growth components is likely to most strongly affect stock returns. Shocks to the dividend growth process should be positively correlated with unexpected returns and shocks to the persistent (predictable) growth component should generally have a larger impact on stock returns than shocks to the temporary growth component. In this section, we test these implications and provide a broader analysis of how the different dividend news components affect not only the mean of stock returns but also the volatility and probability of observing a jump in stock returns.[34]

A. Dividend News and Contemporaneous Stock Returns

To analyze how stock returns are affected contemporaneously by shocks to the dividend growth process, we consider two alternative formulations of the prediction equations underlying the present value model. Although both formulations follow from the present value model and thus are equivalent, considering both allows us to explore the robustness of our results with regard to how we estimate and test the model. In both cases, note that because μ_{dt+1} is highly persistent, $\Delta\mu_{dt+1}$ is essentially identical to the innovation to the μ_{dt+1} process ($\varepsilon_{\mu t+1}$): the two have a correlation of 0.9999. Hence, $\Delta\mu_{dt+1}$ can be interpreted as the news or shock to the predictable (persistent) component of the dividend growth process, and therefore, we include this variable in our model for contemporaneous movements in stock returns.[35]

Our first specification decomposes daily stock returns into a weighted average of the capital gain ($p_{t+1} - p_t$) and dividend yield ($d^{cf}_{t+1} - p_t$),

$$\begin{aligned} r_{t+1} &\approx k + \rho p_{t+1} - p_t + (1-\rho)d^{cf}_{t+1} \\ &= k + \rho(p_{t+1} - p_t) + (1-\rho)(d^{cf}_{t+1} - p_t). \end{aligned} \tag{23}$$

Next, we relate $p_{t+1} - p_t$ and $d^{cf}_{t+1} - p_t$ to news about the persistent dividend growth component, $\Delta\mu_{dt+1}$, controlling for movements in the smoothed log dividend-price ratio, $\overline{d_t - p_t}$, which captures investors' forward-looking expectations,

$$\begin{aligned} p_{t+1} - p_t &= \pi_0 + \pi_\mu \Delta\mu_{dt+1} + \pi_{dp}(\overline{d_t - p_t}) + \varepsilon_{pt+1}, \\ d^{cf}_{t+1} - p_t &= \delta_0 + \delta_{dp}(\overline{d_t - p_t}) + \varepsilon^{cf}_{t+1}. \end{aligned} \tag{24}$$

[34] Internet Appendix Section III reports results from an analysis of the cross-sectional effects of shocks to the persistent component in the dividend growth process.

[35] Empirical results are essentially identical if we use $\varepsilon_{\mu t+1}$ in our analysis instead of $\Delta\mu_{dt+1}$.

We omit $\Delta\mu_{dt+1}$ from the equation for $d^{cf}_{t+1} - p_t$ because dividend announcements on day $t+1$ should not affect dividends paid out on the same day due to the lengthy lag between when dividends are announced and distributed.

The premise of our first specification, (23) and (24), is that daily dividend payments, d^{cf}_{t+1}, are largely known ahead of time. This is unique to the daily frequency and does not hold at longer horizons such as a quarter or a year. Because of daily composition shifts in the dividend-announcing firms, this does not imply that we would expect a high R^2 for the regression of $d^{cf}_{t+1} - p_t$ on $\overline{d_t - p_t}$ in (24). However, it does mean that investors are able to predict growth in next-day dividend distributions quite accurately. As a result, forecasting r_{t+1} at the daily frequency is largely equivalent to forecasting the capital gain, $p_{t+1} - p_t$.

Combining the present value model in (23) with the regressions in (24), we get

$$r_{t+1} = \lambda_0 + \lambda_\mu \Delta\mu_{dt+1} + \lambda_{dp}(\overline{d_t - p_t}) + \varepsilon_{rt+1}, \tag{25}$$

where $\varepsilon_{rt+1} = \rho\varepsilon_{pt+1} + (1-\rho)\varepsilon^{cf}_{t+1}$. It follows from the present value model (23) that the coefficients in (25) must satisfy the cross-equation restrictions

$$\begin{aligned}
\lambda_0 &= k + \rho\pi_0 + (1-\rho)\delta_0, \\
\lambda_\mu &= \rho\pi_\mu, \\
\lambda_{dp} &= \rho\pi_{dp} + (1-\rho)\delta_{dp}.
\end{aligned} \tag{26}$$

These restrictions are analogous to the constraints obtained for the predictive present value model in (22), which only uses prior-day (lagged) information.

Our second specification uses the present value model (15) with the key difference that we now use the contemporaneous news shock to the persistent dividend growth component, $\Delta\mu_{dt+1}$, instead of its lagged value, $\Delta\mu_{dt}$. This model therefore takes the form,

$$\begin{aligned}
\Delta d^{cf}_{t+1} &= \theta_0 + \theta_\mu \Delta\mu_{dt+1} + \theta_{dp}\left(\overline{d_t - p_t}\right) + \theta_d \Delta d^{cf}_t + \varepsilon_{dt+1}, \\
\Delta\left(d^{cf}_{t+1} - p_{t+1}\right) &= \gamma_0 + \gamma_\mu \Delta\mu_{dt+1} + \gamma_{dp}\left(\overline{d_t - p_t}\right) + \gamma_d \Delta d^{cf}_t + \varepsilon^{cf}_{t+1}, \\
r_{t+1} &= k - \Delta\left(d^{cf}_{t+1} - p_{t+1}\right) + \Delta d^{cf}_{t+1} + \varepsilon_{rt+1}.
\end{aligned} \tag{27}$$

Again, the coefficients in (27) are subject to a set of cross-equation restrictions equivalent to those in (22).

B. Empirical Findings

Panel A of Table IX reports estimates of (24) to (25) fitted to returns data for the full sample (i.e., 1927 to 2016). First consider the capital gains equation in (24). At 1.55, the estimated coefficient on $\Delta\mu_{dt+1}$ is highly significant, while

Table IX

Estimates and Diagnostic Tests for the Contemporaneous Present Value Model Fitted to Daily Data

Panel A reports parameter estimates and t-statistics for the contemporaneous present value model fitted to daily data,

$$\begin{aligned} p_{t+1} - p_t &= \pi_0 + \pi_\mu \Delta\mu_{dt+1} + \pi_{dp}(\overline{d_t - p_t}) + \varepsilon_{pt+1}, \\ d^{cf}_{t+1} - p_t &= \delta_0 + \delta_{dp}(\overline{d_t - p_t}) + \varepsilon^{cf}_{t+1}, \\ r_{t+1} &= \lambda_0 + \lambda_\mu \Delta\mu_{dt+1} + \lambda_{dp}(\overline{d_t - p_t}) + \varepsilon_{rt+1}, \end{aligned}$$

with the following restrictions implied by the model:

$$\lambda_0 = k + \rho\pi_0 + (1-\rho)\delta_0,\ \lambda_\mu = \rho\pi_\mu,\ \lambda_{dp} = \rho\pi_{dp} + (1-\rho)\delta_{dp}.$$

The present value model in Panel B follows the same layout as in Table VI, with the only difference that we replace the lagged term $\Delta\mu_{dt}$ with the contemporaneous value $\Delta\mu_{dt+1}$. Panel C reports estimates of the heteroskedasticity test using the squared residuals from the return model, $\varepsilon^2_{rt+1} = a_0 + a_1\varepsilon^2_{rt} + a_2\sigma^2_{t+1} + a_3\Delta\mu_{dt+1} + \varepsilon^\sigma_{t+1}$. Panel D reports the OLS estimates of the jump test $I_{t+1} = b_0 + b_1\xi_{dt+1}J_{dt+1} + b_2\sigma^2_{dt+1} + b_3\Delta\mu_{dt+1} + \varepsilon^I_{t+1}$. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively. In all cases, the sample period is 1927 to 2016. †: coefficient has been multiplied by 100.

Panel A: Estimates and Tests for Present Value Model (I)

	Const.	$\Delta\mu_{d,t+1}$	$\overline{d_t - p_t}$	R^2
$p_{t+1} - p_t$	−0.00	1.55***	−0.00	0.13%
	(−1.26)	(3.94)	(−1.62)	
$d^{cf}_{t+1} - p_t$	−4.13***		0.94***	14.44%
	(−40.49)		(54.80)	
ret_{t+1}	−0.00	1.57***	−0.01†	0.12%
	(−0.09)	(3.97)	(−0.63)	
p-Value of restriction	0.13	0.98	0.60	

Panel B: Estimates and Tests for Present Value Model (II)

	Const.	$\Delta\mu_{dt+1}$	$\overline{d_t - p_t}$	Δd^{cf}_t	R^2
$= \Delta d^{cf}_{t+1}$	−0.24**	30.28	−0.04**	−0.36***	12.87%
	(−2.31)	(0.66)	(−2.54)	(−58.32)	
$\Delta(d^{cf}_{t+1} - p_{t+1})$	−0.23**	28.86	−0.04	−0.36***	12.87%
	(−2.31)	(0.63)	(−2.54)	(−58.32)	
ret_{t+1}	0.00	1.45***	0.00	−0.00	0.11%
	(0.87)	(3.74)	(0.28)	(−0.14)	
p-Value of restriction	0.36	0.94	0.13	0.75	

Panel C: Heteroskedasticity Test

	Coeff.	Std. Err.	t-Stat	p-Value
Intercept	0.00***	0.00	11.81	0.00
$\varepsilon^2_{r,t}$	0.17***	0.06	3.11	0.00
σ^2_{t+1}	0.00***	0.00	5.56	0.00
$\Delta\mu_{dt+1}$	−0.03	0.02	−1.55	0.12

(*Continued*)

Table IX—*Continued*

Panel D: Jump Test				
	Coeff.	Std. Err.	t-Stat	p-Value
Intercept	0.04***	0.00	21.49	0.00
$\xi_{dt+1}J_{dt+1}$	−0.02*	0.01	−1.70	0.09
σ^2_{dt+1}	0.31***	0.05	6.82	0.00
$\Delta\mu_{dt+1}$	−23.98**	9.59	−2.50	0.01

the coefficient on $\overline{d_t - p_t}$ is statistically insignificant. In turn, $\overline{d_t - p_t}$ is highly statistically significant in the prediction equation for $d^{cf}_{t+1} - p_t$ with a coefficient of 0.94, consistent with $\overline{d_t - p_t}$ picking up a persistent component in the daily dividend-price ratio.

Turning to the return regression (25), the coefficient on $\Delta\mu_{dt+1}$ remains highly statistically significant with an estimate of 1.57, suggesting that higher dividend growth news translates into positive stock returns with a strong and accurately estimated effect. The coefficient on $\overline{d_t - p_t}$ is essentially zero with a t-statistic below unity, confirming our finding in Section IV that movements in this predictor are not important for explaining variation in next-day stock returns, r_{t+1}. The R^2 value of 0.12% shows that news about persistent dividend growth is a nontrivial determinant of daily stock returns.

The bottom row of Table IX, Panel A, reports a set of Wald tests of the cross-equation restrictions in (26). With p -values ranging from 0.13 to 0.98, none of the restrictions is rejected, which suggests that the effects of news about dividend growth are consistent with the asset pricing implications that follow from the present value model in (23). Panel B reports estimates from the present value specification in (27). We find coefficient estimates that are very similar to those obtained from the predictive model analyzed in Section IV. In particular, the estimated coefficient on $\Delta\mu_{dt+1}$ in the return equation equals 1.45 with a highly significant t-statistic of 3.74. This estimate is larger than the estimate obtained from the predictive model (1.28) in Table VI , though a little lower than its value (1.57) obtained from equation (25).

A key motivation for constructing a daily dividend growth measure is that it can help identify the drivers of movements in daily stock returns. From a theoretical perspective, we would expect the three dividend growth components in (2) to have a very different impact on stock prices. For example, temporary shocks to dividend growth should have little effect on stock prices, while shocks to the persistent dividend growth component should have a larger impact. To see if this conjecture holds, we consider whether the other components from our dividend growth model (2) help explain movements in mean returns by expanding (25) to

$$\begin{aligned} r_{t+1} = \lambda_0 + \lambda_\mu \Delta\mu_{dt+1} + \lambda_{dp}(\overline{d_t - p_t}) + \lambda_\varepsilon \varepsilon_{dt+1} + \lambda_h \exp(h_{dt+1}) \\ + \lambda_J J_{dt+1}\xi_{dt+1} + \varepsilon_{rt+1}. \end{aligned} \tag{28}$$

In sharp contrast to the significant return effect of $\Delta\mu_{dt+1}$, we fail to reject the hypothesis that $\lambda_\varepsilon = \lambda_h = \lambda_J = 0$ and we also find that the estimated return effect of these additional terms is small. Hence, temporary shocks to the dividend growth process as well as movements in dividend growth volatility or jumps in dividend growth fail to have a significant effect on mean returns.[36]

These results suggest that news about the persistent dividend growth rate affects the first moment (mean) of same-day stock returns, while news about the temporary components does not. However, dividend news need not be limited to affecting mean returns and could also influence the volatility of stock returns on a given day: we might expect positive news about persistent dividend growth to be associated with calmer markets, as such news reduces the likelihood of being in a low-growth ("bad fundamentals") state. Movements in the volatility of dividend growth may also spill over to return volatility.

To explore these possibilities, we estimate the following auxiliary regression, which uses the squared daily return residual from (25) as a proxy for return variance:

$$\varepsilon_{rt+1}^2 = a_0 + a_1\varepsilon_{rt}^2 + a_2\sigma_{dt+1}^2 + a_3\Delta\mu_{dt+1} + \varepsilon_{\sigma t+1}, \tag{29}$$

where $\sigma_{dt+1}^2 = \exp(h_{dt+1})$ is the stochastic variance series extracted from the daily dividend growth process, Δd_{t+1}. We include the lagged squared return residual, ε_{rt}^2, to account for persistence in the return variance process.

Panel C of Table IX reports results from this regression. The significantly positive coefficient on prior-day squared return shocks, ε_{rt}^2, picks up persistence in the daily return variance. The variance of the dividend growth process, σ_{dt+1}^2, is strongly positively correlated with return variance, suggesting that greater uncertainty about dividend growth translates into higher same-day return volatility.[37] Consistent with good news about the persistent dividend growth component being associated with calmer markets, the coefficient on $\Delta\mu_{dt+1}$ is negative, although with a p-value of 0.12, it fails to be statistically significant.[38]

Daily dividend news may also affect the probability of observing jumps in stock returns. To investigate this possibility, we go through our sample to flag days with unusually large price movements and, following Bandi and Renò (2016), introduce a jump indicator that equals 1 on days when the absolute

[36] In Internet Appendix Section III, we report results of an analysis that uses daily dividend distributions rather than dividend announcements. While dividend announcements are significantly positively correlated with same-day stock returns, we find no such effect for dividend distributions.

[37] Using textual analysis, Boudoukh et al. (2019) find that, consistent with our results, public information about firm-level news is important for explaining the variance of stock returns.

[38] As we show in Internet Appendix Section III, μ_{dt} also contains significant predictive power for volatility measures based on high-frequency intradaily price data and option prices (VIX).

value of the (standardized) residuals from the return equation exceeds two. We then estimate the jump test model[39]

$$I_{t+1} = b_0 + b_1 \xi_{dt+1} J_{dt+1} + b_2 \sigma^2_{dt+1} + b_3 \Delta\mu_{dt+1} + \varepsilon^I_{t+1}.$$

Panel D of Table IX shows that positive news about the persistent cash flow growth component, $\Delta\mu_{dt+1}$, is associated with a reduced probability of jumps in returns with a coefficient that is highly significant. Positive jumps in the dividend growth process tend to be associated with a slightly reduced likelihood of observing a jump in stock returns on the same day, although the effect is only marginally significant at the 10% level. Conversely, greater uncertainty about cash flow growth (a larger σ^2_{dt+1}) is associated with a highly significant increase in the likelihood of a jump in stock returns on the same day. Greater uncertainty about cash flow growth thus translates into a larger chance of large movements in stock prices and our dividend growth components help identify cash flow news as a source of jumps in daily stock returns.[40]

VI. Conclusion

This paper develops a new methodology for constructing a daily "bottom-up" measure of aggregate cash flow growth based on firms' dividend announcements. In constructing this measure, we address two key challenges. First, individual firms' announced dividends can change by large amounts from one quarter to the next and display strong heterogeneity across firms. Second, the number and type of firms that announce dividends often changes substantially from day to day, leading to large composition shifts. Both effects lead to lumpiness in the daily cash flow news process.

We address this lumpiness by decomposing news on dividend growth into a transitory "normal" shock whose volatility can vary over time, jumps that occur more rarely but whose magnitude tends to be much larger, and a persistent, smoothly evolving component that captures long-run predictive dynamics in the mean of the cash flow growth process. We find that these components are well identified in the daily dividend growth data. Importantly, the persistent mean component captures predictable dynamics in dividend growth that are easily overlooked in the raw dividend growth series, which is dominated by the highly volatile temporary jump component. We show empirically that this persistent

[39] Estimates of a probit model linking this jump indicator to the three components extracted from our daily dividend growth model, that is, jumps, $\xi_{dt+1} J_{dt+1}$, stochastic variance, σ^2_{dt+1}, and news about persistent dividend growth, $\Delta\mu_{dt+1}$,

$$\Pr\left(I_{t+1} \middle| \xi_{dt+1}, J_{dt+1}, \sigma^2_{dt+1}, \Delta\mu_{dt+1}\right) = \Phi\left(b_0 + b_1 \xi_{dt+1} J_{dt+1} + b_2 \sigma^2_{dt+1} + b_3 \Delta\mu_{dt+1}\right),$$

are very similar.

[40] Our findings are related to an extensive finance literature that finds evidence of time-varying jump risk in stock returns using data on high-frequency intraday returns or out-of-the-money options; see, for example, Bollerslev and Todorov (2011). Similarly, Kelly and Jiang (2014) introduce a new measure of time-varying tail risk extracted from the cross-section of stock returns that they find has strong predictive power for returns on individual stocks and on market returns.

dividend growth component can be used to produce more accurate forecasts of future dividend growth than alternative approaches from the finance literature.

While our empirical analysis uses high-frequency (daily) dividend announcements, it also offers new insights into the determinants of stock price dynamics at longer horizons. Indeed, shocks to the persistent dividend growth component identified by our model are long-lasting—with a half-life exceeding a year—and lead measures of economic activity traditionally used as proxies for "fundamentals" such as GDP and consumption growth. Moreover, our estimates of persistent dividend growth can be used to produce more accurate forecasts of dividend growth than existing alternatives at both monthly and annual horizons.

Finally, we use a present value model to explore implications for asset pricing and return predictability associated with our results on dividend growth predictability. We find strong evidence that the persistent dividend growth component can be used to forecast stock market returns at both short horizons such as a single day and longer horizons such as a month. Moreover, the cross-equation restrictions implied by the present value model are supported by the data. Positive news about the persistent dividend growth component is associated with higher mean stock returns, reduced stock market volatility, and a lower likelihood of a jump in stock returns on the same day. Finally, we find that greater uncertainty about cash flow growth translates into larger volatility for stock prices, and we identify news about our dividend growth components as a determinant of jumps in daily stock returns.

Initial submission: September 17, 2018; Accepted: October 28, 2019
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

Boragan Aruoba, S., Francis X. Diebold, and Chiara Scotti, 2009, Real-time measurement of business conditions, *Journal of Business and Economic Statistics* 27, 417–427.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, 2016, Measuring economic policy uncertainty, *Quarterly Journal of Economics* 131, 1593–1636.

Bandi, Federico M., and Roberto Renò, 2016, Price and volatility co-jumps, *Journal of Financial Economics* 119, 107–146.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears, and risk premia, *Journal of Finance* 66, 2165–2211.

Boudoukh, Jacob, Ronen Feldman, Shimon Kogan, and Matthew Richardson, 2019, Information, trading, and volatility: Evidence from firm-specific news, *Review of Financial Studies* 32, 992–1033.

Boyd, John H., Jian Hu, and Ravi Jagannathan, 2005, The stock market's reaction to unemployment news: Why bad news is usually good for stocks, *Journal of Finance* 60, 649–672.

Brav, Alon, John R. Graham, Campbell R. Harvey, and Roni Michaely, 2005, Payout policy in the 21st century, *Journal of Financial Economics* 77, 483–527.

Campbell, John Y., Stefano Giglio, and Christopher Polk, 2013, Hard times, *Review of Asset Pricing Studies* 3, 95–132.

Campbell, John Y., and Robert J. Shiller, 1988a, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Campbell, John Y., and Robert J. Shiller, 1988b, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.

Cavallo, Alberto, and Roberto Rigobon, 2016, The billion prices project: Using online data for measurement and research, *Journal of Economic Perspectives* 31, 151–178.

Chib, Siddhartha, Federico Nardari, and Neil Shephard, 2002, Markov chain Monte Carlo methods for stochastic volatility models, *Journal of Econometrics* 108, 281–316.

Cochrane, John H., 1992, Explaining the variance of price-dividend ratios, *Review of Financial Studies* 5, 243–280.

Cochrane, John H., 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533–1575.

Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2009, The price is (almost) right, *Journal of Finance* 64, 2739–2782.

Cutler, David M., James M. Poterba, and Lawrence H. Summers, 1989, What moves stock prices?, *Journal of Portfolio Management* 15, 4–12.

Dangl, Thomas, and Michael Halling, 2012, Predictive regressions with time-varying coefficients, *Journal of Financial Economics* 106, 157–181.

Glen Donaldson, R., and Mark Kamstra, 1996, A new dividend forecasting procedure that rejects bubbles in asset prices: The case of 1929's stock crash, *Review of Financial Studies* 9, 333–383.

Eraker, Bjorn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in volatility and returns, *Journal of Finance* 58, 1269–1300.

Ferreira, Miguel A., and Pedro Santa-Clara, 2011, Forecasting stock market returns: The sum of the parts is more than the whole, *Journal of Financial Economics* 100, 514–537.

Gilchrist, Simon, and Egon Zakrajsek, 2012, Credit spreads and business cycle fluctuations, *American Economic Review* 102, 1692–1720.

Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.

Hansen, Lars Peter, and Ravi Jagannathan, 1997, Assessing specification errors in stochastic discount factor models, *Journal of Finance* 52, 557–590.

Hu, Grace Xing, Jun Pan, and Jiang Wang, 2013, Noise as information for illiquidity, *Journal of Finance* 68, 2341–2382.

Jagannathan, Ravi, and Binying Liu, 2019, Dividend dynamics, learning, and expected stock index returns, *Journal of Finance* 74, 401–448.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng, 2015, Measuring uncertainty, *American Economic Review* 105, 1177–1216.

Kelly, Bryan, and Hao Jiang, 2014, Tail risk and asset prices, *Review of Financial Studies* 27, 2841–2871.

Kelly, Bryan, and Seth Pruitt, 2013, Market expectations in the cross-section of present values, *Journal of Finance* 68, 1721–1756.

Koijen, Ralph S.J., and Stijn van Nieuwerburgh, 2011, Predictability of returns and cash flows, *Annual Review of Financial Economics* 3, 467–491.

Lettau, Martin, and Sydney Ludvigson, 2005, Expected returns and expected dividend growth, *Journal of Financial Economics* 76, 583–626.

Lettau, Martin, and Stijn van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607–1652.

Liew, Jimmy, and Maria Vassalou, 2000, Can book-to-market, size and momentum be risk factors that predict economic growth?, *Journal of Financial Economics* 57, 221–245.

Maio, Paulo, and Pedro Santa-Clara, 2015, Dividend yields, dividend growth, and return predictability in the cross-section of stocks, *Journal of Financial and Quantitative Analysis* 50, 33–60.

McQueen, Grant, and V. Vance Roley, 1993, Stock prices, news, and business conditions, *Review of Financial Studies* 6, 683–707.

Patton, Andrew J., and Michela Verardo, 2012, Does beta move with news? Firm-specific information flows and learning about profitability, *Review of Financial Studies* 25, 2789–2839.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2010, Out-of-sample equity premium prediction: Combination forecasts and links to the real economy, *Review of Financial Studies* 23, 821–862.
Sabbatucci, Riccardo, 2017, Are dividends and stock returns predictable? New evidence using M&A cash flows, Working Paper, Stockholm School of Economics.
Savor, Pavel, and Mungo Wilson, 2016, Earnings announcements and systematic risk, *Journal of Finance* 71, 83–138.
van Binsbergen, Jules H., Michael Brandt, and Ralph S. J. Koijen, 2012, On the timing and pricing of dividends, *American Economic Review* 102, 1596–1618.
van Binsbergen, Jules H., and Ralph S. J. Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance* 65, 1439–1471.
van Binsbergen, Jules H., and Ralph S. J. Koijen, 2017, The term structure of returns: Facts and theory, *Journal of Financial Economics* 124, 1–21.
Vuolteenaho, Tuomo, 2002, What drives firm-level stock returns?, *Journal of Finance* 57, 233–264.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication code.