Dividend-Price Ratios and Payout Constraints*

Ivo Welch
UCLA

Amit Goyal
University of Lausanne

August 21, 2025

Abstract

Common models of stock returns and dividend growth have ignored the constraint that dividend payments can never exceed stock prices. If a statistical process finds itself in such a scenario, it is best interpreted as investment termination. In the simplest canonical model, in which stock returns net of dividend growth rates are matched to their historical means and standard deviations, the termination scenario is triggered between 65% and 80% of the time. Termination also has implications regarding the interpretation of predictive OLS coefficients of dividend-price ratios on stock returns. A second important contribution of our paper is to show multi-stock evidence that hints that *very* high dividend-price ratios predicted future dividend reductions rather than higher stock returns, contrary to the association within the narrower dividend-price ratio range observed for the overall market.

*We thank Mark Garmaise, Valentin Haddad, and seminar participants at USC, UCLA, Amsterdam, and Frankfurt for helpful comments and advice.

I Introduction

The dividend-price analysis in Campbell and Shiller (1988) has become foundational for many modern asset-pricing courses. It would seem odd that three decades later, there would still be novel aspects of exactly this canonical relationship that have not yet been considered in the literature. Yet, our paper seeks to do precisely that. It points attention to a constraint that has been ignored so far: the fact that dividends should not exceed prices. The reasonable economic interpretation of a scenario in which a hypothesized process suggests that a stock would pay out more in dividends than it is worth would be investment termination (that is, the "demise" of existing investors). Of course, in real life, in the CRSP US sample, there was no termination. Dividends never reached above the stock price. The highest dividend-price ratio was only 7.2% (though dividend-price ratios of 6%+ occurred four different times: in 1932, 1940, 1950, 1980). This raises the question of why we would imagine that the limit on dividend payments is worth contemplating.

The answer lies in the dynamics of the standard process. Ignore our payout constraint for a moment and just focus on iterating the standard canonical process. Recall that the average stock return (r) in the CRSP sample was about 10%, with a standard deviation of 20%; the average dividend-growth rate (g) was about 5%, with a standard deviation of 15%; and the average difference of these two ($r-g$) was about 5%, with a standard deviation of 15%, too. Now consider a scenario that starts with a 7.2% dividend-price ratio and has a following realization of r and g at their respective means of 10% and 5%, respectively. At a \$100 price with an \$7.20 dividend, the next dividend would turn into \$7.56, the \$100 would turn into a \$110 return, of which \$102.44 would still be trading after the dividend payment. The dividend-price ratio would thus increase from 7.2% to $\$7.56/\$102.44 \approx 7.4\%$. If there is another draw at the means of r and g, the next dividend-ratio would be higher yet, and so on. In the standard investigated historical sample, with empirically calibrated moments, the basic dynamics should have pushed the dividend-price ratio upwards more often than not. It is easy to show that with higher dp ratios, the upward bias accelerates. Once the dividend-price ratio exceeds 10-15%, even quite positive $r-g$ realizations are no longer enough to prevent a rapid further upward slide. (For more intuition, note that a stock with a 20% dividend-price ratio

that experiences a not-rare $r-g$ of 0% would slide to a 25% dividend-price ratio the following year.) Of course, this is not our stipulation but merely the meaning of the standard identity.

Simple simulations with appropriately calibrated randomness suggest that, starting at 1926's dividend-price ratio of 5.3%, given an assumption of iid draws and the historical $r-g$ mean of 4.8% and standard deviation of 14.6%, researchers should have expected termination at $D \geq P$ between **65%** and **80%** of the time. (The exact number depends on assumptions about selection bias, which we will explain more clearly below.) We note also that the termination puzzle is primarily historic. It still matters to our understanding of the processes that drove the last 100 years of dividend and price data, which much finance and economics research relies on. It is less of an issue in the U.S. today, because current dividend-price ratios are hovering about 1.5% per year.

The lack of a clear analysis and awareness in the academic literature about the process when dividend-price ratios approach and exceed 1.0, in turn, muddles the standard dividend-price ratio testing framework. What exactly is the model that is being tested, and is it conditional on non-termination? What exactly is the null hypothesis, then? Can we interpret the slopes in predictive OLS regressions while ignoring the intercepts, considering the latter merely nuisance parameters? (They do influence the observable dividend-price ratio associations.)

Our paper offers some initial answers to these questions. If a process is to have avoided termination, then whenever the dividend-price ratio happened to be high, the subsequent $r-g$ would have been more likely to come in high, too. This requirement does not require any empirical data or analysis, but merely stipulation that we are looking at a 99-year time series in which investors did not terminate.

Note also that our termination constraint is not related to the familiar hypotheses (and disagreements) about whether the dividend-price ratio is stationary or non-stationary (e.g., Campbell and Yogo (2006), Torous, Valkanov, and Yan (2004), and Lewellen (2004)). The termination constraint is relevant regardless of whether the dividend-price ratio is stationary or not.

The non-termination concern suggests exploring one of two analyses:

1. In the first approach, our framework maintains the literature's traditional underlying process. We posit that the log dividend-price ratio (dp) forecasts r, g, and itself (dp_{+1}) in a simple linear fashion over the entire domain. This comes with the familiar hypothesis tests. The most interesting null hypothesis is about whether dp predicts the future r.

   However, our paper also considers that this approach is being tested in the context of a "lucky time-series draw," in which existing investors happened not to have terminated. For consistency, it is reasonable to go a little further and recognize that we are working with dp ratios that ranged historically from 1% to 7%. We can then explore the predictive coefficients under those "lucky ex-post realization" outcomes that match the observed empirical range. Doing so imposes very complex, unintuitive, and path-dependent constraints on stock returns and dividend growths, which are necessary to exclude "unlucky" terminated paths.

   Our paper shows that the empirical predictive OLS coefficient of dividend-price ratios on stock returns was, if anything, *less* than one would have expected if $r-g$ were truly unpredictable *and* the dp ratio matched the empirical range of 1% to 7%. The OLS coefficient predicting stock returns was 0.076 in the data, but should have been expected to be 0.086 under a null hypothesis of randomly drawn unpredictable $r-g$, *given* the empirically observed dp range.[1]

2. In the second approach, we posit a more general process. In particular, we consider that the true underlying process may not only have a uniformly linear association between dp and $r-g$, but a *level-dependent* association. Requisite additional process parameters can determine not only the predictive coefficients, but also the probability of non-termination. For example, such a process could posit that the mean $r-g$ is twice as high whenever the D/P ratio exceeds 10%. With suitably chosen additional parameters in the wider range of D/P ratios never observed, such a process can greatly reduce the probability of termination *and* account for observed empirical patterns in the narrower range of small D/P ratios.

[1] Without the lower bound matched to the empirics, still in the absence of predictability, the coefficient would have been expected to be "only" 0.041. The empirical [non-significant] coefficient of 0.076 was, of course, higher, but the simulations still imply that more than half of the empirical coefficient was just bias.

Unfortunately, the extra degrees of freedom in this second approach raise other thorny questions: What probability of termination should a plausible process tolerate? Although 65-80% termination probability seems too high to be aesthetically pleasing, would 10-20% be reasonable? Or 1-2%? Furthermore, what should be assumed of $r-g$ at high levels of D/P never observed in the historical data? And what should be the null hypothesis: Is it a zero predictive OLS coefficient of D/P for future stock returns at lower dividend-price levels or at higher dividend-price levels? Note that if a higher mean $r-g$ is assumed at high dividend price levels, then we have assumed a negative association between D/P and future $r-g$. Reducing the probability of termination is thus assuming by fiat that the alternative hypothesis is correct.

Nevertheless, despite these difficulties, our paper can explore this second potential resolution to the termination puzzle. Although the stock market overall never had a dividend-price ratio much above 7%, many individual stocks did. The individual stock data shows that high dividend-price ratios predicted strong dividend declines rather than high stock returns. This inference is opposite to the (naïve) interpretation that the predictive coefficient of aggregate dividend-price ratios on aggregate stock market returns was positive, and the predictive coefficient on aggregate dividend growth rates was zero. Our inference is also robust with respect to assumptions about the treatment of firms paying no dividends and firms dropping off the database.

Our paper is intentionally brief to focus the reader's attention on the key insights. Section II reviews the central identity and introduces an exact (rather than approximate) version of the Campbell-Shiller identity. Section III describes the data. Section IV explains the dynamics of the dividend-price ratio and the termination constraint. Section V introduces the "Goldilocks" model in which U.S. market investors were lucky not to terminate, and we observe the conditional data. Section VI explores the identity in the sample of individual stocks to gain insights into behavior when dividend-price ratios exceed 7%. Section VII reviews the literature, including our agreements and disagreements with Cochrane (2008). Section VIII concludes.

II Algebraic Identities and Economic Hypotheses

A Exact and Approximate Identities

It is useful to begin by illuminating the meaning of the updating identities and the corresponding predictive equations. Consider a simple ratio updating equation,

$$\frac{a}{b} \;=\; \frac{c\cdot(a/c)}{d\cdot(b/d)} \quad\Longleftrightarrow\quad \log\left(\frac{a}{b}\right) \;=\; \log\left(\frac{c}{d}\right)+\log\left(\frac{a}{c}\right)-\log\left(\frac{b}{d}\right) \quad .$$

This equation contains no meaning. Its algebra applies anywhere, e.g., if a is the number of stars in the Andromeda galaxy, b the number of bees in Oregon, c the consumption of apples in India, and d the (Roman) number 500. It is not so much a meaningful equation like $E = m\cdot c^2$, as it is like $m\cdot c^2 = m\cdot c^2$. Of course, the equation can also hold for economic quantities, although this does not imply any economic implications either. If a is the wealth of one gambler and b the wealth of another gambler (possibly but not necessarily playing against the first), then the first two terms are what is required to update wealth from one period to the next, given the latter two terms, which are the performances of the two players. In particular, this identity does not contain any implications about whether the gamblers' wealth or wealth ratios today can help forecast their winnings tomorrow.

An equivalent but asymmetric algebraic identity is

$$\begin{aligned} -\log\left(1+\frac{1}{a/b}\right) \;&=\; \log\left(\frac{a}{b}\right)-\log\left(1+\frac{a}{b}\right) \\ &=\; \log\left[\frac{c}{d}\right]+\log\left[\frac{a}{c}\right]-\log\left[\frac{b+a}{d}\right] \quad . \end{aligned} \tag{1}$$

Like the preceding equation, (1) has no meaning, applying to any quantities. We can use a more familiar set of variable names, $a \equiv D_{1,2}$, $c \equiv D_{0,1}$, $d \equiv P_1$, and $b \equiv P_2$, in which case, (1) reads

$$\begin{aligned} -\log\left(1+\frac{1}{D_{1,2}/P_2}\right) \;&=\; \log\left(\frac{D_{1,2}}{P_2}\right)-\log\left(1+\frac{D_{1,2}}{P_2}\right) \\ &=\; \log\left(\frac{D_{0,1}}{P_1}\right)+\underbrace{\log\left(\frac{D_{1,2}}{D_{0,1}}\right)}_{\text{dividend growth}}\underbrace{-\log\left(\frac{P_2+D_{1,2}}{P_1}\right)}_{\text{stock return}} \quad . \end{aligned} \tag{2}$$

This is the mechanical updating equation for dividend-price ratios, where the latter two terms are the familiar dividend growth rate and stock return terms, respectively. (Table 1 provides a glossary of the notation used in our paper.)

[Insert Table 1 here: **Notation Glossary**]

It is conventional to abbreviate $dp_{0,1} \equiv \log\left[D_{0,1}/P_1\right]$, $dp_{1,2} \equiv \log\left[D_{1,2}/P_2\right]$, $g_{1,2} \equiv \log(D_{1,2}/D_{0,1})$, and $r_{1,2} \equiv \log\left[(P_2 + D_{1,2})/P_1\right]$ in the dividend literature. We also add $\eta(\, dp_{1,2}\,) \equiv \log(1 + D_{1,2}/P_2)$, so we can rewrite (2) as

$$dp_{0,1} \;=\; dp_{1,2} \;+\; r_{1,2} \;-\; g_{1,2} \;-\; \eta(\, dp_{1,2}\,) \;\;.$$

Campbell and Shiller (1988) derive a linear Taylor approximation of this identity,[2]

$$dp_{0,1} \;\approx\; \psi \cdot dp_{1,2} \;+\; r_{1,2} \;-\; g_{1,2} \;+\; \omega(\,\psi\,) \qquad .$$

Of course, there is still no theory, hypothesis, or economics in any of the above equations. They are just algebra that applies to any arbitrary variables. They are useful for mechanical process updating and for context discussions below.

The equations can be rearranged to "suggest" a predictive association,

$$\begin{aligned} r_{1,2} - g_{1,2} \;&=\; (dp_{1,2} - dp_{0,1}) \qquad\quad + \;\; \eta(\, dp_{1,2}\,) \\ r_{1,2} - g_{1,2} \;&\approx\; (dp_{1,2} - dp_{0,1}) \qquad\quad - \;\; (1-\psi) \cdot dp_{1,2} - \omega(\,\psi\,) + \epsilon(\, dp_{0,1}, dp_{1,2}, r_{1,2}, g_{1,2}\,) \quad , \end{aligned}$$

where the final terms are typically quite small. Using the latter equation as guidance, Campbell and Shiller (1988) and Cochrane (2008) explore the association of stock returns, dividend growth rates, and dividend-price ratio ($dp_{0,1}$).

B Forecasting Coefficients

Both Goyal and Welch (2003) and Cochrane (2008) introduce a convenient rewrite of the C-S approximation in terms of predictive OLS regression coefficients.[3] The identity

[2] The estimates of their two constants are $\psi \approx 1/(1 + e^{\overline{\log D/P}}) \approx 1/(1 + e^{-3.39}) \approx 0.9675$, and $\omega = -\log(\psi) - (1-\psi)\cdot\log(1/\psi - 1) \approx 0.144$, updated as of 2024.

[3] We have not found this decomposition in earlier work. If we are wrong, please bring the reference to our attention.

can be expressed in terms of OLS betas of bivariate regressions in which the independent variable is always the prevailing dividend-price ratio $dp_{0,1}$. This is accomplished by taking the covariance of the identity with respect to $dp_{0,1}$ and then dividing by the variance of $dp_{0,1}$. The C-S approximation therefore implies (with a 1-period shift in t)

$$b^{ols}(r_1,\ dp_0)\ -\ b^{ols}(g_1,\ dp_0)\ +\ b^{ols}(\psi{\cdot}dp_1,\ dp_0)\ \approx\ 1\ ,$$

where the first subscript on b is the dependent variable and the second subscript is the independent variable, always dp_0. In our current paper, we use the exact identity instead, which implies that

$$b^{ols}(r_1,\ dp_0)\ -\ b^{ols}(g_1,\ dp_0)\ +\ b^{ols}(dp'_1,\ dp_0)\ =\ 1\ .$$

where $dp'_1 \equiv dp'_{0,1} \equiv -\log\left[1+1/(D_{0,1}/P_1)\right]$. The first two predictive OLS coefficients are the same; the third is very similar. Because the same algebra holds for arbitrary time indexes,

$$\lim_{t\to\infty}\ b^{ols}(r_t,\ dp_0)\ -\ b^{ols}(g_t,\ dp_0)\ =\ 1\ -\ b^{ols}(dp'_t,\ dp_0)\ .$$

So far, there is still no economics or testable implication. The above equation holds for arbitrary variables.

C Economics and Potential Termination Biases

It is here that economics finally makes its entrance. Campbell and Shiller (1988) propose their (only) testable implication, which is that dividend-price ratios should be stationary:[4]

$$dp \text{ Stationarity}\ \Rightarrow\ \lim_{t\to\infty} b^{ols}(dp'_t,\ dp_0)\ =\ 0\ \ \Rightarrow\ \ \lim_{t\to\infty} b^{ols}(r_t,\ dp_0)\ -\ b^{ols}(g_t,\ dp_0)\ =\ 1\ .$$

Under the stationarity hypothesis, in the limit as $t\ \to\ \infty$, it is not possible for both $b^{ols}(r_t,\ dp_0)$ and $b^{ols}(g_t,\ dp_0)$ to be zero. (Furthermore, a higher $b^{ols}(r_t,\ dp_0)$ also has to associate with a higher $b^{ols}(g_t,\ dp_0)$.) Campbell-Shiller then explores a more ad hoc

[4] We have rewritten their equation in terms of the exact identity instead of their approximation (i.e., we used $b^{ols}(dp'_1,\ dp_0)$ instead of $b^{ols}(\psi{\cdot}dp_1,\ dp_0)$).

but also more testable analog hypothesis on finite horizons. When they posit that $b^{ols}(\,dp'_1,\,dp_0\,) \neq 1$ in

$$b^{ols}(\,r_1,\,dp_0\,) \;-\; b^{ols}(\,g_1,\,dp_0\,) \;=\; 1 - b^{ols}(\,dp'_1,\,dp_0\,)\;,$$

they can test the prediction that $b^{ols}(\,r_1,\,dp_0\,)$ and $b^{ols}(\,g_1,\,dp_0\,)$ cannot both be zero.

Furthermore, it is also clear that the more specific $b^{ols}(\,dp'_1,\,dp_0\,)$ is, the tighter must be the linkage between $b^{ols}(\,r_1,\,dp_0\,)$ and $b^{ols}(\,g_1,\,dp_0\,)$. Goyal and Welch (2003, Table 6) use empirically estimated sample $b^{ols}(\,dp'_1,\,dp_0\,)$'s to instrument the sample association between $b^{ols}(\,r_1,\,dp_0\,)$ and $b^{ols}(\,g_1,\,dp_0\,)$. Over longer horizons, the sample $b^{ols}(\,r_1,\,dp_0\,)$ increased. Cochrane (2008) primarily posits two specific values, $b^{ols}(\,dp_1,\,dp_0\,)$=0.941 (p.1541), and $b^{ols}(\,dp_1,\,dp_0\,)$=0.99 (p.1543,1553).[5]

Of course, the identity has to hold in every valid sample. However, there is a different concern now. It is that $b^{ols}(\,r_1,\,dp_0\,) - b^{ols}(\,g_1,\,dp_0\,)$ can be larger when conditioned on samples that did not terminate. The degree to which the selection bias matters depends on the true underlying mean (and path-dependent configuration) of $r-g$, and not just the slopes. For example, given the 1926 starting value of dp (equivalent to a D/P ratio of 5.3%), if the mean of $r-g$ had been 20% per year, then the termination constraint would probably never have been binding, and the researcher would have observed the full range of $b^{ols}(\,r_1,\,dp_0\,) - b^{ols}(\,g_1,\,dp_0\,)$. The mean $r-g$ and mean intercepts in the predictive regressions would have been just nuisance parameters. The low probabilities of termination would have justified neglecting selectively biased draws (lack of termination, e.g., as assumed in effect in Cochrane (2008)). In contrast, if the mean of $r-g$ had been 1% per year, then the termination constraint would have almost always been binding. The only possible series still observable would have been those in which returns were high, and especially high dp's would have happened to have been followed by high $r-g$'s to avert termination. (Strong mean reversion in dp would have been necessary to avoid termination.) In such cases, the mean $r-g$ and mean intercepts in the predictive regressions would no longer appropriately be considered merely nuisance parameters that could not have biased the empirical data sample.

[5]Neither of these two papers fully estimated a population $b^{ols}(\,r_1,\,dp_0\,)$, corrected for bias. We will discuss these interpretations further in our literature review in Section VII.

D Alternative Fundings and Aggregate Constraints

It is not necessarily the case that investor termination means market termination. In real life, it is possible (and indeed common) for investors to infuse capital into the stock market via new issuances. This does not invalidate the Campbell-Shiller identity, but it changes the $D \leq P$ constraint interpretation. Existing investors would have suffered a complete loss, but new investors could have, in effect, created in all-but-name a new stock market. For illustration, in the extreme scenario, imagine that existing firms could pay out their entire value as dividends (perhaps because they experienced terrible stock returns), which would thus end their investors' participation in the market. However, simultaneously, new investors could have purchased 100% of newly issued shares in the empty shells and thereby recapitalized them. Of course, this would not have helped previous investors avoid the termination constraint that their shares should never be worth more than the dividends being paid.

This raises the question of whether it is possible to integrate net payouts (based not only on payouts but also payins) into a Campbell-Shiller-like framework. Unfortunately, the answer is no. The equivalents of the stock return (the percent change in the market capitalization) and the dividend-price ratio (the payout ratio) are straightforward enough, and there are indeed many years in which the net payout ratio was large enough to raise similar exhaustion concerns as the dividend-price ratio. However, the equivalent of dividend growth becomes nonsensical because net payouts can be negative in years in which firms issue more than they pay out.

III Data and Summary Statistics

Our paper reanalyzes the familiar association between annual dividends and annual prices in the CRSP database. Our prices, dividends, and stock returns originate from the monthly CRSP file (version 2) in March 2025, resulting in coverage from 1926 to 2024. Because stocks can have unusual dividends and stock splits, we follow CRSP's choices and work only with `MthRet` and `MthRetx` to identify dividend growth rates, stock returns, and dividend-price ratios. We first calculate the compounded annual returns with and without dividends ($r_t = -1 + \prod_{t=1}^{12}(1 + \texttt{MthRet}_t)$ and $x_t = -1 + \prod_{t=1}^{12}(1 + \texttt{MthRetx}_t)$),[6] and then define

$$\begin{aligned} \text{Dividend-Price Ratio:} \quad \frac{D_{t,t+1}}{P_{t+1}} &\equiv \frac{1 + r_{t,t+1}}{1 + x_{t,t+1}} - 1 \ , \\ \text{Dividend Growth:} \quad \frac{D_{t,t+1}}{D_{t-1,t}} &\equiv \frac{(1 + x_{t-1,t}) \cdot (r_{t,t+1} - x_{t,t+1})}{r_{t-1,t} - x_{t-1,t}} - 1 \ . \end{aligned} \tag{3}$$

We always use the full CRSP sample unless otherwise indicated, with one year lost due to the definition of the dividend growth rate and one year lost due to the lead-lag nature of the prediction.

[Insert Table 2 here: **Summary Statistics for the Value-Weighted CRSP Market Portfolio, 1926-2024**]

Table 2 shows the descriptive statistics for the source input stock returns, specifically for the value-weighted stock market portfolio (Panel A), and for the derived key variables used in our paper. Panel B shows that at a one-year horizon, the log of the dividend-price ratio ($dp_{0,1}$) was highly autocorrelated at 0.94.[7] Panel C shows that stock returns were, on average, twice as high as dividend growth rates. Both had low autocorrelation but strong cross-correlation (66%). There is a natural linkage between contemporaneous

[6] Note that these are not true annual buy-and-hold portfolios, but entail some monthly rebalancing. Furthermore, note that the easiest way to work with other frequencies (e.g., multi-year identities) would be to aggregate returns with and without dividends first into the desired frequency and then proceed as if there were only the two periods.

[7] Its first difference was well-behaved, with an autocorrelation of about 8% at its first lag. As one could have guessed from the autocorrelations and unit-root tests offered by Campbell and Shiller (1988), Goyal and Welch (2003), Cochrane (2008), Boudoukh, Richardson, and Whitelaw (2008), and many others, the dividend-price ratio data have strongly pointed to a process that seems more like a random walk than one that seems stationary.

r and g.[8] The $(r-g)_{1,2}$ series (which by definition is also $dp'_{1,2} - dp_{0,1}$) has no strong autocorrelation pattern. Panel D shows that standard tests suggest that the log dividend-price ratio was not stationary using standard tests assuming normal distributions. Of course, economics tells us that these tests could only be halfway reasonably close to valid if the drift was strongly negative. A positive drift would quickly run into the terminating termination constraint, and the tests were derived for unconstrained normally distributed process variables.

[Insert Figure 1 here: **Log D/P Ratio**]

Figure 1 plots the historical log dividend-price ratio. All ratios before 1994 were above 2.6%, and all ratios after 1994 were below. The dividend-price ratio reached local maxima above 5% four times: in 1931-32, in 1941, in 1949-50, and in 1978-82. The first dividend-price ratio in 1926 (and our starting point for our simulations later) was 5.3%, i.e., a dp of -2.93. Its last value in 2024 was 1.57% (-4.15), but only the 2023 value of 1.82% (-4.01) could be used to predict subsequent (2024) values.

[Insert Figure 2 here: **Log Rates of Return Net of Log Dividend Growth Rates ($r-g$)**]

Not shown, stock rates of return and dividend growth rates seem like well-behaved normal iid draws. Both rates fell over the sample period, leaving the difference with only the slightest of downward trends. This is shown in Figure 2, which plots the historical difference of log returns and log dividend growth rates. There were no strong secular trends. The $r-g$ mean seems to have been steady around 5%, perhaps modestly lower in recent years.

[8] Our null hypothesis draws r and g together, both independent of dp. One could go a different route. One could first draw an iid r, and then construct g as a linear combination of the contemporaneous r and the past dp, designed to match the correlation between r and g. This would allow for a less restrictive null hypothesis.

[Insert Figure 3 here: **Log D/P Ratio and Future Variables**]

Figure 3 shows the well-known empirical relation between lagged log dividend-price ratios and subsequent variables, as, e.g., in Goyal and Welch (2003) and Cochrane (2008), in "spider web" form. All values in the right halves of the panels occurred before 1994, all values in the left halves thereafter. A (loess) line fit, shown in red, suggests a positive relationship between dividend-price ratios and subsequent stock returns before 1994, an inverse U-shaped relationship thereafter, and a positive relationship over the entire range. The equivalent loess line for dividend growth shows no predictive relationship. The bottom left panel shows that the stock-return predictive association carried over to the difference between stock returns and dividend-growth rates. However, it must be emphasized that it is well known that, even in the absence of the termination constraint, the sample coefficient estimates here are not only noisy but *biased* reflections of equivalent underlying relationships. The bottom right panel shows that the subsequent dividend-price ratio did not have a strong mean reversion tendency. The red (loess) line fit lies nearly perfectly on the diagonal — except perhaps for the very left-most and very right-most dividend-price ratios.[9]

[Insert Table 3 here: **Predictive OLS Coefficients**]

Table 3 shows the empirical sample estimates for the predictive OLS coefficients. They suggest a positive (in-sample) association between prevailing dividend-price ratios and subsequent stock returns. Of course, these coefficients cannot be interpreted as unbiased estimates of an underlying process that has the same functional form. The independent variable is highly autocorrelated. (Stambaugh (1999) suggests a correction if the independent variable follows a non-terminating normal AR1 process.) The association between dividend-price ratios and subsequent stock returns is solid but not very strong, just as observed in the figure. Even the bias-uncorrected coefficient is "only" 0.076, with a T-statistic of 1.79 (i.e., not statistically significant at the conventional 5% level). The final two rows in Table 3 show that the approximation and the exact identity result in almost identical predictive OLS coefficients (of 0.922 and 0.923) *in the empirical*

[9] Fitting a second polynomial degree to the process results in $dp_{t+1} \approx -1.17 + 0.36 \cdot dp_t - 0.09 \cdot dp_t^2$. It indeed suggests a more S-shaped relation, but this is *not* statistically significant. It is merely a visual confirmation.

sample. (Not shown, this will no longer be the case for dividend-price ratio higher than empirically observed dp domain.)

IV Dynamic Instability

Sections II and III have mostly followed well-tread ground. Our own paper is related to the asymmetry in equation 1 between the role of stock returns and the role of dividend ratios. It is the fact that dividends are paid out and thereby increase the subsequent price-dividend ratio. For example, if a stock has a price of $100 and pays out dividends of $20, for a prevailing dividend-price ratio of 20%, and both stock returns and dividend growth are zero, then the next dividend-price ratio will creep up to 25% ($20/$80).

[Insert Figure 4 here: **Dynamics of Dividend-Price Ratios**]

Figure 4 illustrates the dynamics of a dividend-price process as a function of the prevailing dividend-price ratio and the subsequent stock return net of dividend growth. Rearranging (2) gives the line of stability, which is $y(D/P)=1+D/P$. This line divides the area into one region in which a specific draw of the difference increases the next ratio realization, and another region where it decreases it. For a high dividend-price ratio, even a modest positive difference between stock returns and dividend growth rates induces a quick acceleration towards higher and higher dividend-price ratios. The right panel shows that, given where the CRSP sample started out in 1926 (about 5.3%), the average net difference between stock returns and dividend growths of about 5.05% would have neither greatly increased nor decreased the next dividend ratio. It lies very close to the stable line.

Dividend-price ratios cannot exceed 1.0, which means that their logs cannot exceed 0.0. The exact updating equation breaks down algebraically, offering no solution.[10]

[10]The negative-price problem manifests itself differently depending on the updating formula used. In the Campbell-Shiller version, whenever dividend-price ratios exceed, say, 20%, the linearization "just" turns into a terrible approximation. (Again, this is not an issue in the empirical sample, which never reaches above 7.2%.) An underlying model based on such a CS linearization (e.g., via the common VAR formulation) then quietly and mistakenly soldiers on. Eventually, the dividend-price ratio may or may not step back into the physically viable domain. In contrast, in the exact identity, the algebra fails immediately and appropriately. It becomes impossible to invert $dp'_{1,2}$ to match a $dp_{1,2}$.

In such a case, the resulting price would turn negative. At this point, further stock returns would be impossible to interpret in economic terms. The only reasonable interpretation of such a scenario would be that existing investors' portfolios are liquidated because as much is paid out as was available. As noted, we call this investor "termination."

Of course, such termination did not happen in our U.S. sample. However, Jorion and Goetzmann (1999) documents that not only investor termination but outright demise did occur in other markets. Reflecting on the periods in the United States when dividend-price ratios were above 5%, as in 1932, 1941, 1950, and 1978, suggests that the long-term non-termination of U.S. market investors was, at those times, perhaps not as clear as it was after 1994 — and especially as it seems clear today, when dividend-price ratios are hovering around 1.5%.

The non-termination constraint (a selection bias) is not linear and thus not intuitive. It imposes complex conditions on valid combinations of stock returns, dividend growth, and dividend-price ratios. A 50% decrease in stock returns and 33% increase in dividend growth is perfectly valid if the prevailing dividend-price ratio is 0.1% *at that moment*, but not if it is 30%. Even the valid domain of stock returns and dividend growth rates is path-dependent, too. The same set of stock return and dividend growth pairs, when reordered, can violate the identity in one ordering but not another.

We must also stress again that the concern about the stationarity of D/P[11] is not the same as our concern that $D/P<1$. Our concern applies not just to explosive but also to stationary processes.

The natural immediate question already briefly discussed in the introduction is about why researchers should be concerned about a dividend-price ratio of 100% suggesting termination, when dividend-price ratios never exceeded 7%. After all, academic research also often works with normally distributed plain stock returns, although stock returns cannot exceed -100%. One difference is that the market could have sustained an infinite number of annual rate of return draws of four standard deviations below the mean. The

[11]The importance of that premise of the prior constraint about the dividend-price ratio stationarity has also been highlighted before. Campbell and Yogo (2006) write that "since statistical tests do not reject an explosive root for the dividend–price ratio, we have evidence for return predictability only if we are willing to rule out an explosive root based on prior knowledge. This reconciles the 'contradictory' findings by Torous, Valkanov, and Yan (2004, Table 3), who report that the dividend–price ratio does not predict monthly returns in the postwar sample, and Lewellen (2004, Table 2), who reports strong evidence for predictability."

next stock return draws would always be relative to the then-lower price. This is not the case for the dividend-price ratio. A few draws of $r-g$ just a few standard deviations below the mean can push the market far enough into a dynamic where even unusually large positive draws might no longer avoid the termination scenario.

[Insert Figure 5 here: **Markovian Process: Demise Outcomes**]

Figure 5 considers scenarios in which the difference between the log stock price and the log dividend growth rate ($r-g$) is drawn iid from a stable normal distribution (a process which seems empirically well), with a standard deviation of 14.6% per year (as in the data). Even when we assume that $r-g$ had its remarkably high mean of 4.8%, starting at a dividend-price ratio of 5.3% in 1926 would have resulted in termination about **65%** of the time. At its 7.2% dividend-price ratio apex in 1950, the stock market would have been on a 90% likely path toward termination over its 99-year interval. It is only very recently that dividend-price ratios have fallen so far that the probability of termination has declined into single percentage points (however, ignoring the fact that $r-g$ may also have fallen; see Figure 2). We conclude that for research using the historical 1926-2024 data, we cannot assume that historical means of $r-g$ under the null hypothesis were unimportant as far as selection bias is concerned.

Note also that the underlying process specification in both Stambaugh (1999) and Cochrane (2008) (as well as in standard stationarity tests) cannot be truly applicable here. These processes allow for realizations that are not inadmissible. Instead, we need to interpret the empirical moments in the context of some alternative feasible processes. Our paper explores the two aforementioned approaches:

1. We can posit (model) that we are living in a Goldilocks world, in which investors were not terminated. In this case, we want to explore a situation in which the observed statistics were conditional on non-termination, or more specifically, conditional on a range of D/P ratios that matched the empirically observed range.

2. We can posit (model) that we are living in a world in which the $r-g$ process was not iid. In particular, it could have had a non-linear association between dp and $r-g$ such that termination was less likely *a priori*, and thus our observation of non-termination was less surprising. However, an ex-ante prior imposition of such

> a process with a lower probability of termination effectively already implicitly states (or at least biases) that the null hypothesis does not hold. When, then, is a test of the null hypothesis?

Thus, the constraint raises the thorny question of exactly what the null hypothesis is. If the null hypothesis is presumed to be the case in which stock returns are independent of lagged dividend-price ratios over the entire domain, then we don't even need data: it is simply not algebraically possible. No test is necessary: As formulated, it's not an economic hypothesis but algebra. The dividend-price ratio could not have been a random walk, either, because it could have ended up in termination. We need to be more precise about what we mean by a hypothesis about an underlying process in which dividend-price ratios do not anticipate stock returns.

One reasonable economic null hypothesis could be that *conditional on non-termination* dividend-price ratios could not predict stock returns. (Of course, D/P ratios could then still predict termination.) Or, it could be that, *conditional on observing dividend-price ratios within the empirically observed range of dividend-price ratios (from 1.1% to 7.2%)*, dividend-price ratios could not predict stock returns. The latter null has the appealing property that it is grounded in empirically observable (matchable) quantities, and that it winsorizes scenarios not just at the top end of dividend-price ratios, but also at the bottom end. Again, both of these null hypotheses still impose non-linear restrictions between dividend-price ratios, stock returns, and dividend growth. If the imposition of such constraints is interpreted as part of the "predictive ability of the D/P ratio" hypothesis, i.e., as part of the alternative rather than the null hypothesis, then no empirical tests are required. The null hypothesis is algebraically false, and the alternative hypothesis must be true.

It is perhaps both simpler and more intuitive to argue less about hypotheses and more about process estimates. We turn there next.

V The Conventional Model in a Goldilocks Scenario

[Insert Table 4 here: **Five-Parameter Canonical Goldilocks Model Simulations**]

The first alternative specification is to stick with the simplest and most common model, in which stock returns and dividend growth are random normal draws, but also take into account the selection condition. This keeps our approach similar to Cochrane (2008), although we simulate stock returns and dividend growth rates (rather than the approximation-derived beta system).[12] We call it our "canonical core null model" — a data-generating process, with all quantities in logs and both epsilons (ϵ^g and ϵ^r) drawn from iid standard normal distributions. Consider the following process:

$$\begin{aligned} r_t &\approx 0.081 + 0.1944 \cdot \epsilon_t^r \\ g_t &\approx 0.048 + 0.1056 \cdot \epsilon_t^g + 0.0936 \cdot \epsilon_t^r \end{aligned}$$

The five numerical parameters of this process were chosen to match in "surviving" simulations the five core empirical statistical moments in the data that are analogous to the generating process: the means, and standard deviations, plus their contemporaneous correlation of both log stock returns and log dividend growth.[13] (The chosen coefficients in the underlying canonical model are *not* unbiased estimates of their empirical sample analogs.) In this process, stock returns are unpredictable using *any* lagged variables, including dividend-price ratios.

Panel A of Table 4 describes the averages of these five simulated moments. Being matched, the five parameters fit the five moments perfectly.

To take the model seriously means not to allow simulations to soldier on when the log dividend-price ratio exceeds 0.0. Instead, it is necessary to consider such cases to represent termination. In practice, this means ignoring model simulations in which termination would have occurred. Panel B shows that only 17% of the draws from this

[12] Cochrane's simulations are under a mean-zero assumption (dispensing with intercepts). Because of the positivity constraint on dividend-price ratios, the demeaning is not as innocuous as it would normally have been.

[13] It is of course impossible to test whether dividend-price ratios predict r in the presence of a large predictive association of gd *and* preserve the correlation between r and g. If dividend-price ratios have a high correlation with g, and g retains the high empirical correlation with r, then the dividend-price ratio will also have a high correlation with r.

model never stepped into the infeasible domain; the other 83% did. Thus, the canonical null model's most surprising new insight for us was the frequent termination of the stock market.

Note that the selection bias had already manifested itself in Panel A. The optimizer matched the empirical means, average standard deviations, and average correlation of log stock returns and log dividend growth rates, as designed, *conditional on non-termination*. Thus, the mean stock return of **9.5%** in the empirical data was higher than the model's **8.1%** expected mean stock return in the underlying process. This discrepancy is required to match conditional stock returns and dividend growth with simulated draws in the presence of the selection (feasibility) constraint. In turn, the lower model mean stock return increases the probability of termination.

Panel C shows the 100-year range of dividend-price ratios in the canonical null model conditional on non-termination. The model predicts (on average) a lowest log dividend-price ratio of -6.1 (0.2%) per year, much lower than the empirically observed lowest ratio of -4.5 (1.1%). The model predicts a highest log dividend-price ratio of -2.73 (6.5%), close to the -2.63 (7.2%) empirically observed ratio. At the end of the sample, the model would have predicted a final log D/P of -5.65 (0.35%), much lower than the empirically observed -4.01 (1.8%).

Panel D explores the predictive OLS coefficients, conditional on non-termination. The simulated mean predictive coefficient of D/P on stock returns, $b^{\text{ols}}(r_1,\ dp_0)$, is 0.041. This is about 0.8 standard deviations lower than the empirically observed $b^{\text{ols}}(r_1,\ dp_0)$ of 0.076. This empirical coefficient of 0.076 lies at the 84th percentile of the simulated distribution. The simulated mean predictive coefficient of D/P on dividend growth is zero, as predicted. Note that it should not be surprising that the coefficients for $b^{\text{ols}}(r_1,\ dp_0)$ and $b^{\text{ols}}(g_1,\ dp_0)$ are quite different. They were never symmetric in the defining identity in (1). The next two predictive coefficients are auto-coefficients, which have to be mirror images of the difference of the other two variables. Thus, we omit repeating the discussion. The final two predictive coefficients work with longer intervals. They confirm the insight of Boudoukh, Richardson, and Whitelaw (2008) that little statistical power is to be gained from them. Although the coefficients increase, this is merely an artifact of the specification and provides no extra power. (We discuss Cochrane's important suggestion about the joint hypothesis test on page VII.)

Panel E considers the use of the Stambaugh beta, using his approximation based on the Marriott and Pope (1954) and Kendall (1954) AR1 bias correction. The simulated data shows that it is unbiased, unlike the $b^{\text{ols}}(r_1, dp_0)$ coefficient with its 3.6% bias. It performs admirably, although it was not designed to work with this process (and of course was not previously validated before). Unfortunately, it is still too optimistic by a factor of about 1.5 (which is better than the uncorrected OLS beta, which is off by a factor of 2.3.[14] Both the OLS and the Stambaugh empirical coefficient estimates suggest that our null process generates higher estimates in simulated samples about 15% of the time.

So far, we have introduced the absolute minimum algebraic admissibility constraint, rather than anything grounded in data. That is, we have not yet considered matching the empirically observed low and high dividend-price ratios of -4.50 and -2.63. Because price-dividend ratios deteriorate rather quickly once they exceed 10%, we effectively only had an upper constraint roughly matching the empirical data, but not an equivalently empirically justified lower constraint. It should be equally bothersome that the canonical null model predicts a much lower dividend-yield today (of $\exp(-5.99) \approx 0.25\%$) than what has been empirically observed in the data (of $\exp(-4.11) \approx 1.64\%$). Finally, it is also the case that the predictive coefficient of 0.070, while not statistically significant, is economically sufficiently above the model-simulated coefficient under the null hypothesis of 0.036 to leave us with a nagging doubt whether to interpret it as tantalizing or not. It turns out that all these issues are linked.

[Insert Figure 6 here: **Goldilocks Density of Minimum and Maximum dp Realizations**]

First, we want to illustrate the mismatch between the canonical null model and empirical reality. Figure 6 shows that the model wants to generate not only a lot of cases in which the maximum dp is above 1.0, but also a lot of cases in which the minimum dp is -6.0 (D/P of 0.25%) or lower. The empirical reality was a market that operated in a much tighter range than predicted by the canonical model.

[14] The Stambaugh 10% level is really the 15% level, the 5% level is really the 7.5% level, and so on.

[Insert Figure 7 here: **Goldilocks minimum dp, maximum dp, and Predictive Coefficients in The Canonical Model**]

[Insert Table 5 here: **Goldilocks Expected Identity Coefficients Given Historical dp Range**]

Figure 7 uses the same coordinate system as Figure 6, but overlays the contour plots for the three predictive coefficients. At the x-y location where the empirical *range* of the dividend-price ratio matches the canonical null model, the empirical predictive coefficients are also well matched to the canonical null model's predictive coefficients. Table 5 shows this numerically. Conditional on the dividend-price ratio staying in the interval from 1.1% to 7.2%, the canonical null model's predictive OLS coefficients were expected to turn out at $b^{\text{ols}}(r_1,\ dp_0) \approx 0.088$, $b^{\text{ols}}(g_1,\ dp_0) \approx -0.012$, and $b^{\text{ols}}(dp'_1,\ dp_0) \approx 0.900$. If anything, the predictive (biased) OLS coefficient of $b^{\text{ols}}(r_1,\ dp_0) \approx 0.070$ seems too low. For the empirical range of the dividend-price ratio in the real world, it is no longer about arguing about how we should interpret a realization at the 85th percentile.

We can summarize the evidence as follows: in a Goldilocks model matching the limited 99-year range of dividend-price ratios just as observed in the data, and in the absence of any predictive ability of dividend-price ratios whatsoever, we would have expected researchers to observe a $b^{\text{ols}}(r_1,\ dp_0)$ coefficient greater (not smaller) than the one observed. Stating the historical dividend-price range thus seems to be enough. It subsumes any further statements about dividend-price ratios having the ability to predict stock returns *somehow*. If there is evidence in favor of true predictability — not as in a linear relationship but as evidenced by the fact that the dividend-price ratio today is not under 0.5% — the statement about the range not having gone below lower than 1.1% can cover it all.

What we are *not* claiming is that the empirical predictive coefficients have a point estimate exactly as expected under a null hypothesis of no association between dividend-price ratios and subsequent stock returns and dividend growths. In fact, we don't even understand the meaning of such a null hypothesis. Instead, we are stating that there is something about the process that has caused D/P not to dip below 1% and above 7.2%, and it was also sufficient to produce the same forecasting coefficient point estimate in the absence of any true underlying predictive association.

VI An Alternative Process

Our second alternative specification dispenses with attempts at embedding a canonical null hypothesis about log stock returns and log dividend growth.[15] Instead, we explore whether we can reduce the frequency of termination by leaning on a model in which the $r_{t,t+1} - g_{t,t+1}$ increases for high dividend-price ratios.

For example, the simplest such model could, by fiat, declare a 100-fold reduction of dividends whenever the dividend-price ratio exceeded, say, 10%. This would eliminate termination scenarios. However, because we have never observed so high a dividend-price ratio for the stock market, in itself, this model is empirically empty. It is merely a mathematical restatement of a philosophically desirable association that firms would not let the dividend-price ratio get out of hand.

There is, however, one possible, meaningful, empirically grounded approach. There were individual stocks that had dividend-price ratios much higher than 7%. We can explore how they forecast stock returns and dividend growth rates under those circumstances. Of course, it is *not* a given that the overall market would behave just like its individual components if it ever reached the higher D/P ratios that we have observed among a few individual stocks.

A Individual Stocks

Our first approach is to work with *all* individual dividend-paying stocks. We first create equivalent annual series of $dp_{t-1,t}$, $r_{t,t+1}$, and $g_{t,t+1}$ from the full CRSP monthly file. We do not eliminate observations in which the subsequent dividend-price ratio was zero, but we do ignore observations in which the prevailing dividend-price ratio was zero. It is important to emphasize that we are *not* running primarily cross-sectional regressions (why one firm's D/P ratio is larger than another's, as in Vuolteenaho (2002)), but are working with time-series (why one firm's D/P ratio changes over time, just as it does in the overall stock market index).

[15] The canonical model's process in effect induces an unconstrained (near) random walk for log dividend-price ratios *a priori*, subject to the termination constraint (an absorbing state).

[Insert Figure 8 here: **Availability of Firm-Specific D/P's in Individual Firms**]

Figure 8 shows the frequency of prevailing (log) dividend-price ratios ($dp_{t-1,t}$) in the data. There are sufficiently many observations for dividend-price ratios between about 0.15% and 22.3% to allow for a meaningful analysis. Thus, the following analysis is based on the sample of 17,903 firms with 217,258 firm-years that had prevailing log dividend-price ratios between -6.5 and -1. We are interested primarily in firm-years with high contemporaneous dividend-price ratios, i.e., those between -2.5 and -1. With thousands of observations, all the discussed associations are statistically significant. We also no longer have to worry about small-sample biases from time-series regressions, because many observations are from the cross-section rather than the time-series.

[Insert Figure 9 here: **Dividend Reductions and Disappearances in Individual Firms**]

Figure 8 shows that about 20% of firms with dividend-price ratios of 15-20% either omitted dividends or disappeared from the CRSP tape the following year. We can interpret the former as a negative infinity on the dividend growth rate. The latter is more difficult to interpret, but it seems reasonable to consider this as due to a strong decline in the stock return (as in a delisting), or a stark increase in the payout (as in an acquisition).

[Insert Figure 10 here: **Log D/P Ratio and Future Log Dividend Growth in Individual Firms**]

Figure 10 shows the association between dividend-price ratios and subsequent dividend growth rates. Each square is the frequency of observations at a vertical and horizontal slice. The advantage of considering the data in vertical intervals non-parametrically is that it is possible to calculate a *median* (rather than a mean) value for the y variable. This is the blue-white non-parametric ragged line in the graph. Given that omission of dividends in the future can be interpreted as negative infinity, the mean would (almost) always be negative infinity. The median avoids these concerns, as it is not impacted by (extreme) winsorization of dependent values at the edges (here, -1 and $+1$). (The shape of a graph based on winsorized means would depend on the winsorization constants chosen.) Finally, the graph shows the spider graph for the value-weighted market index (as in Panel B of Figure 3) in orange.

It is visually obvious that the association between dividend-price ratios and subsequent dividend-price growth is strongly negative within individual stocks. Though the result bears similarity to the finding in Vuolteenaho (2002)), it is very different. Vuolteenaho (2002) documents that firms that have higher D/P ratios have higher growth rates than their peers in the cross-section. We document that firms that have higher D/P ratios tend to have slower dividend growth themselves.

The effect is especially pronounced for dividend-price ratios above 7% (as observed for the overall market). This is a feature that can contribute to reducing the incidence of termination. Stocks with log-dividend-price ratios above -2.0 (about 13% raw dividend-price ratio) are dramatically reducing their dividends. For the few observations where log dividend-price ratios approach -1 (raw dividend-price ratio of 37%), the indicated dividend decline is about 65%. Note further that from the perspective of our theory, it makes little difference whether firms subsequently reduce their dividends because a high dividend payment in one year was extraordinary and unusual, or because the price had declined greatly, thereby leaving the previous dividends too high to be sustained by the now lower price. The effect is the same.

The strong negative slope among individual stocks, even in the range in which the overall stock market remained (and where it showed no association with subsequent dividend-growth rates), is a warning that our premise that the overall stock market would behave analogously to how individual stocks behaved is not necessary. It must be left to the reader to decide how much transferable guidance individual stocks give as far as the overall market is concerned.

[Insert Figure 11 here: **Log D/P Ratio and Future Log Stock Return in Individual Firms**]

Figure 11 shows the equivalent association between dividend-price ratios and subsequent stock returns. Here, we interpret the disappearance of a firm to suggest a higher rate of return that reaches above the (median) dividend-price ratio association. Interestingly, the stock return prediction is almost perfectly linear over the entire range of dividend-price ratios, with a coefficient of about 0.02 — considerably smaller than the 0.07 coefficient observed for the value-weighted stock market, but positive.

Note that the relative importance of dividend growth rate prediction and stock return prediction is the opposite of what we observed for the overall stock market, where

dividend-price ratios had a (visually stronger) association with subsequent stock returns than with subsequent dividend growth rates. The dividend-price ratio predicts dividend growth rates more than it predicts stock returns in the sample of all firms — though it predicts both.

[Insert Figure 12 here: **Log D/P Ratio and Future D/P Ratio for Individual Dividend-Paying Firms**]

Figure 12 contains the most interesting associations. Dividend disappearance the following year is interpreted as a realization more negative than the median, and firm disappearance the following year is interpreted as a realization more positive than the median. Over the domain relevant to the stock market (about 1% to 7% yield), the plot shows a very similar 1-to-1 slope for individual stocks as it showed for the overall stock market. However, in the domain beyond 1% to 7%, the association is no longer 1-to-1. On the left side, it is associated with a higher future dividend-price ratio; on the more interesting right side, it is associated with a lower dividend-price ratio. This is precisely what would be needed to reduce the probability of termination.

However, before concluding that a moderated relationship just beyond the observed dividend-price ratio is a good solution to the termination puzzle, it must be emphasized that the presumption that the overall value-weighted stock-market index would behave like individual firms is just that — a reasonable presumption but not a necessary one.

B Portfolios

Our second approach is to work with portfolios. This avoids having to interpret dividend omissions and delistings. In order not to destroy the variation in lagged D/P ratios (which is what we need to retain to show an extended perspective), we first sort dividend-paying stocks every year, and then form equal-weighted portfolios of 10 consecutive stocks. (This is done by equally averaging the annual portfolio returns of the ten stocks with and without dividends, and then recalculating the exact identity components.) It is also analogous to how researchers work with overall market indexes, except ours is equal-weighted. This results now in 20,475 portfolio-year observations, each with a dividend-price ratio, and a subsequent stock return and dividend growth rate. Note

that we are not tracking the same portfolio over several years. Instead, in each year, the set of portfolios is different. Furthermore, note that we not only gain a wider range in terms of lagged dividend-price ratios, but also a higher variance in terms of stock returns and dividend growth rates.

[Insert Figure 13 here: **Log D/P Ratio and Future Variables for Portfolios**]

Figure 13 shows that the results look very similar to those in Figure 10 and 11. There is a strong non-linear predictive association with log dividend growth rates and a weaker linear predictive association with log stock returns.

[Insert Table 6 here: **Log D/P Ratio and Future Variables under Simulated Null**]

Table 6 explores whether the imputed OLS slope coefficients (of –0.28 and 0.02, respectively) are statistical artifacts. Note that with so many observations, there is no longer a concern for the kind of small-sample time-series bias explored in (Stambaugh (1999)). There is, however, still the concern that we are observing only selected realizations.

We now run a different simulation experiment. We simply reshuffle all stock returns and dividend growth rates, keeping the two together in the same year in order not to destroy their full bivariate distribution (including their covariation). Then, beginning at a log dividend-price ratio of -2.93, we randomly select 99 random portfolio-year observations for r and g, and iterate forward to recover the dividend-price ratio series. With the higher variances, the termination constraint becomes even more binding. Only about 2% of all simulations do not result in strategy termination. A much higher probability of termination for portfolios smaller than the entire market is not unappealing, because 10% of all stocks disappear in any given year. Our analysis is conservative in that by reshuffling, we are ignoring the fact that the true underlying population mean stock returns are likely lower than the observed mean stock returns in the sample, as suggested in Section V.

Table 6 suggests that the positive linear relationship between dividend-price ratios and subsequent stock returns was a statistical artifact. Although the empirical coefficient of 0.020 exceeds the simulation coefficient of 0.010, it only lies at the 73.9th percentile of the simulation coefficients.

VII Prior Literature

Our different underlying process assumptions do not allow our analyses to fit into the frameworks of prior literature. However, it is still useful to compare our analyses of the same variables to those made in the past.

Stambaugh (1999) offers a seminal analysis of the effects of autocorrelated independent variables in predictive regressions. However, his process does not embed the case where the independent variable can step into an invalid region. The analysis does not seem easily transferable to the Goldilocks process with its feasibility and potential selection concerns, and its augmentation with the simple AR1 approximate correction from Kendall (1954) is also quite different from the suggested pattern in Section VI. (It is not clear how to interpret the behavior towards the edges of the observed range in light of an underlying process — other than what we have done in our Goldilocks model with the upper and lower constraints, in which case the bias is much stronger and indeed enough to more than explain the entire predictive coefficient.) As we pointed out above, within the empirically observed domain, the Stambaugh inference performed admirably if imperfectly.

The focus of our analysis was the dividend identity. The seminal Campbell and Shiller (1988), as well as Goyal and Welch (2003), Cochrane (2008), and many others, have conducted related empirical analyses — all, however, ignoring the necessary non-termination constraint. The most articulate defense of the use of an identity-derived model appears in the equally seminal and brilliant Cochrane (2008). Thus, we are focusing the discussion on Cochrane.

As already described, the identity $r_{1,2} - g_{1,2} = (dp_{1,2} - dp_{0,1}) + \eta(\, dp_{1,2}\,)$ is an algebraic necessity without meaning, having to hold for bees and stars just as having to stock returns and dividend growth rates. There is no equilibrium concept in the equation itself, as (e.g.) in a Gordon dividend growth model. It is a simple updating identity. This identity (and its simpler symmetric cousin) also has to hold for two gamblers, A and B. This does not imply that their current wealth ratio must predict either the winnings of A, or B, or the difference in A and B. Their next wealth ratio can absorb all the slack.

As stated clearly by Cochrane and others but often overlooked, the economics are entirely contained in the "auxiliary" statement about whether $(dp_{1,2} - dp_{0,1}) + \eta(\, dp_{1,2}\,)$

is anchored or not. He anchors it by assuming that dp is stationary. In this case, when dp is above (below) its long-term mean, we should see a relatively higher (lower) net difference for $r_{1,2} - g_{1,2}$. The importance of the assumption of the prior constraint about the dividend-price ratio stationarity has also been highlighted before. Campbell and Yogo (2006) write that "since statistical tests do not reject an explosive root for the dividend–price ratio, we have evidence for return predictability only if we are willing to rule out an explosive root based on prior knowledge. This reconciles the 'contradictory' findings by Torous, Valkanov, and Yan (2004, Table 3), who report that the dividend–price ratio does not predict monthly returns in the postwar sample, and Lewellen (2004, Table 2), who reports strong evidence for predictability."

Cochrane defends his stance by writing that "the null should be a coherent, economically sensible view of the world, and explosive price-dividend ratios do not qualify." Further "since the dividend yield is stationary, one of dividend growth or price growth must be forecastable to bring the dividend yield back following a shock."

However, it is not entirely clear to us why explosive price-dividend ratios do not qualify. In particular, besides the fact that this renders economic series only meaningless over infinite time (and there could be a slow migration from explosive towards non-explosive process in the future), we cannot understand why a variable in logs could not trend to negative infinity. An explosive process would only imply that dividend-price ratios would trend towards zero. It is not difficult to write down equilibrium models in which this can occur. But this is perhaps merely a side observation. Most importantly, from the perspective of our paper, even stationary processes are vulnerable to violating the $D<P$ constraint, too.

Because we disagree with Cochrane's statement that a random walk (within the empirically observed range of D/P) does not qualify on a priori grounds, and because the empirical relationship between prevailing and future dividend-price ratio appears to have been nearly perfectly 1-to-1 (within the empirically observed range of D/P), we also disagree with his statement that dividend-price ratios must have predicted either stock returns or dividend growth rates. Instead, it seems primarily that the dividend-price ratio had merely predicted itself in the past. It is like our analogy of a gambler, where the original wealth ratio must neither predict the future winnings of one nor the other gambler. The dividend-price ratio took random steps and predicted only itself. That is,

even if the process had been stationary but with a coefficient close to 1,[16] one still does not have to ask *on annual horizons*, in Cochrane's words, *Which of dividend growth or returns is forecastable? (Or really, "How much of each?")*. The next dividend-price ratio could easily absorb most of the annual realizations of g and r. In this case, it would be only over very long horizons, perhaps centuries, that such a requirement would have bite — and only if it doesn't trend towards negative infinity. And, of course, our analysis is for a finite time series. At some point, infinite series arguments become irrelevant to the human experience.

Cochrane (2008) also suggests that a joint hypothesis test is more powerful in rejecting a null that dividend-price ratios cannot predict stock returns. Our investigations show that this is correct, but further conditions on any model should also themselves be adjusted for when it comes to rejecting the null hypothesis (of no predictive power).

Consider a null hypothesis that is correct — that draws are from a standard normal distribution with a mean of 0 and a standard deviation of 1. Under the null hypothesis, $p(Z_1 > 0) = 50\%$. More importantly, a single hypothesis test should reject the null at the 5% level if the draw is greater than 1.96. Simulating the model and counting up incidences where the null hypothesis generates this high a draw, if we observe a draw of 1.0, we would find that it lies at the 84th percentile, $p(Z_1 > 1.0) = 16\%$. We would not reject the null. But the null further predicts that a second draw should also be zero. The joint hypothesis test, based on two draws, would leave us with $p(Z_1 > 0 \cap Z_2 > 0) = 25\%$, and $p(Z_1 > 1.0 \cap Z_2 > 0) = 8\%$. Repeat with three, and the equivalent numbers are 12.5% and 4%. A fairer benchmark to test the null hypothesis would be to divide the rejected frequency by the area of the quadrant under the null. In this case, with all subsequent draws smack on the value predicted by the null hypothesis, with one variable, we would reject 16/100; with two variables, we would reject 8/50; and with three variables, we would reject 4/25, where the divisor depends on the number of conditions. In other words, a fairer benchmark to reject the null hypothesis would be relative to the frequency in the quadrant of interest.

Cochrane is correct that this is indeed a more powerful test, but it is still not sufficient to raise even the non-bias-adjusted probability to an acceptable statistical significance

[16] If one (incorrectly) ignores the $D<P$ constraint, then standard econometric tests strongly suggest that dividend-price ratios followed a random walk with negative drift, of course only in the observed 1% to 7% interval (see Table 2).

level. Properly adjusted for additional conditions, in our Goldilocks model, adding the model condition on g reduces the probability that stock returns are *not* predictable **from 15% to 10%**. It does not reduce it from 15% to 5%. (Of course, as shown above, accounting for the empirical range of dp from 1.1% to 7.2% further reduces the significance to the point where there is no longer any question.)

As Boudoukh, Richardson, and Whitelaw (2008) (and our own Table 4) point out, adding longer-term predictive associations does not improve our ability to reject the model *conditional* on non-termination. Cochrane suggests working with much longer time horizons. We did not explore this further.

Finally, it is worth mentioning an interesting conceptually related problem. Jiang et al. (2024) note the ever-increasing and potentially unsustainable US debt-to-GDP ratio. There is plausibly a similar "death spiral" at work in their context. However, their concern was less pressing in the past than it is today. Past data continues to be easily interpretable. Our concern is the opposite. It was more pressing in the past than it is today (with dividend yields under 2% per year). This makes past stock return data more difficult to interpret, but also less worrisome.

VIII Conclusion

Our paper has raised a constraint that has hitherto been ignored. Dividends cannot exceed prices. We interpret the draw of a dividend in a process that would violate this condition as "termination." Thus, our paper has raised a novel puzzle. If stock returns are normally iid distributed (as the evidence suggests they are) and if dividend growth rates are normally iid distributed (as the evidence suggests they are), and with the observed cross-correlation, then starting at the 1926 dividend-price ratio of 5.3%, we should have expected the stock market to have disappeared about 65% of the time. This is even though stock returns exceeded dividend growth rates by a remarkable 5% per year. However, this may be too low an estimate of the termination probability because of selection bias. If we interpret the world as a Goldilocks scenario that happened to have avoided termination, then a better estimate of the termination probability is 80%. Non-termination of investors to whom the identity was applied was not a foregone

conclusion. They survived the Great Depression, World War II, and rampant stagflation in the Carter years. Non-termination occurred only under complex and path-dependent constraints on dividend-price ratios, stock returns, and dividend growth rates that were required to make our presence and measurement as of 2025 ultimately possible.

On the most basic level, it seems impossible even to formulate an empirically plausible null hypothesis in which dividend-price ratios would not have predicted stock returns (net of dividend growth rates) *and* one would not have had a concern that the stock market process would have come to an end as the price could no longer sustain the dividends. Thus, perhaps most simply, our paper suggests that the literature should be clearer about what the underlying model is that it is estimating and testing, including what would happen in the model if dividend-price ratios happened to reach higher than the hitherto highest observed empirical D/P level of 7.2%.

Our paper could not offer one conclusive answer, but it could offer a first and early exploration of two potential approaches.

If we interpret the world as a Goldilocks conditional scenario, in which the stock market just happened to have experienced a D/P range from about 1.1% to about 7.2% over the last 99 years, then the predictive OLS coefficient estimate for stock returns is not only not statistically significant, but, if anything, *lower* than expected under a zero relationship between lagged dp and future r and g. The observation of the 1.1% lower bound is more than enough to render the 0.076% predictive coefficient a bias artifact that one would likely have observed in the absence of any true predictive power. It also does not imply an unconditional absence of association between dp and subsequent variables — an absence relation that we repeatedly explained we don't even know how to specify appropriately.

If we interpret the world as following a model in which dividend-price ratios can have level-dependent predictive power, and *this level-dependent association is the same as it is in individual stocks*, then we find that dividend-price ratios predicted primarily (or only) dividend growth rates, especially at high dividend-price ratios. Stocks aggressively reduced their dividend payments the subsequent year when their dividend-price ratios moved above 7% — possibly due to high dividends or low stock returns (from the perspective of the theory, the reason need not matter).

Although our two approaches — the Goldilocks view and the analogy with individual stocks — sketch viable directions for working with dividend-price ratio constraints, they are not the only venues for future research. Macro-finance analysis could expand the theory to utilize other instrumental variables that could help offer suggestions as to how stock returns and dividend growth rates should behave at higher dividend-price ratios than 7%. Empirical analyses could expand towards other measures (such as dividends plus repurchases, though there is no easy analogy when it comes to overall net inflows and net outflows), as well as to other financial markets.

[Insert Table 7 here: **MSCI Dividend Yields, March 2025**]

On a final note, the termination concern may seem less urgent today when dividend-price ratios hover around 1.5%. However, a better understanding of stock price and dividend processes would seem necessary if we want to interpret the last 100 years of observed history — and it is of course this history that we usually muster when we want to determine whether dividend-price ratios predicted stock returns. Furthermore, the termination concern is also not completely outdated in other locations. Table 7 shows that the U.S. is an outlier in terms of low dividend-price ratios today. Australia, the United Kingdom, and Brazil have higher dividend-price ratios and are thus potentially more vulnerable today to future termination concerns than the United States.

IX Appendixes

References

Boudoukh, Jacob, Matthew Richardson, and Robert F Whitelaw (2008). "The myth of long-horizon predictability". In: *The Review of Financial Studies* 21.4, pp. 1577–1605.

Campbell, John Y and Motohiro Yogo (2006). "Efficient tests of stock return predictability". In: *Journal of financial economics* 81.1, pp. 27–60.

Campbell, John Y. and Robert J. Shiller (1988). "The dividend-price ratio and expectations of future dividends and discount factors". In: *The review of financial studies* 1.3, pp. 195–228.

Cochrane, John H. (2008). "The dog that did not bark: A defense of return predictability". In: *The Review of Financial Studies* 21.4, pp. 1533–1575.

Dickey, David A. and Wayne A. Fuller (1979). "Distribution of the estimators for autoregressive time series with a unit root". In: *Journal of the American Statistical Association* 74.366a, pp. 427–431.

Goyal, Amit and Ivo Welch (2003). "Predicting the equity premium with dividend ratios". In: *Management Science* 49.5, pp. 639–654.

Jiang, Zhengyang, Hanno Lustig, Stijn Van Nieuwerburgh, and Mindy Z Xiaolan (2024). "The US public debt valuation puzzle". In: *Econometrica* 92.4, pp. 1309–1347.

Jorion, Philippe and William N. Goetzmann (1999). "Global stock markets in the twentieth century". In: *The journal of finance* 54.3, pp. 953–980.

Kendall, Maurice G (1954). "Note on bias in the estimation of autocorrelation". In: *Biometrika* 41.3-4, pp. 403–404.

Kwiatkowski, Denis, Peter C.B. Phillips, Peter Schmidt, and Yongcheol Shin (1992). "Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?" In: *Journal of econometrics* 54.1-3, pp. 159–178.

Lewellen, Jonathan (2004). "Predicting returns with financial ratios". In: *Journal of Financial Economics* 74.2, pp. 209–235.

Marriott, FHC and JA Pope (1954). "Bias in the estimation of autocorrelations". In: *Biometrika* 41.3/4, pp. 390–402.

Stambaugh, Robert F. (1999). "Predictive regressions". In: *Journal of financial economics* 54.3, pp. 375–421.

Torous, Walter, Rossen Valkanov, and Shu Yan (2004). “On predicting stock returns with nearly integrated explanatory variables”. In: *The Journal of Business* 77.4, pp. 937–966.

Vuolteenaho, Tuomo (2002). “What drives firm-level stock returns?” In: *The Journal of Finance* 57.1, pp. 233–264.

Table 1: Notation Glossary

$r_{t,t+y}$	plain rate of return from time t to time $t+y$, including compounded dividends
$x_{t,t+y}$	plain rate of return from time t to time $t+y$, excluding all dividends, compounded

(Page 10 gives formulas to calculate quantities below.)

$D_{t,t+y}$	dividends paid from time t to time $t+y$
P_t	price at the end of year t.
$dp_{0,1}$	$\log\left[D_{0,1}/P_1\right]$: dividend-price ratio with all dividends from time 0 to time 1.
$r_{1,2}$	$\log\left[(P_2+D_{1,2})/P_1\right] = \log\left[1+R_{1,2}\right]$: Log stock-return with all dividends.
$g_{1,2}$	$\log(D_{1,2}/D_{0,1})$: Log dividend growth rate. Needs dividend from *three* periods.
$dp'_{1,2}$	$-\log\left[1+1/(D_{1,2}/P_2)\right]$: Used only in exact identity. Usually similarly to $dp_{1,2}$.
$CS\Delta$	$dp_{1,2} - dp_{0,1}$.
$\Delta \equiv EI\Delta$	$dp'_{1,2} - dp_{0,1}$.
ψ	parameter used in CS approximation. $\psi^{-1} \approx (1+e^{\overline{\log D/P}}) \approx (1+e^{-3.384}) \approx 1.0339$, $\psi \approx 0.9672$
k	constant used in CS approximation. $k_1 = -\log\psi - (1-\psi)\log(1/\psi - 1) \approx 0.144$.

OLS Coefficients Explaining Future Quantities With Dividend-Price Ratios

$b^{ols}(\, r_1,\ dp_0\,)$	future log returns with prevailing log dividend-price ratio.
$b^{ols}(\, g_1,\ dp_0\,)$	future log dividend growth with prevailing log dividend-price ratio.
$b^{ols}(\, dp'_1,\ dp_0\,)$	future transformed log dividend-price ratio with prevailing log dividend-price ratio.

Note: When time subscripts are quoted as integers, they are relative to an arbitary starting point t. (This shortens notation in subscripts from "t–1,t" to "0,1.") In the predictive coefficients, there is only one subscript for a two-subscript quantity, because only the second subscript (time) is quoted.

Table 2: Summary Statistics for the Value-Weighted CRSP Market Portfolio, 1926-2024

Panel A: CRSP Inputs

	#years	mean	sd	pmin	10%	50%	90%	pmost	autocor
yyyy	99	1975		1926		1975		2024	
vwretx	99	0.080	0.192	–0.477	–0.145	0.110	0.285	0.509	–0.018
vwretd	99	0.119	0.198	–0.444	–0.112	0.156	0.332	0.574	–0.016

Panel B: Dividend-Price Ratios

	#years	mean	sd	pmin	10%	50%	90%	pmost	autocor
D/P	99	0.0365	0.0155	0.011	0.018	0.035	0.057	0.072	0.912
dp	99	–3.406	0.451	–4.495	–4.010	–3.359	–2.865	–2.626	0.945
D/P	98	0.0367	0.0154	0.011	0.018	0.035	0.057	0.072	0.910
dp'	98	–3.447	0.436	–4.506	–4.029	–3.399	–2.920	–2.696	0.945

First D/P (1926): **5.3% (**$-2.93\log$**)**.
Last D/P (2024): 1.6% ($-4.15\log$). [last used for prediction, 2023: 1.8% ($-4.01\log$)])

Panel C: Returns and Dividend Growth

	#years	mean	sd	pmin	10%	50%	90%	pmost	autocor
$r_{t-1,t}$	98	0.095	0.194	–0.587	–0.119	0.146	0.286	0.454	0.034
$g_{t-1,t}$	98	0.047	0.141	–0.368	–0.099	0.051	0.216	0.370	–0.079
$r_{,} - g_{,}$	98	0.048	0.146	–0.340	–0.160	0.059	0.209	0.545	0.062

Correlations: $g_{t-1,t}$ and $r_{t-1,t}$: **66.0%**

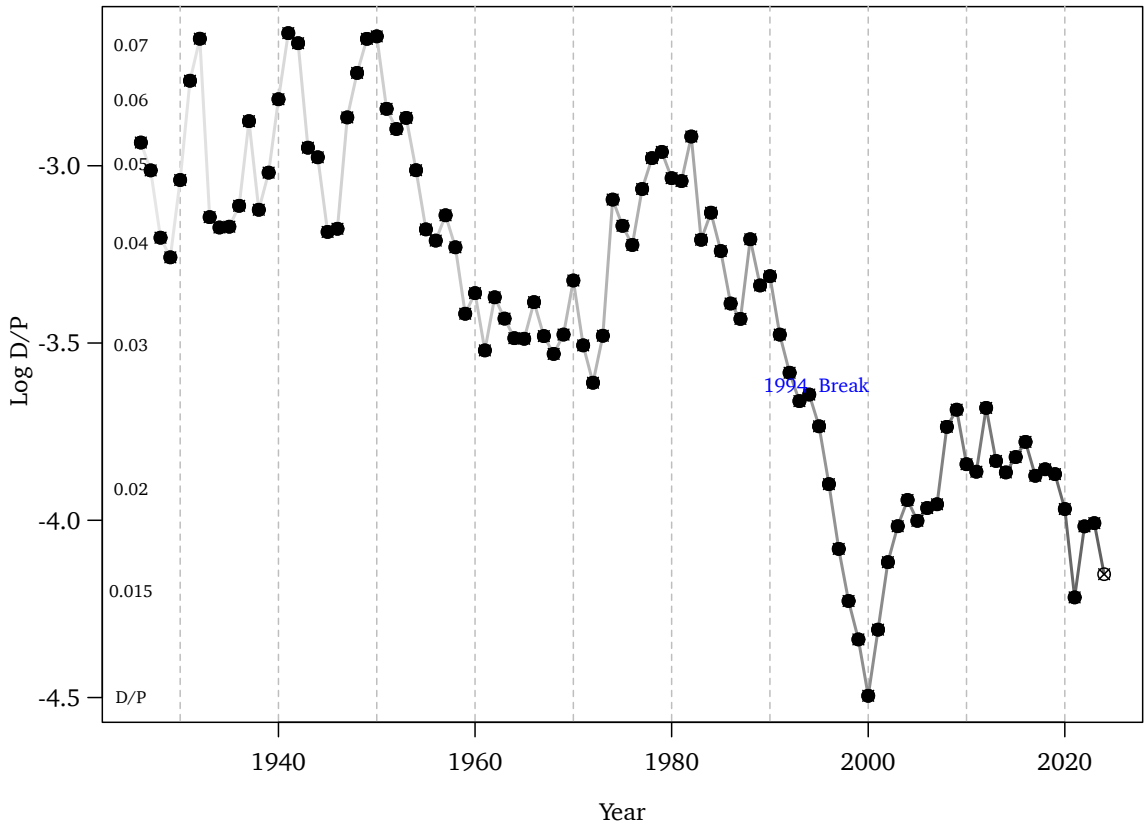
Panel D: Stationarity Statistics

	ADF	p(ADF)	Interpret	KPSS	p(KPSS)	Interpret
$p_t \equiv \sum r_{t-1,t}$	–2.53	0.36	Non-Stationary	+2.53	<0.01	Non-Stationary
$d_t \equiv \sum g_{t-1,t}$	–4.20	<0.01	Stationary	+2.50	<0.01	Non-Stationary
$dp_{t-1,t}$	–2.64	0.31	Non-Stationary	+1.88	<0.01	Non-Stationary
$dp'_{t-1,t}$	–2.49	0.38	Non-Stationary	+1.85	<0.01	Non-Stationary

Explanations: The statistics are calculated from annual returns with and without dividends (vwretd and vwretx, respectively), as described in equation 3. To calculate the first growth rate of dividend consumes one extra year. Series with 98 observations exclude the final 2024 D/P ratio, because it cannot be used for prediction. Panel D: ADF is the augmented Dickey and Fuller (1979) statistic. KPSS is the Kwiatkowski et al. (1992) statistic.

Source: 1mkexactidentity.R: Panels A, B, C.
2stationarity-tests.R: Panel D.

Figure 1: Log D/P Ratio

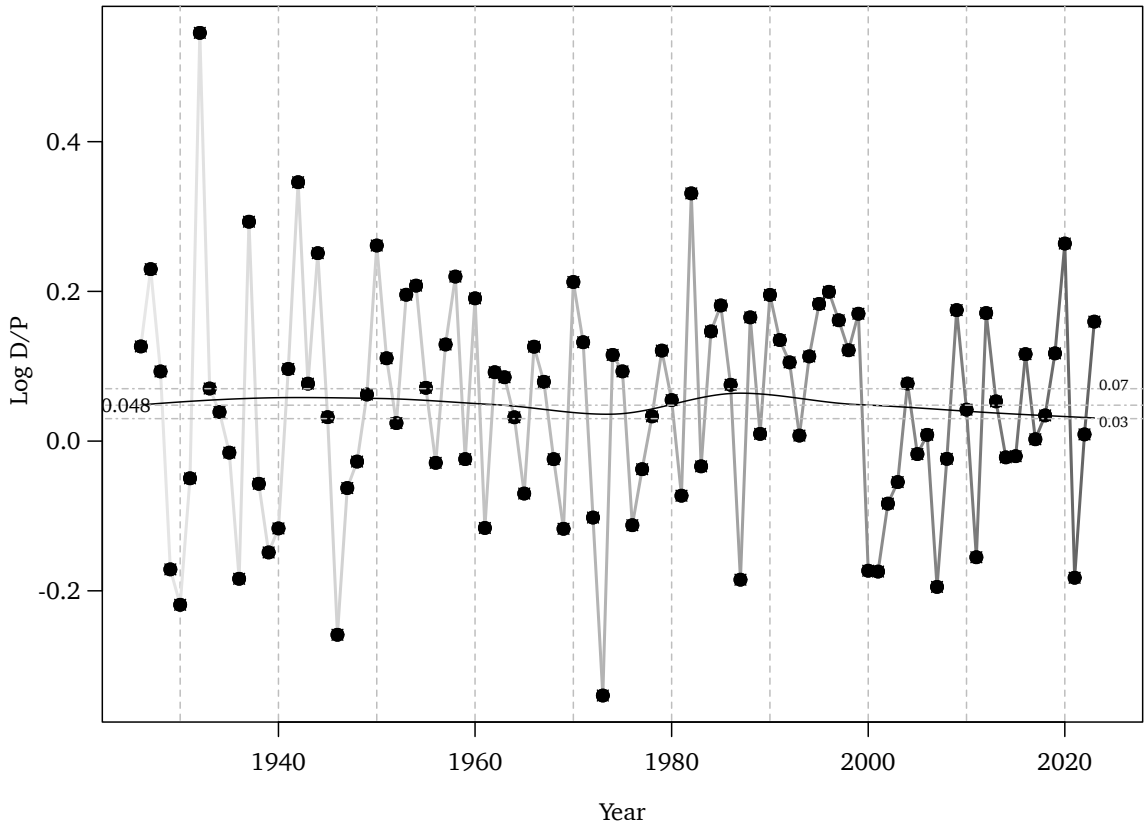


Explanations: This is the familiar log-dividend price ratio for the value-weighted stock market index in the CRSP sample (1926-2024). The connecting lines are lighter for observations longer ago.

Interpretation: There were four peak dividend-price ratios: 1931-32, 1941, 1949-1950, and 1978-1982. All observations before 1994 had dividend-price ratios higher than those after 1994.

Source: 2plot-descriptives.R

Figure 2: Log Rates of Return Net of Log Dividend Growth Rates ($r-g$)

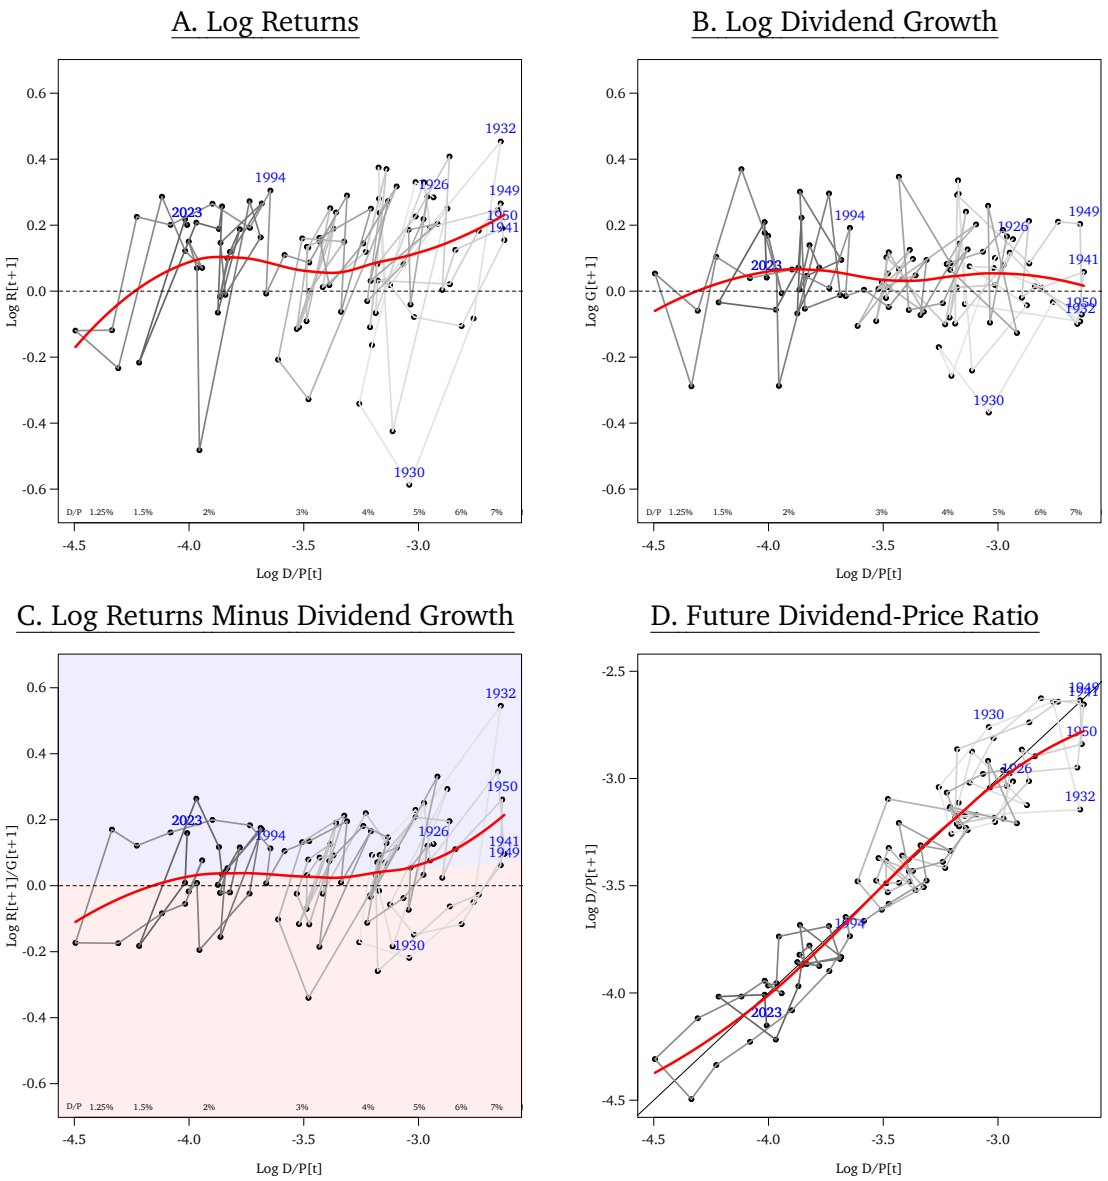


Explanations: This is the familiar difference between log-returns and log-dividend growth rates for the value-weighted stock market index in the CRSP sample (1926-2024). The connecting lines are lighter for observations longer ago. The smooth line around 0.05 is based on loess.

Interpretation: There were no obvious secular trends in the sample (at least until 1990). The mean remained tight around 5%, with the lowest values in the 1970s around 4.2% and the highest values in the 1980s and early 1990s around 6.9%.

Source: 2plot-descriptives.R

Figure 3: Log D/P Ratio and Future Variables



Explanations: These panels show future realizations based on the prevailing log D/P ratio for the value-weighted stock market index in the CRSP sample. The connecting lines are lighter for observations longer ago. The red lines are best-loess fits. In Panel C, the background color indicates whether, given the $(r - g)$, realization, the D/P ratio would increase or decrease. In Panel D, the diagonal line indicates the one-to-one association.

Interpretation: Panel A shows the familiar *known-to-be-biased* positive association between log D/P and subsequent log stock returns. Panel B shows the weak association between log D/P and subsequent log dividend growth. Panel C shows the net effect. Panel D shows that the D/P ratio was nearly perfectly on the diagonal line between -4 and -3, but it frayed at both ends, thus taking on a mild S shape.

Source: 2plot-descriptives.R

Table 3: Predictive OLS Coefficients

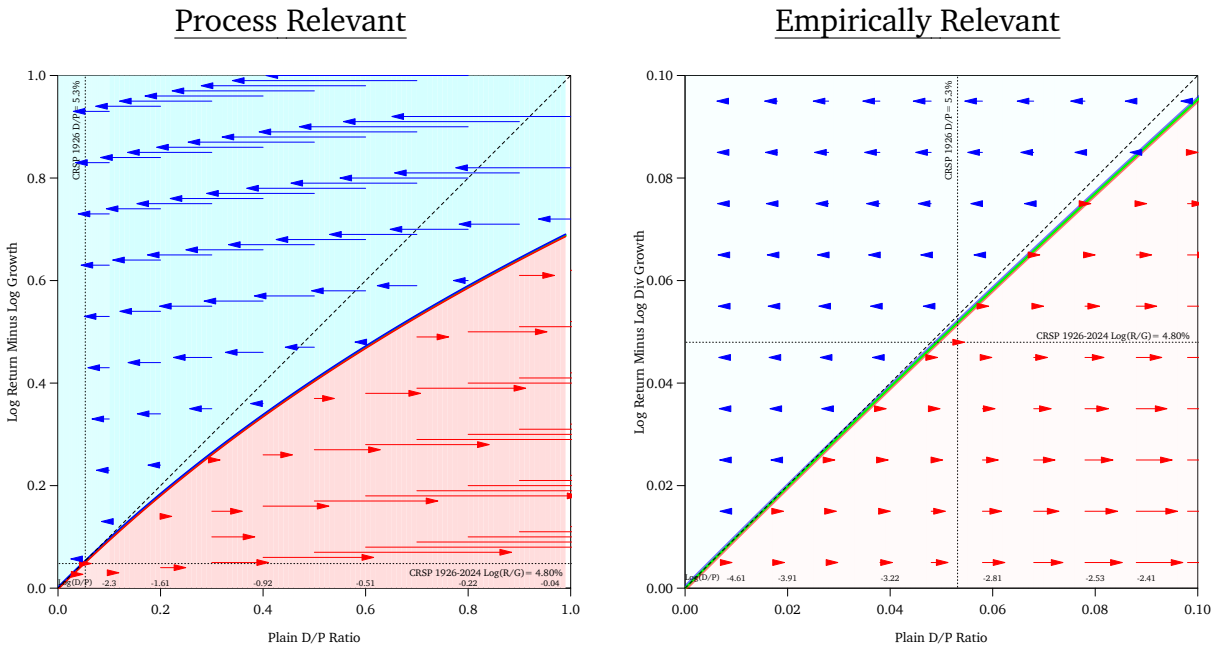
	Dependent	a	T(a)	b	T(b)	N
A	$r_{t,t+1}$	0.353	2.38	0.076	1.79	98
B	$g_{t,t+1}$	0.043	0.40	–0.001	–0.03	98
C	$r_{t,t+1} - g_{t,t+1}$ $\equiv dp'_{t,t+1} - dp_{t-1,t}$	0.310	2.42	0.077	2.10	98
D	$dp_{t,t+1}$	–0.172	–1.49	0.953	25.09	98
E	$\psi \cdot dp_{t,t+1}$	–0.166	–1.49	0.922	25.09	98
F	$dp'_{t,t+1}$	–0.310	–2.78	0.923	25.13	98

Explanations: The sample is the CRSP data described in Table 2. The independent variable is always $dp_{t-1,t}$ (the preceding year's log dividend-price ratio). The standard errors are Newey-West adjusted with one-year lag. By the identity, the slope coefficients (b) add to $[(A) - (B)] + (F) = 1$.

Interpretation: The C-S approximation error is negligible in the range in which the empirical dividend-price ratio drifted. The difference between stock returns and dividend growth is just about statistically significant at conventional levels.

Source: 1mkexactidentity.R

Figure 4: Dynamics of Dividend-Price Ratios

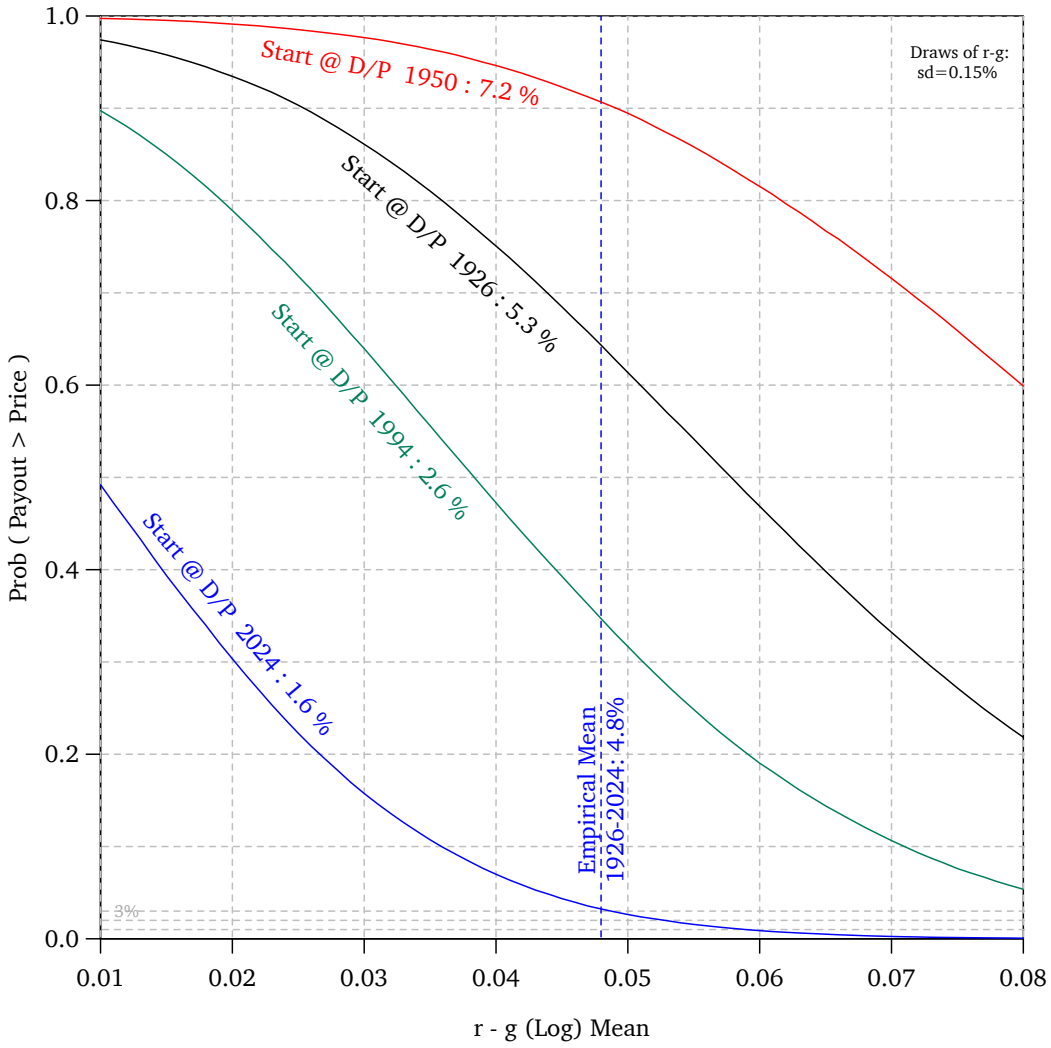


Explanations: This figure illustrates where an $(r-g)_{t,t+1}$ draw will push the next dividend-price ratio (from the tail to the head of the arrow). For example, if the D/P ratio is 0.3, and the log return minus log dividend growth is 10%, the next D/P ratio will be about 0.37. The D/P ratio remains stable at $(r-g)_{,} = \log(1 + \text{D/P})$.

Interpretation: At the CRSP average 1926-2024 $(r-g)_{,}$ of 4.8%, the dividend-price ratio of 5.3% in 1926 sat almost exactly on the critical line.

Source: 5plot-dynamics.R

Figure 5: Markovian Process: Demise Outcomes



Explanations: These are 99-year outcomes in which $D>P$, when draws of $(r-g)$, are from iid unconditional normal distributions with a mean shown on the X-axis and a standard deviation of 0.146 (as observed in the 1926-2024 data). For example, starting at the 1926 D/P ratio of 5.3%, if $(r-g)$, are iid draws based on the 4.8% mean (as was their mean in the 1926-2024 data), then $D>P$ occurs about 65% of the time in simulated 99-year samples.

Interpretation: AR1 process violations in the historical 1926-2024 data should not be viewed as obscure. However, forward-looking as of 2024, if historical stock returns net of dividend growth rates continue at 4.8%, then the probability of demise is no longer central, because the dividend-ratio is so low. However, (a) this low a dividend-price ratio has been a fairly recent development and (b) Figure 2 offered some visual indication that the mean $(r-g)$, may have declined.

Source: 5plot-markov.R

Table 4: Five-Parameter Canonical Goldilocks Model Simulations

Panel A: Matched Empirical Moments

		col-to-be-deleted	Empirical	Simulated Mean	Simulated Sdv	Simulated Prctl
Log Return	Mean	.mn.r	0.095	0.095	0.018	50.3
	Sdv	.sd.r	0.194	0.194	0.014	50.6
Log Div Growth	Mean	.mn.g	0.047	0.047	0.014	50.0
	Sdv	.sd.g	0.141	0.141	0.010	50.4
Correlation		.cor.r.g	0.660	0.660	0.058	47.4
Covariance		.cov.r.g	0.018	0.018	0.003	50.1
Log Ret – DivGro	Mean	.mn.rmg	0.048	0.048	0.011	51.6
Log Ret – DivGro	Sdv	.sd.rmg	0.146	0.146	0.011	51.7

Panel B: Simulation Information

	Freqency	Fraction
Feasible Simulations	17.5%	(175,297/1,000,000)

Panel C: Dividend-Price Ratio Range

	col-to-be-deleted	Empirical	Simulated Mean	Simulated Sdv	Simulated Prctl	(not in logarithm) Empirical	(not in logarithm) SimMean
min dp	dpr.min	–4.50	–6.07	1.35	86.8	1.1%	0.2%
max dp	dpr.max	–2.63	–2.73	0.31	77.0	7.2%	6.5%
$dp_{,2024}$	dpr.last	–4.01	–5.65	1.69	83.9	1.8%	0.4%
$\Delta dp_{1926,2024}$	dpr.decline	1.07	2.72	1.69	16.1		

Panel D: Dividend Identity (Forecasting) Moments

	col-to-be-deleted	Empirical	Simulated Mean	Simulated Sdv	Simulated Prctl
$b^{ols}(r_1,\ dp_0)$	beta.r.ldp	0.076	0.041	0.043	84.1
$b^{ols}(g_1,\ dp_0)$	beta.g.ldp	–0.001	–0.004	0.023	53.9
$b^{ols}(dp'_1,\ dp_0)$	beta.dpp.ldp	0.923	0.955	0.040	16.7
$b^{ols}(dp_1,\ dp_0)$	beta.dp.ldp	0.948	0.972	0.036	18.1
$b^{ols}(r_1+r_2+r_3,\ dp_0)$	beta.r3.ldp.x	0.179	0.120	0.121	77.1
$b^{ols}(g_1+g_2+g_3,\ dp_0)$	beta.r3.lr.x	–0.172	–0.025	0.173	20.1

Panel E: Rejecting (Correct) Null Hypothesis (No Predictive Ability)

			Simulated			Sim Percentile			
	col-to-be-deleted	Empirical	Mean	Sdv	Prctl	90%	95%	97.5%	99%
$b^{\text{ols}}(r_1,\ dp_0)$	beta.r.ldp	0.070	0.036	0.039	85.8	77.2%	88.4%	94.1%	97.7%
$b^{\text{stam}}(r_1,\ dp_0)$	bstambaugh.r.ldp	0.042	0.005	0.044	84.4	85.0%	92.5%	96.2%	98.5%

Explanations: Based on 1,000,000 simulations, each starting at a dp ratio of –2.93 (as in 1926) and running for 99 years, the simulation model parameters were chosen to be a best fit of means in Panel A. This resulted in an underlying conjectured "canonical null model" (in which not only stock returns but also dividend growth rates were unpredictable) of

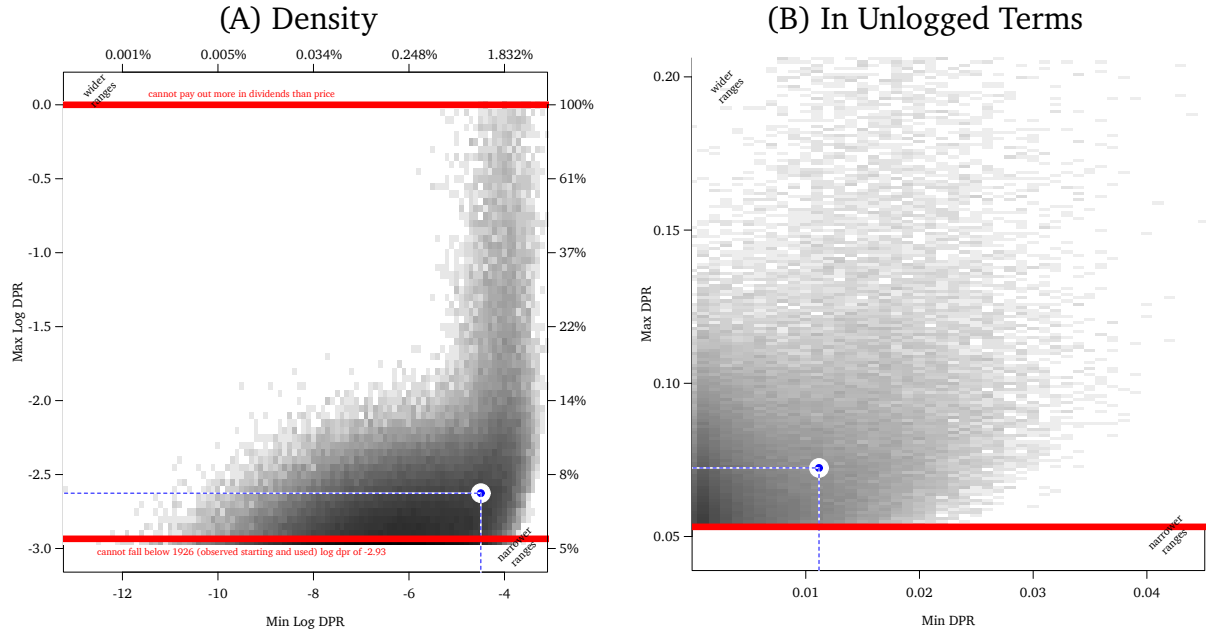
$$\begin{aligned} r_t &\approx 0.081 + 0.1944 \cdot \epsilon_t^r \\ g_t &\approx 0.048 + 0.1056 \cdot \epsilon_t^g + 0.0936 \cdot \epsilon_t^r \end{aligned}$$

and all epsilons were iid unit-normal draws. Simulated moments (in all panels, except of course B) are based on the 175,297 simulations in which the model did not want to step out of the valid domain ($dp < 0$). "Mean" and "Sdv" apply to the simulations, while "pctl" is the location of the empirical estimate in the distribution of simulated estimates. The "Sim Percentile" in Panel E refers to the p-statistic of the coefficient.

Interpretation: (1) Survivorship bias induces the empirical mean return (9.5%) to be above the simulated mean return (8.1%) Other model moments are closer to their empirical moments. (2) Far more often than not (about 4 in 5 times), this canonical model would step into the demise region. *Conditional on feasible runs*, the model fits $b^{\text{ols}}(r_1,\ dp_0)$ and $b^{\text{ols}}(g_1,\ dp_0)$ quite well (at the 15.9th and 46.1th percentiles). (3) The commonly used asymptotic-bias correction of Stambaugh (1999) reduces the gap between the estimate and the correct coefficient estimate and the rejection probability by about half.

Source: optmodel5/optmodel.R. The .Rout file has nice kable output, except for Panel E.

Figure 6: Goldilocks Density of Minimum and Maximum *dp* Realizations



Explanations: The canonical data-generating model explored here is

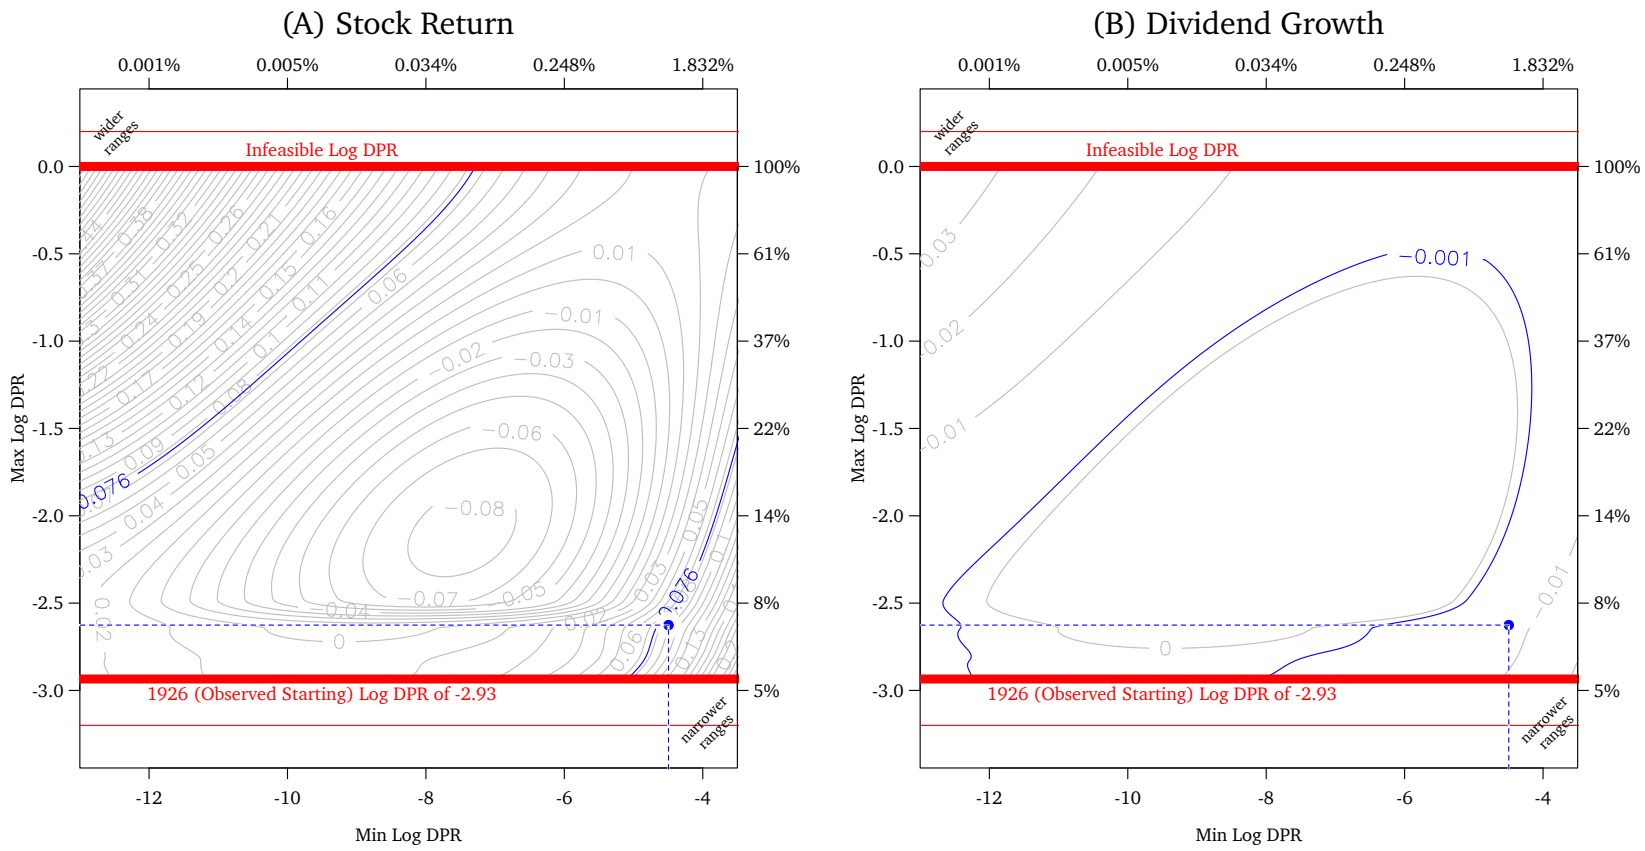
$$\begin{aligned} r_t &\approx 0.081 + 0.1944 \cdot \epsilon_t^r \\ g_t &\approx 0.048 + 0.1056 \cdot \epsilon_t^g + 0.0936 \cdot \epsilon_t^r \end{aligned}$$

with t from 1 to 98. The five numerical parameters were chosen to match the empirically observed stock returns and dividend growths equivalent moments in all simulations that did not step out of the physically possible space. The frequency of each dual-draw (per 98-year draw), i.e., observed *dp* ranges, is indicated by the darkness of the gray rectangle in the coordinate system. The blue dot indicates the actually-observed empirical range of the dividend-price ratio in the CRSP data, which was from -4.50 ($D/P \approx 1.1\%$ in 2000) to -2.63 ($D/P \approx 7.2\%$ in 1932, 1941, and 1950). There would have been a large mass of points ($> 80\%$ of the sample) on the top right of the left graph that have not been plotted, because they would have stepped into the demise region.

Interpretation: The CRSP pair (minimum, maximum) is modestly unusual. In particular, the minimum is not as low as the model suggests.

Source: optmodel/3ddensity.R

Figure 7: Goldilocks minimum dp, maximum dp, and Predictive Coefficients in The Canonical Model



Explanations: These are contourplots overlaid on top of the surface of Figure 6 for the identity coefficients. The blue dot again indicates the actually-observed empirical range of the dividend-price ratio in the CRSP data, which was from –4.5 ($D/P \approx 1.1\%$) to –2.6 ($D/P \approx 7.2\%$). The predictive coefficient contour line that matches the empirical predictive OLS predictive coefficient is in blue.

Interpretation: The dividend growth coefficient contour is very flat with respect to the dividend-price ratio range. Given the observed range of historical dividend-price ratios, the model seems fits the predictive coefficients well. The empirical $b^{\text{ols}}(r_1, dp_0)$ seems too low for a model in which dividend-price ratios do not predict stock returns.

Source: optmodel/3dcontours-1calc.R, optmodel/3dcontours-1plot.R

Table 5: Goldilocks Expected Identity Coefficients Given Historical dp Range

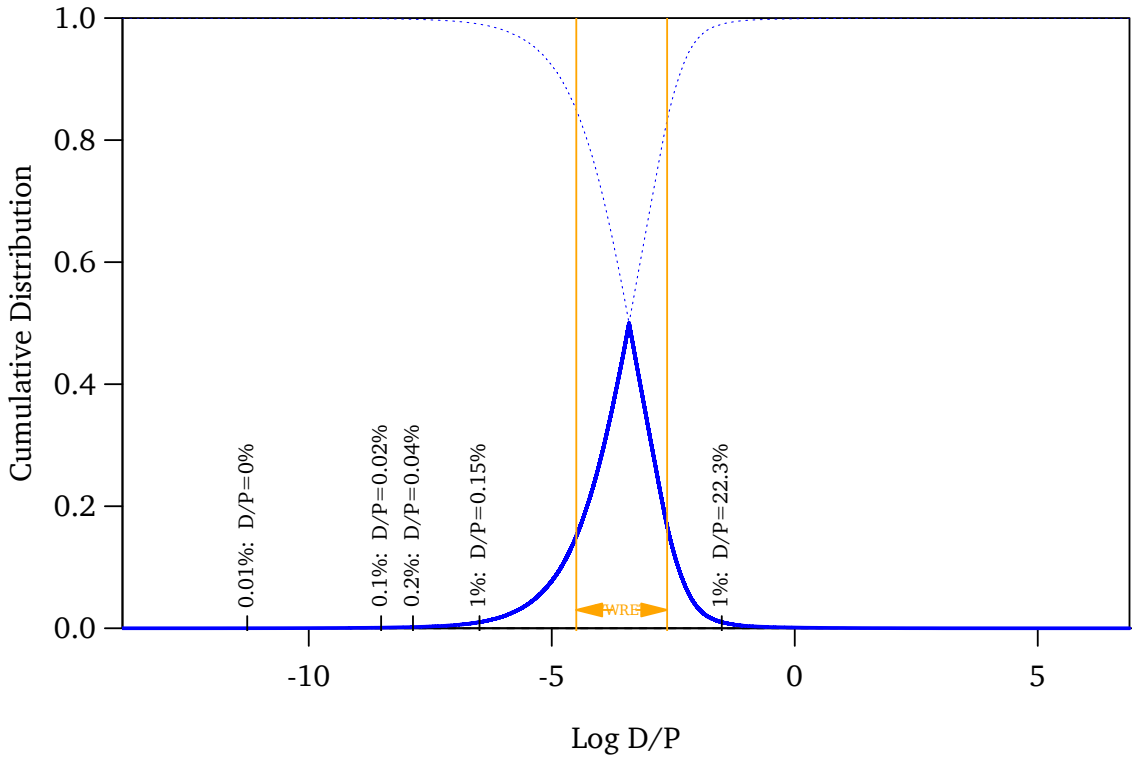
		dp Range		Identity Coefficients		
		$\min(dp)$	$\max(dp)$	$b^{ols}(\, dp'_1,\ dp_0\,)$	$b^{ols}(\, r_1,\ dp_0\,)$	$b^{ols}(\, g_1,\ dp_0\,)$
Sim	All feasible	$-\infty$	0.0	0.955	0.041	–0.004
	Loess-Fitted Near	–4.50	–2.63	0.906	0.086	–0.008
			Empirical:	0.923	0.076	–0.001

Explanations: The constraints on dp imposed on the canonical null model are in columns 2 and 3. The first row imposes only a non-demise constraint. The second row imposes a "match empirical reality" constraint (i.e., a D/P ratio ranging from 1.1% to 8.0%).

Interpretation: Given [a] a historical range of dp from 1.1% to 8.0%, [b] in which the stock market survived, and [c] in which the dividend-price ratio did not forecast stock returns or dividend growth: we should have seen an OLS $b^{ols}(\, r_1,\ dp_0\,)$ coefficient of 0.086 — higher than the empirically observed 0.076.

Source: optmodel/3dcontour-1calc.R(out)

Figure 8: Availability of Firm-Specific D/P's in Individual Firms



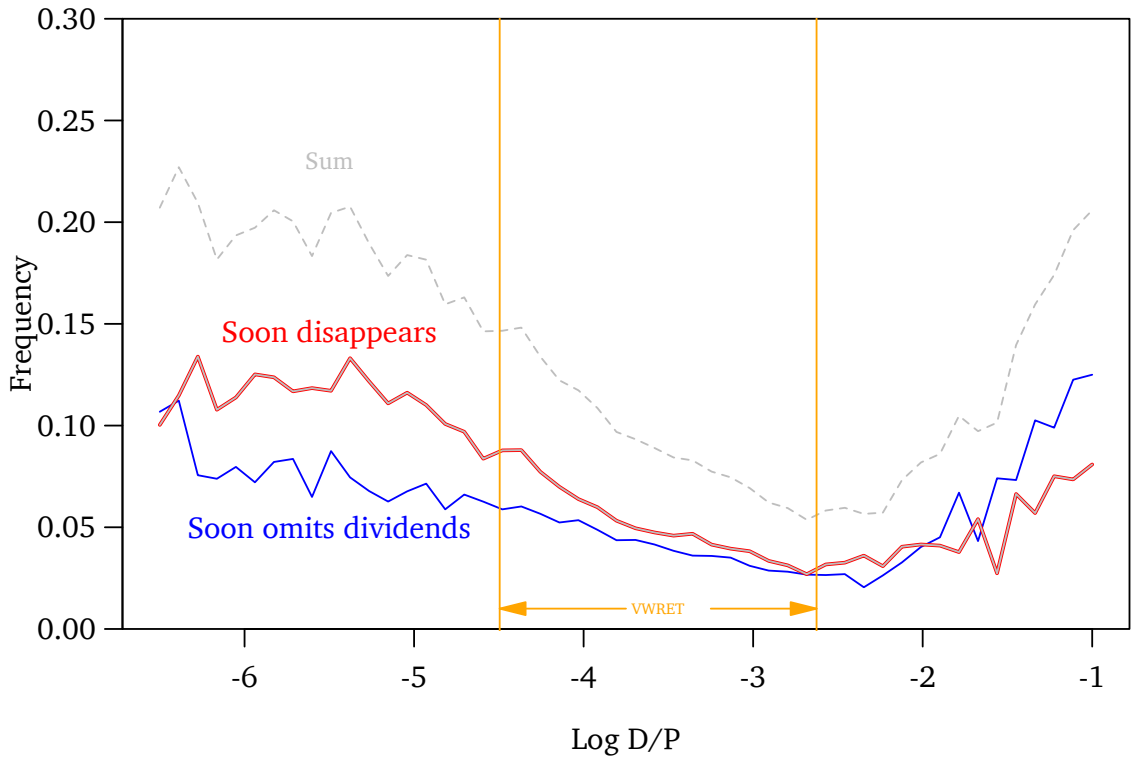
	N	Mean	10%	Med	90%
$dp_{t-1,t}$	217,258	—3.52	–4.81	–3.41	–2.40

Explanations: These are the empirical cumulative distribution functions for the universe of all CRSP firms paying dividends from 1926-2024.

Interpretation: It is reasonable to explore the 98% of observations that lie between the 1st and the 99th percentile of dividend price ratios, which in logs means between -6.5 to -1. Below a dividend price ratio of 0.15% or above 37%, there are simply no longer enough observations to draw reliable inference.

Source: 2plot-many-hist.R

Figure 9: Dividend Reductions and Disappearances in Individual Firms

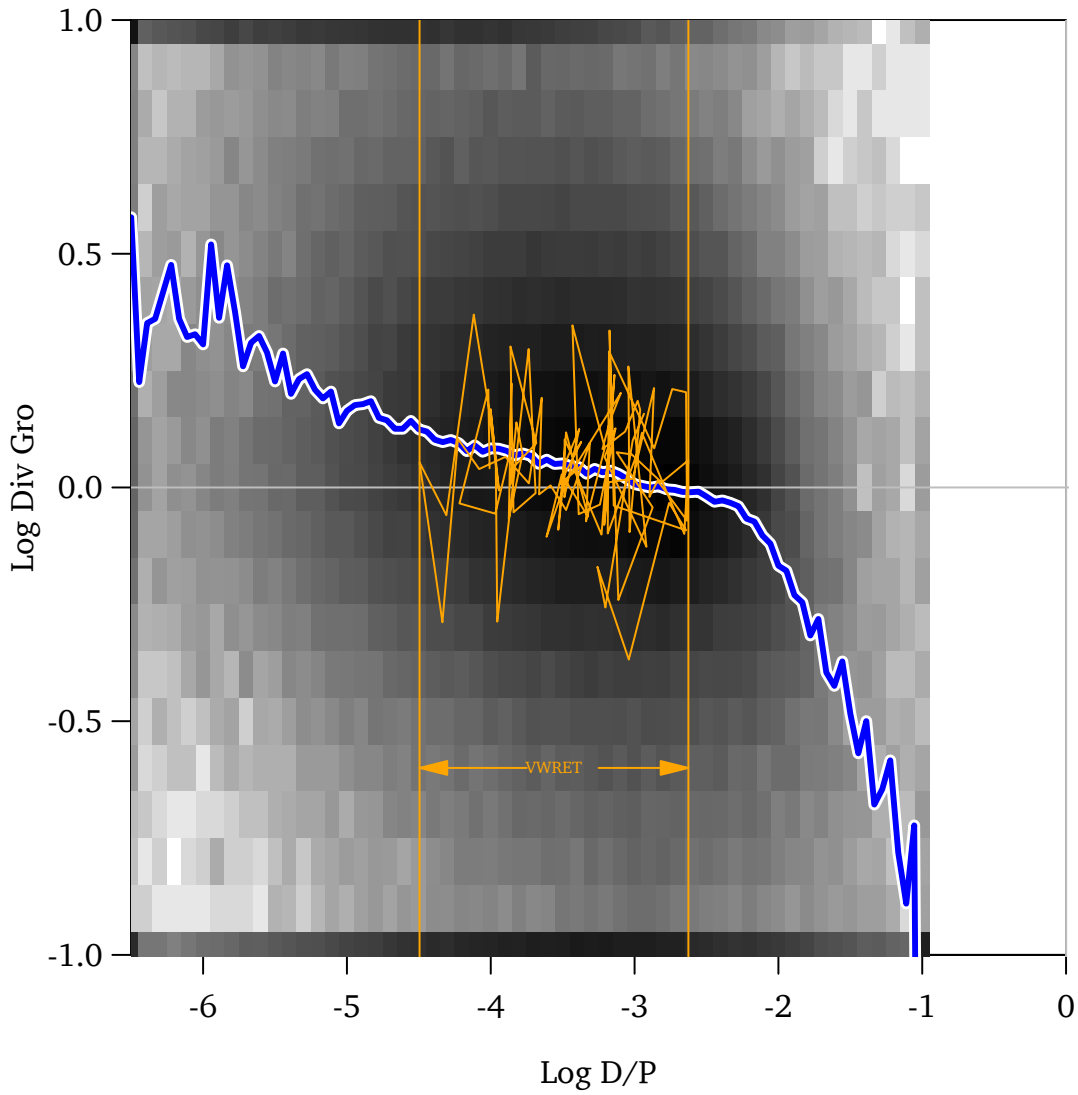


Explanations: This plot is based on the universe of all CRSP firms paying dividends from 1926-2024. The blue line are firms that omit dividends the following year. The red line are firms that disappear from CRSP the following year.

Interpretation: Firms paying dividends but at an extremely low level were more likely to disappear than they were to eliminate their dividends. Firms paying dividends at an extremely high level were more likely to eliminate their dividends than they were to disappear.

Source: 2plot-many-hist.R

Figure 10: Log D/P Ratio and Future Log Dividend Growth in Individual Firms

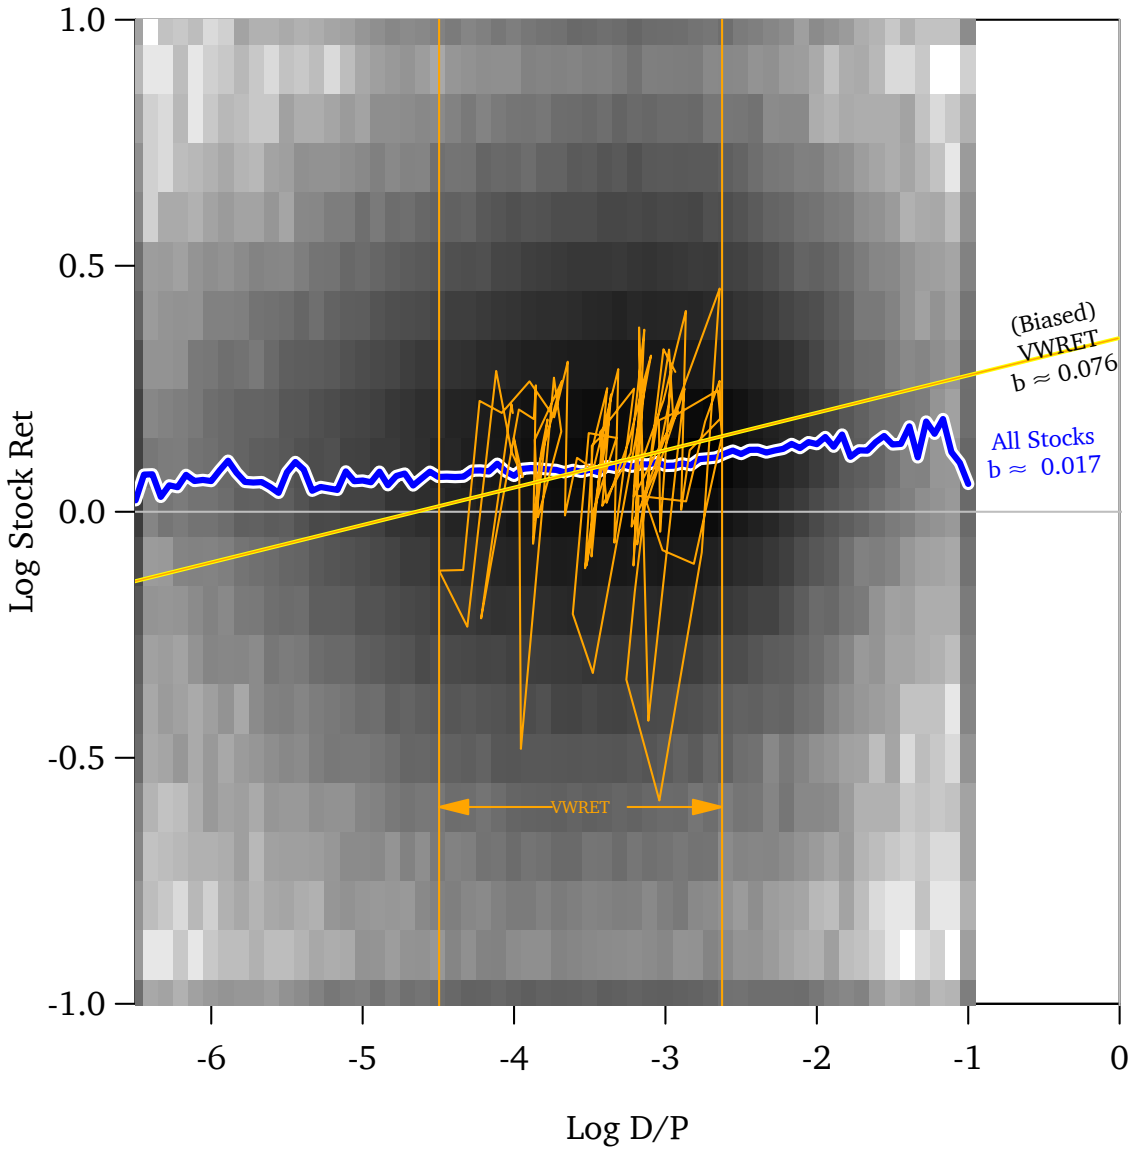


Explanations: The gray level indicates the number of observations in each rectangle. The blue-white line connects the *medians* (not means) in each vertical slice. Note that $\log(D/P)$ observations are winsorized at -1 (D/P of 37%), which is why the fitted blue-white curve declines steeply at its right edge (also at -1). (Unfortunately, there are too few observations with D/P ratios in excess of 37% to allow more reliable estimation above –1.) The orange spider graph are the observations for the value-weighted stock market index.

Interpretation: Higher dividend-price ratios solidly predicted lower dividend growth rates. The relationship was steeper the farther away from -3 to -2.5.

Source: 2plot-many-xy.R

Figure 11: Log D/P Ratio and Future Log Stock Return in Individual Firms

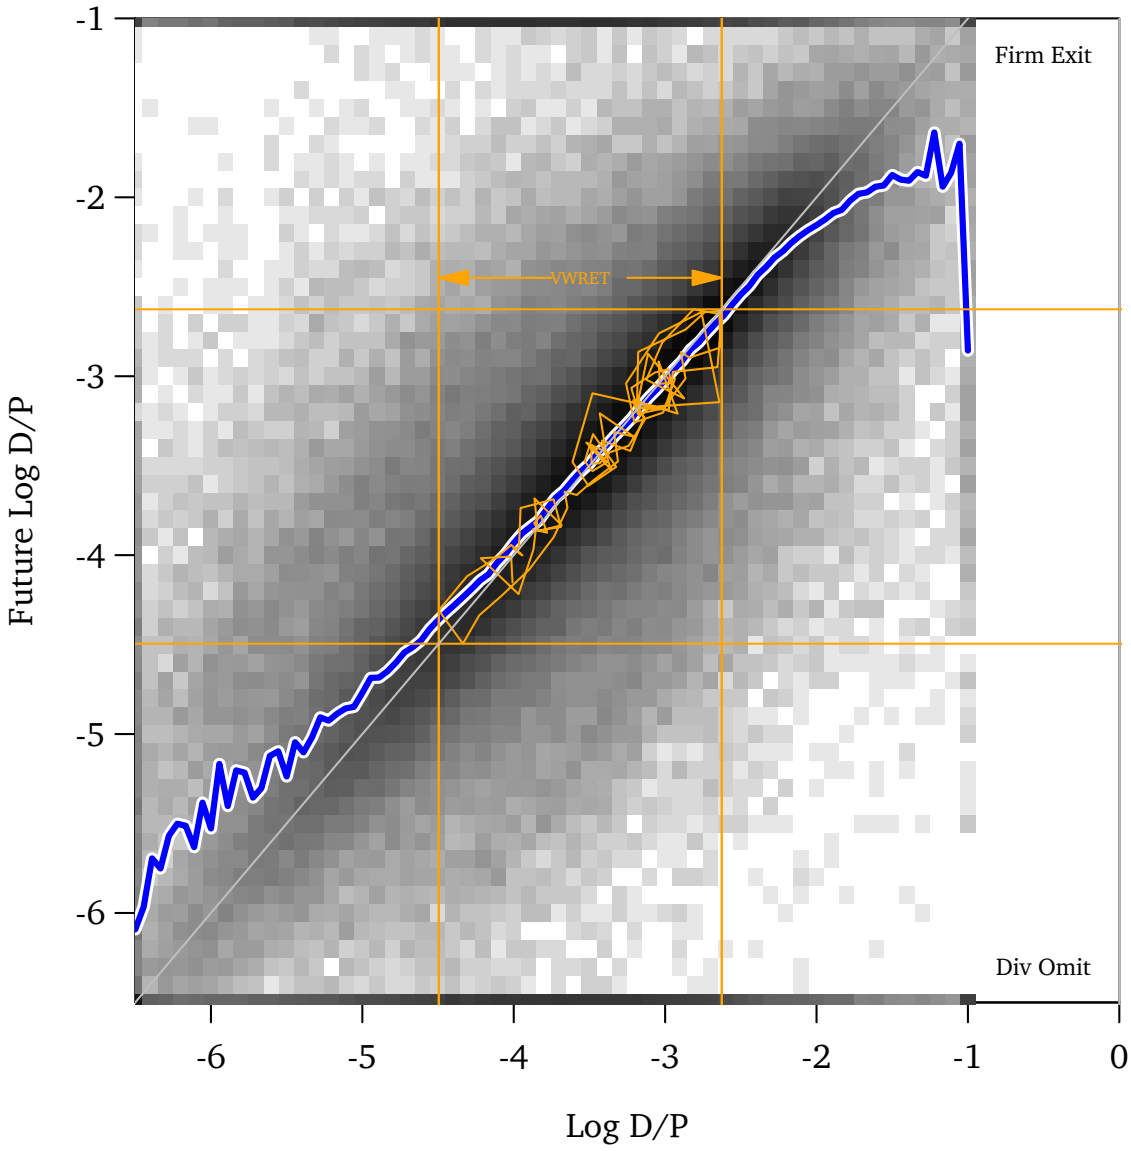


Explanations: This figure is the equivalent to Figure 10 for future log stock returns.

Interpretation: (1) The approximate linearity of the "all stocks" sample with a slope of about 0.02 is apparent. (2) This slope is flatter than its equivalent in Figure 10, except possibly from -2.5 to -3. (3) The (bias-uncorrected) slope for the value-weighted stock market index (of 0.07) was higher than that of the "all stocks" sample.

Source: 2plot-many-xy.R

Figure 12: Log D/P Ratio and Future D/P Ratio for Individual Dividend-Paying Firms

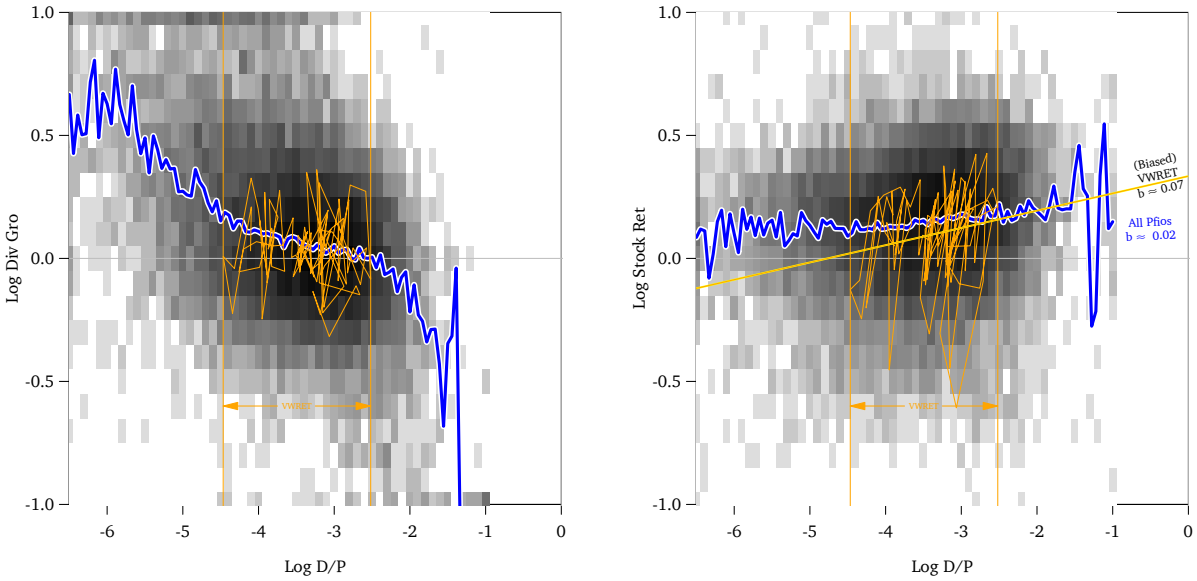


Explanations: This figure is the equivalent to Figure 10 for future log dividend price ratios. In addition, there is a diagonal line. To be included, a firm had to have paid dividends in year t, but not in t+1. A firm that ceased paying dividends is coded as below median (and winsorized to be shown at the lower edge of the graph); a firm that disappeared from CRSP is ceded as above median (and winsorized to be shown on the upper edge).

Interpretation: There is a distinct "S-shape" especially at D/P ratios above 8%.

Source: 2plot-many-xy.R

Figure 13: Log D/P Ratio and Future Variables for Portfolios



Explanations: This figure is the equivalent of Figures 10-11, but for portfolios formed from 10 stocks based on sorted consecutive lagged dividend-price ratio. The portfolios are formed anew in every year. These portfolios do not suffer from complete disappearances or dividend omissions. The OLS line slopes are -0.2774 (-0.2992 with year dummies) for dividend growth rates; and 0.0203 (0.0052 with year dummies) for stock returns.

Interpretation: The earlier inference remains robust: there is a strong association of lagged log dividend-price ratios with log dividend growth rates; and a mild linear association with stock returns.

Source: 2plot-pfios-xy.R

Table 6: Log D/P Ratio and Future Variables under Simulated Null

			Simulated		
		Empirical	Mean	Sdv	Pctl
Predicting r	Intercept	0.181	0.172	0.100	56.4%
	Slope	0.020	0.010	0.020	73.9%
Predicting g	Intercept	–0.083	–0.265	0.257	2.7%
	Slope	–0.282	–0.067	0.045	0.0%

Explanations: This is based on the 10-stock lag-dpr-sorted portfolios from Table 13. The empirical coefficients are from a grand panel regression (mixing different years). The simulations are from random (unconditional) (10 million) resamplings of stock returns and dividend growth rates, from which the dividend-price ratio series are reconstructed. Each resampling is for 99 years.

Interpretation: Among a broader set of portfolios, the dividend-price ratio forecasts only dividend-growth rates, not stock returns. However, survivorship bias is severe, as only 2% of simulations do not result in demise. Such high demise rates are not even unexpected or undesirable, because each of these portfolios contain only 10 stocks, and 1 in 10 stocks indeed disappear every year.

Table 7: MSCI Dividend Yields, March 2025

USA	Germany	UK	Australia	Japan	Brazil	China	South Korea	Thailand
1.3%	2.4%	3.7%	4.0%	2.2%	7.0%	2.4%	2.3%	3.7%

Source: MacroMicro.me, obtained 2025/03/20.

Superfluous Appendices

Alternative Process Modeling Possibilities:

- $g = \theta_0 + \theta_1 \cdot r + \theta_2 \cdot (d/p_{-1}) + \epsilon$
- Pull g and r from the loess curve.
- Pull g and r from perturbations instead of the normal distribution.

Figure 14: Flowchart

