

There Is No Excess Volatility Puzzle*

Andrew Atkeson Jonathan Heathcote Fabrizio Perri

UCLA

Federal Reserve Bank of Minneapolis

August 2024

Abstract

We present two valuation models which we use to account for the annual data on price per share and dividends per share for the CRSP Value-Weighted Index from 1929 to 2023. We show that it is a simple matter to account for these data based purely on a model of variation over time in the expected ratio of dividends per share to aggregate consumption under two conditions. First, investors must receive news shocks regarding the expected ratio of dividends per share to aggregate consumption in the long run. Second, the discount rate used to evaluate the impact of this news on the current price per share must be low. We use the approach of [Campbell and Shiller \(1987\)](#) and [Campbell and Shiller \(1988\)](#) to argue that the cash flow news in our model is not a stand-in for changes in expected returns: with our model parameters, returns are not predictable and price dividend spreads and ratios predict dividend growth at model-implied magnitudes. We illustrate which parameter choices account for differences between our results and prior findings in the literature. We conclude that the answer to Shiller's (1981) question "Do stock prices move too much to be justified by subsequent movements in dividends?" is "not necessarily."

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

1 Introduction

Shiller (1981) famously asked, “Do stock prices move too much to be justified by subsequent movements in dividends?” An important body of work in finance argues that the answer to this question is “yes.” See, for example, Leroy and Porter (1981), Campbell and Shiller (1987), Campbell and Shiller (1988), Cochrane (2011), and Shiller (2014).

In this paper, we call into question the conclusions drawn from this prior work. In particular, we present two simple valuation models which we use to account for the annual data on price per share and dividends per share for the CRSP Value-Weighted Index from 1929 to 2023. These two valuation models are based on two different models of the dynamics of agents’ expectations of future dividends per share valued at constant discount rates. We use these two models to show that it is a straightforward exercise to account for these aggregate stock market data based purely on a model of variation in the expected ratio of dividends per share to aggregate consumption over time under two conditions:

1. First, investors must receive news shocks regarding the expected ratio of dividends per share to aggregate consumption in the long run.
2. Second, the discount rate that determines the impact of this news on the current price per share must be low.

In contrast, we find that if these conditions are not satisfied, in particular the second condition regarding discounting, then the standard results pointing to stock prices being excessively volatile hold.

Thus, we conclude that the answer to the question of whether stock prices are excessively volatile depends on how valuation models are parameterized. As a result, we do not see our analysis as the final word on the question of what drives stock market volatility. More research is needed to discern what are the appropriate parameters to use in valuing the stock market. Instead, we see our paper as breathing new life into the old hypothesis that changing expectations of future dividends play an important and perhaps dominant role in driving aggregate stock market volatility.

1.1 A Decomposition of Asset Prices

We make use of the framework laid out in Campbell and Shiller (1987) to uncover the drivers of stock price volatility.

We study the following identity decomposing any asset price p_t into two components:

$$p_t = p_t^* + \phi_t \tag{1}$$

where

$$p_t^* \equiv \sum_{k=1}^{\infty} \beta^k \mathbb{E}_t d_{t+k} \quad (2)$$

and $\beta < 1$.

In this decomposition, p_t^* is the expected present value of dividends discounted at a constant rate β . We label p_t^* the “fundamental” component of price.¹ The component ϕ_t is simply the difference between the observed price p_t and the fundamental component p_t^* . We label ϕ_t the “residual” component of price.

In this paper, we use this decomposition of price into a fundamental and a residual component to ask whether one can account for observed fluctuations in prices $\{p_t\}_{t=0}^T$ based on a model of fluctuations in the fundamental component of price $\{p_t^*\}_{t=0}^T$ with a constant residual term $\phi_t = \bar{\phi}$. The alternative hypothesis is that variation over time in the residual term $\{\phi_t\}_{t=0}^T$ is required to account for the data on prices. In what follows, we refer to the hypothesis that one can account for fluctuations in observed prices p_t based on a model of fluctuations in the fundamental price p_t^* with constant $\phi_t = \bar{\phi}$ as the *Dividends Hypothesis*. We label the alternative hypothesis that variation over time in the residual term ϕ_t is required to account for observed prices the *Excess Volatility Hypothesis*.

A closely related framing of the question of what drives fluctuations in asset prices asks whether periods in which prices are high relative to an appropriate measure of dividends tend to be followed by higher than average dividend growth, or by lower than average returns. We now state a proposition that links the dynamics of our residual component of price ϕ_t to expected returns. This proposition illustrates that our framing of the question of what drives fluctuations in the asset price p_t is closely related to the common framing of that question that emphasizes changes in expected returns as a potential driver of fluctuations in price.

Proposition 1: Given equations (1) and (2) and a parameter β , we have

$$\beta \mathbb{E}_t \phi_{t+1} - \phi_t = \beta \mathbb{E}_t [p_{t+1} + d_{t+1}] - p_t \quad (3)$$

We refer to the expression on the right side of this equation as a *quasi-return*.

For a security whose price is always positive, this can be written as

$$\beta \mathbb{E}_t \left[\frac{\phi_{t+1}}{p_t} - \frac{\phi_t}{p_t} \right] + (\beta - 1) \frac{\phi_t}{p_t} = \beta \mathbb{E}_t R_{t+1} - 1 \quad (4)$$

where returns are given by

$$R_{t+1} \equiv \frac{p_{t+1} + d_{t+1}}{p_t}.$$

¹We will assume that β is small enough that the infinite sum in equation 2 is defined.

Equation (4) establishes a direct link between the dynamics of the residual component of price ϕ_t and the dynamics of expected returns. The proof of this proposition follows directly from the definition of the fundamental price and the Law of Iterated Expectations.²

To understand what we do in this paper, consider two versions of the Dividends Hypothesis.

First, suppose ϕ_t is constant at $\bar{\phi}$ and equal to zero. In this case, the observed price will always equal the fundamental price, and, from equation (4), the expected return to buying the asset $\mathbb{E}_t R_{t+1}$ will be constant and equal to $1/\beta$ and the expected quasi-return is zero. This is the case considered in [Campbell and Shiller \(1987\)](#) and [Campbell and Shiller \(1988\)](#) and the subsequent literature when evaluating the Dividends Hypothesis. In this case, the observation that realized returns on equity have been high suggests that a relatively low value for the parameter β should be used in modeling the fundamental price. We confirm in Section 9 of this paper the prior finding that the Dividends Hypothesis is rejected when stated in this form with $\bar{\phi} = 0$ and $1/\beta$ set to be consistent with the observed high returns on equity.

Second, suppose ϕ_t is constant at some negative value $\bar{\phi}$. This is the baseline case that we consider. Then Proposition 1 indicates that the expected quasi-return is $(\beta - 1)\bar{\phi}$

$$\mathbb{E}_t R_{t+1} = \frac{1}{\beta} \left(1 + \frac{(\beta - 1)\bar{\phi}}{p_t} \right). \quad (5)$$

In this scenario, the residual component $\bar{\phi}$ depresses observed prices relative to the fundamental price by a constant additive amount, which translates into a model-implied expected return for equity in excess of $1/\beta$. Thus, in this case, one can entertain the possibility that the β that enters into the calculation of fundamental price is high. With a high value of β news about dividends in the far future can have a large impact on current price. We use our two models of the dynamics of dividends to show that the Dividends Hypothesis fits the data very well when stated in this looser form.

²Start with the observation that

$$\beta p_{t+1}^* - p_t^* = \sum_{k=2}^{\infty} \beta^k [\mathbb{E}_{t+1} d_{t+k} - \mathbb{E}_t d_{t+k}] - \beta \mathbb{E}_t d_{t+1}$$

Thus, by the Law of Iterated Expectations

$$\beta \mathbb{E}_t p_{t+1}^* - p_t^* = -\beta \mathbb{E}_t d_{t+1}$$

Equation (1) then implies that

$$\beta \mathbb{E}_t p_{t+1} - p_t = \beta \mathbb{E}_t \phi_{t+1} - \phi_t - \beta \mathbb{E}_t d_{t+1}$$

which gives us equation (3).

1.2 Data

The data we use for our analysis is annual data from 1929 through 2023 on price per share for the CRSP Value-Weighted Total Market Index, which we denote by P_t and the corresponding measure of dividends per share for this index D_t . We scale these measures by annual data on Personal Consumption Expenditures (PCE) from the National Income and Product Accounts. Thus, in our applications below, p_t represents the ratio of price per share to PCE, with its value normalized to one in 1929. The variable d_t represents the ratio of dividends per share to PCE with this variable normalized so that p_t/d_t corresponds to the ratio of price per share to dividends per share at every date.

We discuss in Section 4 that our choice of consumption as a variable with which to scale dividends and price is motivated by an assumption that the price of a perpetual claim to aggregate consumption relative to current consumption is close to constant over time. This is equivalent to assuming in a Gordon growth model for aggregate consumption that movements in the discount rate relevant for a claim to aggregate consumption and movements in the expected growth rate of aggregate consumption offset, leaving the price-dividend ratio for such a claim constant over time. In this way, we do not need to assume that real interest rates and real growth rates are constant over time.

As we discuss in Section 4, for our baseline calibration we set $\beta/(1 - \beta) = 80$, which is consistent with estimates of this price-dividend ratio for a claim to aggregate consumption found by [Lustig, Van Nieuwerburgh, and Verdelhan \(2013\)](#). In Section 9 we show how our results vary with alternative values of β .

1.3 Return and Dividend Forecasting Regressions

To evaluate whether fluctuations in stock prices are driven by fluctuations in the residual component ϕ_t rather than fluctuations in the fundamental price p_t^* , we develop two parsimonious models of the dynamics of dividends that we use to construct two models of the dynamics of the fundamental component of price p_t^* . In the first of these two models, we focus on the dynamics of the level of dividends as in [Campbell and Shiller \(1987\)](#). We refer to this first model as our *linear model* of dividends. In the second, we focus on the dynamics of the log of dividends as in [Campbell and Shiller \(1988\)](#). We refer to this second model as our *log-linear model* of dividends.

In both of these models, we assume that dividends (or log dividends) follow ARIMA processes integrated of order one and thus have a [Beveridge and Nelson \(1981\)](#) decomposition into an unobserved trend component x_t (or $\log x_t$) and a transitory component $d_t - x_t$ (or $\log d_t - \log x_t$). By the definition of a Beveridge-Nelson decomposition, this trend component

corresponds to the expected value of dividends relative to consumption in the long run

$$x_t = \lim_{k \rightarrow \infty} \mathbb{E}_t d_{t+k}.$$

Thus, by construction, this trend component of dividends x_t is a martingale.

The question we seek to address is whether it is movements in this unobserved trend component of dividends x_t or fluctuations in the unobserved residual component of price ϕ_t that are driving price p_t . The observation that x_t is a martingale while ϕ_t need not be is the central implication of the model which we use in our testing of the Dividends Hypothesis. To that end, with each of our models of the dynamics of dividends, we develop a suite of forecasting regressions for quasi-returns (or log returns), dividend growth, and price growth to evaluate the fit of our model to the data under the Dividends Hypothesis relative to the alternative Excess Volatility Hypothesis. We introduce these forecasting regressions for the linear model in Section 2.1 and for the log-linear model in Section 7.2.

From our Proposition 1, one can interpret the first of our forecasting regressions with the linear model as analogous to the [Campbell and Shiller \(1988\)](#) regressions widely used in the prior literature. To see this connection, consider the implications of equation (3) from Proposition 1. We see from that equation that, given a choice of the parameter β , under the Dividends Hypothesis, this quasi-return should be constant over time. Thus, the quasi-return should not be predictable with any variable known at time t . Evidence that the quasi-return were predictable would favor the Excess Volatility Hypothesis.

For our first forecasting regression in our linear model we ask whether quasi returns in the data can be forecast using a valuation statistic constructed from data on price and dividends and shown to be stationary in [Campbell and Shiller \(1987\)](#) given by

$$p_t^T = p_t - \frac{\beta}{1 - \beta} d_t$$

This statistic is the analog in our linear model of a price-dividend ratio in a log-linear model. We discuss the pitfalls of using non-stationary valuation metrics to evaluate the Dividends Hypothesis in Section D.

The flip side of the prediction of the Dividends Hypothesis that the valuation statistic p_t^T should not forecast future quasi-returns is that p_t^T should forecast future growth in dividends, with coefficients consistent with the specific model of the transitory component of dividends being used. At the same time, proxies for the trend component x_t should be unforecastable. Results from all these regressions with our baseline parameters are reported in Section 5.

One appealing property of our linear model is that it is fully tractable without approximations. With our log-linear model, we must use first-order approximations to the model

solution to use the results in Proposition 1. We follow the approximations in [Campbell and Shiller \(1988\)](#). The central insight we use in constructing these forecasting regressions is that in a model with $\phi_t = \bar{\phi} < 0$, the appropriate measure of the log price-dividend ratio to use as the independent variable in evaluating the Dividends Hypothesis is $\log(p_t^*) - \log(d_t)$. This is because, by definition, the expected return on a claim to the fundamental price is constant over time. All variation in the log ratio of the fundamental price to dividends should be driven by changing expectations of future dividend growth.³

Unfortunately, the log ratio $\log(p_t^*) - \log(d_t)$ is not directly observed in data. Given a choice of the parameter $\bar{\phi}$, however, and taking the Dividends Hypothesis as the null hypothesis, this valuation statistic can be estimated as $\log(p_t - \bar{\phi}) - \log(d_t)$. We use this adjusted log price-dividend ratio in our forecasting regressions with this second, log-linear model. It is this adjustment that leads us to find different results to the prior literature (which imposes $\bar{\phi} = 0$) when we set $\bar{\phi} < 0$. In this case, both log-return and dividend growth forecasting regressions results are consistent with the Dividends Hypothesis and not with the Excess Volatility Hypothesis. Regression results in the second model with our baseline parameters are reported in Section [7.3](#).

1.4 Risk Adjustments to Price

We see our assumption that an additive constant risk adjustment $\bar{\phi}$ is present in stock prices as a key difference between our analysis and that in [Campbell and Shiller \(1987\)](#) and [Campbell and Shiller \(1988\)](#) and the literature that follows these papers. This raises the question then of what is the theoretical foundation for such a constant additive risk adjustment between observed prices p_t and fundamental prices p_t^* ?

In Section [4](#) and Appendix [C](#) we show that, in the context of our linear model, a constant additive risk adjustment between observed price p_t and fundamental price p_t^* arises naturally. Specifically, when the pricing kernel $M_{t,t+1}$ and consumption growth C_{t+1}/C_t are both conditionally lognormal while innovations to d_t are conditionally normal, then an application of Stein's Lemma delivers this additive risk adjustment. In contrast, if innovations to dividends are also lognormal, as in our second model of the dynamics of dividends, the risk adjustment between price and fundamental price is multiplicative rather than additive. Thus, we see the theoretical foundations for such an additive risk adjustment as a matter for future research. We return to discuss this question in our conclusion.

³Note from Proposition 1 that with $\bar{\phi} < 0$ there is time variation in conventionally measured expected returns due to movements in the ratio $\bar{\phi}/p_t$. But these movements in expected returns do not contribute to movements in price p_t .

1.5 Accounting for Observed Stock Prices

We take from our regression results the conclusion that with a high value of β and a correspondingly low value of $\bar{\phi}$, the data on stock prices and dividends are compatible with the Dividends Hypothesis. Under this hypothesis, we can use each of our two models of the dynamics of dividends to uncover the dynamics year-by-year of the unobserved trend component of dividends (x_t or $\log x_t$) needed to reconcile the model with observed stock prices and dividends.

We present results from our linear model in Section 6. In Figure 3, we see that with a high value of β , only modest variation over time in the long-run trend component of dividends x_t is required to account for observed stock prices. In Figure 4, we see that these fluctuations in the trend component x_t account for the large majority of observed fluctuations in stock prices. In Figure 5, we see that the expected returns implied by our linear model with $\bar{\phi} < 0$ do vary modestly over time. This variation in expected returns, however, does not contribute to volatility in prices p_t since this model has $\phi_t = \bar{\phi}$ by construction.

We present results from our log-linear model in Section 8. In Figure 6, we see again that with high β and low $\bar{\phi}$, only modest variation in the long-run trend in dividends $\log(x_t)$ is required to account for observed stock prices.

We note that our second, log-linear model of dividend dynamics nests a simple long run risks model in which the expected trend growth rate of log dividends is subject to persistent shocks. In Figure 7 we show the path of the expected trend growth rate of log dividends needed to account for observed prices. One striking feature of that figure is that the model-implied expected trend growth rate of dividends needed to account for the stock market boom of 2000 is in line with the ex-post realized growth of dividends over the next 20 years. Thus, in hindsight, it is not clear if the exuberance of the 2000 stock market boom was irrational.

1.6 Related Literature

To account for observed stock prices under the Dividends Hypothesis, we must have that the ratio of dividends per share relative to aggregate consumption is not stationary. A reader who is not familiar with the construction of measures of price per share and dividends per share for equity indices such as the CRSP Value-Weighted Index might wonder why one would have significant uncertainty about either the level or the growth rate of the ratio of dividends per share to aggregate consumption in the long run. We argue that much of the uncertainty about the long-run values of these two ratios is not driven by economic fundamentals but is instead driven by what are called *corporate actions* that impact the number of shares outstanding for firms included in the index. These corporate actions include entry of new firms into

the index, exit of firms from the index, mergers and acquisitions, new equity issuances and repurchases of shares by incumbent firms in the index.

As noted by [Dichev \(2006\)](#), [Boudoukh et al. \(2007\)](#), [Larrain and Yogo \(2008\)](#), [Gârleanu and Panageas \(2023\)](#), and [Davydiuk et al. \(2023\)](#), among others, these corporate actions generate large differences in the dynamics of dividends per share relative to aggregate dividends, in aggregate dividends relative to total cash flows to owners of equity, and in the dynamics of price per share relative to aggregate equity market capitalization as shown in [Figure 2](#). In particular, it appears from [Figure 2](#) that a large share of the movements in the ratio of dividends per share to aggregate consumption that we see in the data over the past century has been driven by the dynamics of corporate actions. Thus, a large portion of the uncertainty about the long-run ratio or growth rate of dividends per share to aggregate consumption, which is so important for the volatility of the ratio of price per share to aggregate consumption, likely reflects uncertainty about future corporate actions.

The impact of corporate actions on measures of stock market value and dividends extends to the observed ratio of price per share to dividends per share, which is typically used as a measure of the price-dividend ratio in empirical asset pricing. As argued by [Miller and Modigliani \(1961\)](#), total cash flows to equity holders are likely fundamental to valuing equity. They note that firms can use corporate actions to alter the dynamics of their dividends while holding fixed these total cash flows to equity holders. Thus, they argue that firms can alter the dynamics of their price-dividend ratio simply through changing their policy for paying dividends. [Larrain and Yogo \(2008\)](#) and [Atkeson, Heathcote, and Perri \(2024\)](#) argue that ratios of total payouts to value do not show the same trends as the ratio of price per share to dividends per share. We infer that corporate actions have played a large role in driving the observed dynamics of the ratio of price per share to dividends per share over the past century.

In terms of the literature, we see [Barsky and De Long \(1993\)](#) as the precursor closest to our work. That paper emphasizes the role of shocks to the dividend growth rate in the long run in accounting for the volatility of stock prices. [Bansal and Lundblad \(2002\)](#) and [Bansal and Yaron \(2004\)](#) also point to low-frequency movements in expected growth in dividends as an important source of changes in the price-dividend ratio for the aggregate stock market. We note that this type of long run risks model is nested in our second, log-linear model of the dynamics of dividends.

[Greenwald, Lettau, and Ludvigson \(2023\)](#) is an important precursor to our work in that it uses a model in which shocks to the ratio of earnings to output play a major role in accounting for the data on the evolution of the value of the stock market over time. We also follow them in using a valuation model to uncover the innovations needed to account for

the data. In contrast to their work, however, we use the standard data on price per share and dividends per share that are used in many asset pricing studies. Their work—as well as work by [Larrain and Yogo \(2008\)](#) and our companion paper, [Atkeson, Heathcote, and Perri \(2024\)](#)—develops valuation models using alternative data on cash flows to owners of U.S. corporations.

We note that it is standard in the asset pricing literature to build models with separate dynamics for dividends and aggregate consumption, and thus these models implicitly incorporate shocks to expectations of the ratio of dividends to consumption in the long-run. See, for example, [Campbell and Cochrane \(1999\)](#) and [Bansal and Yaron \(2004\)](#). But these alternative models do not appear to put these shocks to long-run expectations of the ratio of dividends per share to consumption at the center of their analysis. Given our results, it is unclear whether the other model elements those papers emphasize are needed to account for stock market data once one allows for news about the ratio of dividends to consumption in the long-run.

2 A First Linear Model Specification

Our first model specification is based on the following assumptions.

Assumption 1: We assume that the observed dividend process d_t is integrated of order one (an ARIMA process). Then, as [Beveridge and Nelson \(1981\)](#) showed, its dynamics can be decomposed into an unobserved trend component x_t defined as the long run expected value of dividends

$$x_t = \lim_{k \rightarrow \infty} \mathbb{E}_t d_{t+k},$$

and a transitory component, $d_t - x_t$. By construction, the process for x_t is a martingale. We model this trend component as

$$x_{t+1} = x_t + \epsilon_{x,t+1} \tag{6}$$

where $\epsilon_{x,t+1}$ should not be predictable by any variables known at t .

Note that we nest the case with d_t stationary with an assumption that $x_t = \bar{x}$ is constant.

Assumption 2: The ARMA representation of the transitory component $d_t - x_t$ can be quite general. To keep our model simple, we will assume that it is an AR(1). Specifically, we posit

$$(d_{t+1} - x_{t+1}) = \rho(d_t - x_t) + \epsilon_{d,t+1} \tag{7}$$

where $\rho < 1$. Note that the permanent and transitory shocks to dividends $\epsilon_{x,t+1}$ and $\epsilon_{d,t+1}$ may be correlated.

Assumptions 1 and 2 deliver the following model of the dynamics of expected dividends

$$\mathbb{E}_t d_{t+k} = \rho^k d_t + (1 - \rho^k) x_t. \quad (8)$$

Given these expectations, we can solve for the fundamental component of price as

$$p_t^* = \frac{\beta\rho}{1 - \beta\rho}(d_t - x_t) + \frac{\beta}{1 - \beta}x_t. \quad (9)$$

This expression is intuitive. If $d_t = x_t$, the fundamental value is just the present value of receiving x_t at every future date. If $d_t > x_t$, the fundamental value is boosted by the expectation that d_t will temporarily exceed x_t , with the differential decaying at rate ρ .

Note that innovations to the long-run expected value of dividends x_t have a larger impact on the fundamental price than equally sized innovations to the transitory component $d_t - x_t$, and this impact is increasing in the discount factor β . Thus, a positive innovation to $\epsilon_{x,t+1}$ coupled with an equal size negative innovation to $\epsilon_{d,t+1}$ will boost the fundamental price without any contemporaneous change in dividends.

Assumptions 1, and 2 give the following model of prices under both hypotheses that we consider:

$$p_t = \phi_t + \frac{\beta\rho}{1 - \beta\rho}(d_t - x_t) + \frac{\beta}{1 - \beta}x_t. \quad (10)$$

The corresponding expected growth in the price is given by

$$\mathbb{E}_t p_{t+1} - p_t = \mathbb{E}_t \phi_{t+1} - \phi_t + \rho \frac{\beta\rho}{1 - \beta\rho}(d_t - x_t). \quad (11)$$

Thus, realized price growth has an unpredictable random walk component inherited from time variation in the long-run expected value for dividends x_t , and two additional components due to (i) temporary deviations of dividends from their long run expected value ($d_t - x_t$) and (ii) movements in ϕ_t which, from equation (4), correspond to time variation in the expected rate of return. Expected price growth retains only the last two of these driving forces.

2.1 Forecasting Regressions

We now seek to test the Dividends Hypothesis, which is the hypothesis that assets are priced as the sum of (i) the expected discounted present value of dividends given a constant discount factor β (the fundamental price) plus (ii) a constant additive risk adjustment term $\bar{\phi}$. We start by constructing a measure of returns between t and $t + s$, $r_{t,t+s}$, which we label *quasi-*

returns, defined as

$$r_{t,t+s} \equiv \sum_{k=1}^s \beta^k d_{t+k} + \beta^s p_{t+s} - p_t. \quad (12)$$

From equations (1) and (2), we have that

$$\mathbb{E}_t r_{t,t+s} = \beta^s \mathbb{E}_t \phi_{t+s} - \phi_t. \quad (13)$$

Thus, under the Dividends Hypothesis, expected quasi-returns over horizon s should be constant and equal to $(\beta^s - 1)\bar{\phi}$, and time variation in realized quasi-returns should not be predictable. Hence, the first question we ask is whether $r_{t,t+s}$ can be forecast in the data. If quasi returns can be predicted, that would provide evidence against the Dividends Hypothesis, and specifically against the hypothesis that the residual term ϕ_t is constant.

Conceptually, one could use any variable known at time t in forecasting $r_{t,t+s}$. Here we consider a more limited forecasting exercise. Specifically, we define a valuation measure – a *price-dividend spread* in the language of [Campbell and Shiller \(1987\)](#) – that eliminates the valuation impact of the trend dividend component x_t . This price-dividend spread is defined as

$$p_t^T \equiv p_t - \frac{\beta}{1-\beta} d_t. \quad (14)$$

Note that p_t^T is measurable from the data given knowledge only of the parameter β . As shown in [Campbell and Shiller \(1987\)](#), this valuation metric is a stationary random variable under the Dividends Hypothesis given a general specification of the dynamics of the transitory component of dividends $d_t - x_t$. To see this point in our specific model, observe that, given equation (10), we have

$$p_t^T = \phi_t - \Gamma(d_t - x_t) \quad (15)$$

where

$$\Gamma \equiv \frac{\beta}{1-\beta} - \frac{\beta\rho}{1-\beta\rho}. \quad (16)$$

Thus, under the Dividends Hypothesis and our Assumption 1 regarding the dynamics of dividends, this valuation metric p_t^T is a stationary random variable.

We evaluate the predictability of $r_{t,t+s}$ by p_t^T using OLS regressions of the form

$$r_{t,t+s} = \alpha_{r,s} + \gamma_{r,s} p_t^T + error_{r,t+s} \quad (\text{regression 1})$$

and ask whether the estimated slope coefficients $\hat{\gamma}_{r,s}$ at different horizons s equal zero as predicted by the Dividends Hypothesis.

Note that given a value of the parameter β , the statistic $r_{t,t+s}$ and the valuation metric p_t^T

can be constructed directly from data on price p_t and dividends d_t . Thus, to explore whether quasi-returns are predictable, one does not need to specify a process for dividends.

We next go beyond this general test of the Dividends Hypothesis by exploring whether we can reject the explicit process for dividends outlined above, according to which dividends have a trend component x_t that is a martingale, and a transitory component $d_t - x_t$ that follows an AR(1) process. Once we posit this particular model, we are effectively testing a stricter version of the Dividends Hypothesis. The additional restriction that now comes into play is that the fundamental component of price reflects the expected present value of dividends when those expectations are formed by rational agents who take as given the dividend process specified in equations (6) and (7). Given that process, the expected present value of dividends is given by equation (9).

From that fundamental price expression, it is immediate that given values for the parameters β , ρ and ϕ , one can always infer a time series $\{x_t\}$ that perfectly replicates the observed sequence for $\{p_t\}$. But the stricter version of the Dividends Hypothesis asserts that the resulting x_t series is a martingale (and also that the series $d_t - x_t$ is an AR(1) process). If the series for x_t that replicates the observed price history appears to be forecastable, that would be evidence against the posited process for dividends and thus against the stricter version of the Dividends Hypothesis. And if x_t appears forecastable, then equation (9) is no longer interpretable as the rational expectations value for the present value of dividends.

Given the pricing equation (10), expected price growth between t and $t + s$ is given by

$$\mathbb{E}_t p_{t+s} - p_t = \frac{\beta}{1 - \beta} (\mathbb{E}_t x_{t+s} - x_t) + \frac{\beta\rho}{1 - \beta\rho} [\mathbb{E}_t (d_{t+s} - x_{t+s}) - (d_t - x_t)] + \mathbb{E}_t \phi_{t+s} - \phi_t$$

Under the Dividends Hypothesis, $\mathbb{E}_t \phi_{t+s} = \phi_t = \bar{\phi}$, and given our model for dividends, $\mathbb{E}_t x_{t+s} - x_t$. Thus, the Dividends Hypothesis implies the following relationship between expected price growth and expected dividend growth:

$$\mathbb{E}_t p_{t+s} - p_t = \frac{\beta\rho}{1 - \beta\rho} [\mathbb{E}_t d_{t+s} - d_t]. \quad (17)$$

Thus, under the Dividends Hypothesis, if investors expect price growth, it must be that they expect proportional dividend growth. If the Dividends Hypothesis were false, then investors could expect price growth without dividend growth if they expect ϕ_t to rise (i.e. they expect high returns).

Given our AR(1) process for dividends, we can solve for expected dividend growth under

the Dividends Hypothesis:

$$\begin{aligned}
\mathbb{E}_t d_{t+s} - d_t &= \mathbb{E}_t(d_{t+s} - x_{t+s}) - (d_t - x_t) \\
&= (\rho^s - 1)(d_t - x_t) \\
&= (\rho^s - 1) \left(\frac{-(p_t - \frac{\beta}{1-\beta}d_t - \bar{\phi})}{\Gamma} \right)
\end{aligned}$$

where the first equality follows from equation (6), the second from (7), and the third from (10). Thus, we can estimate the parameter ρ by running regressions of the form

$$d_{t+s} - d_t = \alpha_{d,s} + \gamma_{d,s} p_t^T + \varepsilon_{d,s}. \quad (\text{regression 2})$$

The estimate $\hat{\gamma}_{d,s}$ identifies $(1 - \rho^s)/\Gamma$ which, given a value for β , delivers an estimate for the persistence parameter ρ .

Our second test of Dividends Hypothesis will be to test equation (17) by exploring whether price growth and dividend growth (scaled by $\beta\rho/(1 - \beta\rho)$) are equally predictable. Thus, we will run a forecasting regression of the form:

$$p_{t+s} - p_t - \frac{\beta\rho}{1 - \beta\rho}(d_{t+s} - d_t) = \alpha_{p,s} + \gamma_{p,s} p_t^T + \varepsilon_{p,s} \quad (\text{regression 3})$$

and check whether $\hat{\gamma}_{p,s} = 0$. Note that this test is equivalent to testing whether growth in the spread $p_t - \frac{\beta\rho}{1-\beta\rho}d_t$ is predictable. That spread, given our pricing equation (10), is given by $\phi_t + \Gamma x_t$. Under the stricter version of the Dividends hypothesis, $\phi_t = \bar{\phi}$ and x_t is a martingale, and thus growth in the spread $p_t - \frac{\beta\rho}{1-\beta\rho}d_t$ is supposed to be unpredictable. Thus, if we find, $\hat{\gamma}_{p,s} \neq 0$, that would indicate that the spread is predictable, which would be evidence against the Dividends Hypothesis null.

If we do find parameters β and ρ such that our regression evidence is consistent with the Dividends Hypothesis, then, given an estimate of the constant $\bar{\phi}$, we can construct the estimates of the unobserved values of $\{x_t\}_{t=0}^T$ directly from the data on prices $\{p_t\}_{t=0}^T$ and dividends $\{d_t\}_{t=0}^T$ using equation (10) with the restriction that $\phi_t = \bar{\phi}$. In this way, we can offer an accounting of forces driving movements in stock prices year-by-year from 1929 to 2023.

3 Our Data

The data we use for this analysis is annual data from 1929 through 2023 on price per share for the CRSP Value-Weighted Total Market Index, which we denote by P_t and the corresponding




Figure 1: Left panel: The ratio of price per share for the CRSP Value-Weighted Total Market Index to Personal Consumption Expenditures 1929-2023. The value of this ratio in 1929 is normalized to one. Right panel: The ratio of dividends per share for the CRSP Value-Weighted Total Market Index to Personal Consumption Expenditures 1929-2023. This series for dividends per share is normalized so that the ratio of the two series equals the ratio of dividends per share to price per share at every date.

measure of dividends per share for this index D_t . We scale these measures by annual data on Personal Consumption Expenditures (PCE) from the National Income and Product Accounts. Thus, in our applications below, p_t represents the ratio of price per share to PCE, with its value normalized to one in 1929. The variable d_t represents the ratio of dividends per share to PCE with this variable normalized so that p_t/d_t corresponds to the ratio of price per share to dividends per share at every date.

We plot these two series in the left and right panels of Figure 1. We see that p_t is quite volatile and d_t has low frequency movements that we model as a process that is integrated of order one. We discuss the construction of these data in greater detail in Appendix B.

One striking feature of the data on the ratio of dividends per share to PCE shown in the right panel of Figure 1 is that it does not appear to be stationary. Instead, it shows a marked downward trend since 1929. The first key assumption in our valuation models is that investors do not have a fixed expectation of the value of this ratio in the long run. Instead, they receive news each period that leads them to revise their expectation of the long-run value of this ratio. We argue that uncertainty about the long-run value of this ratio is plausible as a matter of econometrics given its historical path as shown in the figure – it is not at all clear what long-run value the series is converging to.

There are also multiple economic reasons why the ratio of dividends per share to PCE

might vary in the long run.

First, it may be the case that the total cash flows paid to owners of U.S. corporations relative to consumption might vary over time. [Greenwald, Lettau, and Ludvigson \(2023\)](#) and [Atkeson, Heathcote, and Perri \(2024\)](#) argue that this is indeed the case.

Second, the fraction of economic activity carried out in publicly-traded corporations relative to all corporations might vary over time. This is likely the case as well over the course of the past century.

Third, the methodology used in the construction of the indices of price per share and dividends per share implies that these measures track neither the total value of the stock market as measured by total market capitalization of the stocks in the index nor the total value of cash payouts to owners of these equities. Instead, the ratio of the index of price per share to total market capitalization varies over time as a result of a large number of actions that result in changes in the number of shares outstanding for the firms in the index. These corporate actions include entry and exit of firms in public markets, mergers of firms, issuance of new shares or repurchases of shares by incumbent firms, and so on. What these corporate actions imply is that an investor who maintained a portfolio to track the CRSP Value-Weighted Index would hold a share of the total market that varies over time. We give further details on these points in [Appendix B](#).

We show the variation of the ratio of the index of price per share to total market capitalization for the CRSP Value-Weighted Index over the period 1929-2023 in [Figure 2](#). This ratio represents the fraction of the total stock market held by an investor tracking the CRSP Value-Weighted Index. In this figure, we normalized 1929 price per share so that the fraction is equal to one in 1929.




Figure 2: The fraction of the total market capitalization of the stocks in the CRSP Value-Weighted Total Market Index held by an investor tracking that index, 1929-2023. The initial value of this fraction has been normalized to one.

We see in this figure a sharp downward trend in the share of the total market held by an index investor over this time period. An investor managing his or her portfolio to track the CRSP Value-Weighted Index would end up holding a shrinking share of the total stock market because he or she would not be purchasing the new shares being issued on net from corporate actions.⁴ Thus, it is natural that the ratio of dividends per share to consumption would also fall over time, as an investor tracking the index of price per share would have claims to a shrinking share of total dividends. This figure also makes clear that part of investors' uncertainty regarding the long-run value of the ratio of dividends per share to consumption is also driven by uncertainty regarding the future course of corporate actions.

We capture this uncertainty about the long-run ratio of dividends per share to PCE in our model variable x_t .

4 Choosing β , ρ , and $\bar{\phi}$

The parameter β plays an important role in our analysis. It defines the cointegrating vector linking prices and dividends in equation (14). As we will show, the value for β also plays an important role in determining whether or not prices appear excessively volatile.

A straightforward way to discipline the choice for β is to note that $\beta/(1-\beta)$ corresponds

⁴Gârleanu and Panageas (2023) document that most of this decline in the share of the total market held by an index investor is driven by the entry of new firms into public markets through initial public offerings.

to the price of a claim to future aggregate consumption in units of current consumption. On this basis, we set our baseline value to $\beta = 80/81$ to target a value for that relative price of 80, as estimated by [Lustig, Van Nieuwerburgh, and Verdelhan \(2013\)](#).

We now expand on these points.

We begin with the standard valuation equation for the level of the stock index before rescaling:

$$P_t = \sum_{k=1}^{\infty} \mathbb{E}_t [M_{t,t+k} D_{t+k}], \quad (18)$$

where $M_{t,t+k}$ is the nominal pricing kernel between periods t and $t+k$, P_t is the stock index and D_t is the index of dividends per share.

We work with ratios of price per share and dividends per share to PCE, using lower case variables p_t and d_t to represent these ratios. We therefore rewrite this pricing equation (18) as

$$p_t = \sum_{k=1}^{\infty} \mathbb{E}_t \left[M_{t,t+k} \frac{C_{t+k}}{C_t} d_{t+k} \right].$$

Using the result that the expectation of a product of two random variables is the product of the expectations plus the covariance between these variables, we have

$$p_t = \sum_{k=1}^{\infty} \mathbb{E}_t \left[M_{t,t+k} \frac{C_{t+k}}{C_t} \right] \mathbb{E}_t d_{t+k} + \sum_{k=1}^{\infty} \text{Cov}_t \left(M_{t,t+k} \frac{C_{t+k}}{C_t}, d_{t+k} \right). \quad (19)$$

Note that the term

$$p_{C,t}^{(k)} \equiv \mathbb{E}_t \left[M_{t,t+k} \frac{C_{t+k}}{C_t} \right]$$

is the price at t of a claim to aggregate consumption at delivered at $t+k$ relative to aggregate consumption at t . We define the price at t of a claim to aggregate consumption in perpetuity relative to the current level of aggregate consumption as

$$p_{C,t} \equiv \sum_{k=1}^{\infty} p_{C,t}^{(k)}.$$

The terms

$$H_t^{(k)} \equiv \text{Cov}_t \left(M_{t,t+k} \frac{C_{t+k}}{C_t}, d_{t+k} \right) \quad (20)$$

constitute a risk adjustment to the price of claims to dividends due to risk associated with fluctuations in the ratio d_t .

Our choice of β and our choice of consumption as a variable with which to scale dividends and price is motivated by an assumption that the prices of consumption claims relative to

current consumption are constant over time, or

$$p_{C,t}^{(k)} = \beta^k.$$

This is equivalent to assuming in a Gordon growth model for aggregate consumption that movements in the discount rate relevant for a claim to aggregate consumption and movements in the expected growth rate of aggregate consumption offset, leaving the price-dividend ratio for such a claim constant over time. In this way, we do not need to assume that real interest rates and real growth rates are constant over time.

Given that assumption, the first term in equation (19) corresponds to the fundamental price p_t^* defined in equation (2). In our baseline calibration we set the price of a perpetual claim to consumption relative to current consumption to $p_{C,t} = \sum_{k=1}^{\infty} \beta^k = \frac{\beta}{1-\beta} = 80$.

For this pricing model to be consistent with the Dividends Hypothesis, we also require that the risk adjustment terms $H_t^{(k)}$ be constant over time, with

$$\bar{\phi} = \sum_{k=1}^{\infty} H^{(k)}.$$

In Appendix C we show that the $H_t^{(k)}$ terms are indeed time invariant under the assumptions that the pricing kernel $M_{t,t+1}$ and consumption growth C_{t+1}/C_t are both conditionally log-normal while innovations to d_t are conditionally normal. This configuration of shocks holds under our Assumptions 1 and 2 if we impose that $\epsilon_{d,t+1}$ and $\epsilon_{x,t+1}$ are both normal. We see this result as a start towards a micro-foundation for our assumption that ϕ_t is constant over time.

To choose $\bar{\phi}$ given β , we compare our model's implications to the data on mean quasi-returns at different horizons. Table 1 reports mean realized quasi-returns at various horizons, together with the standard errors for these estimates of mean quasi-returns. Observe from equations (3) and (13) that with $\bar{\phi} = 0$, our model would imply that expected quasi-returns at all horizons should be zero. At a horizon of $s = 1$ years, mean realized quasi-returns in the data are $\mathbb{E}r_{t,t+1} = 0.0101$ which is within 0.86 standard errors of zero. Longer horizon mean realized quasi-returns, however, such as for $s = 5, 10$, and 15 , are all significantly greater than zero. Thus, to come close to matching the data on mean realized quasi-returns at longer horizons with a value of $\beta = 80/81$, we must choose a value of $\bar{\phi}$ that is negative. In the first row of Table 1 we show model-implied expected quasi-returns $(\beta^s - 1)\bar{\phi}$ with $\beta = 80/81$ and $\bar{\phi} = -0.9$. This parameter configuration replicates mean observed quasi-returns closely at all horizons.

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$(\beta^s - 1)\bar{\phi}$	0.0111	0.0542	0.1051	0.1530
mean($r_{t,t+s}$)	0.0101	0.0658	0.1096	0.1530
S.E.	(0.0118)	(0.0194)	(0.0246)	(0.0279)

Table 1: Realized mean and model-implied quasi-returns at horizon s with $\beta = 80/81$ and $\bar{\phi} = -0.9$.

Note that, in the data, the mean arithmetic realized return on equity in excess of consumption growth is 1.0530, with a standard error of that mean of 0.0199. Using equation (4), with our baseline parameters, the mean one-period ahead expected return on equity in excess of consumption growth is 1.0360. Thus, we see our model-implied expected arithmetic returns with these parameters as being close to mean realized returns in the data.

Given our choice of β , we estimate ρ using [regression 2](#). Table 2 shows the estimates $\hat{\gamma}_{d,s}$ for different values for the horizon s . The row labelled $(1 - \rho)^s/\Gamma$ shows the Dividends Hypothesis predicted value for that expression evaluated at $\rho = 0.96273$ and $\beta = 80/81$, which replicates the regression estimate at horizon $s = 5$. Note that regression predicted dividend growth matches theory predicted growth quite well at all horizons given these parameter values, offering support for our AR(1) model for the transitory component of dividends.

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$(1 - \rho)^s/\Gamma$	0.0006144	0.002852	0.005209	0.007160
$\hat{\gamma}_{d,s}$	0.00093581	0.0028518	0.005122	0.0067141
S.E.	(0.00047511)	(0.00069836)	(0.00071879)	(0.00068034)
t-Stat	1.9697	4.0836	7.1258	9.8687
R^2	0.0405	0.159	0.38	0.555

Table 2: Estimates from regressions of the form in [regression 2](#) with $\beta = 80/81$.

In what follows, we use as our baseline parameters $\beta = 80/81$, $\bar{\phi} = -0.9$, and $\rho = 0.96273$.

5 Forecasting Regression Results

In this section, we report results from forecasting quasi returns ([regression 1](#)) and price growth ([regression 3](#)) using the stationary price dividend spread p_t^T as the predictor under our linear model of dividends with parameters $\beta = 80/81$ and $\rho = 0.96273$. Note that we have used the dividend growth regression ([regression 2](#)) to estimate the parameter ρ . Thus, the model mechanically fits the relationship between dividend growth and p_t^T .

We begin with results from [regression 1](#) at various horizons s shown in [Table 3](#). As we see in this table, there is no evidence that the quasi return $r_{t,t+s}$ can be forecast with the price-dividend spread p_t^T . The estimated coefficients $\hat{\gamma}_{r,s}$ and the regression R^2 are all very close to zero. Thus, these regressions do not reject the central property of the Dividends Hypothesis, which is that ϕ_t and expected quasi-returns are constant.

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$\hat{\gamma}_{r,s}$	0.0078894	-0.00062418	0.014361	0.0005422
S.E.	(0.019114)	(0.031424)	(0.039455)	(0.044795)
t-Stat	0.41275	-0.019863	0.36398	0.012104
R^2	0.00185	4.48e-06	0.00159	1.88e-06

Table 3: Estimates from regressions of the form in [regression 1](#) with $\beta = 80/81$.

We next consider results from [regression 3](#) with $\beta = 80/81$ and $\rho = 0.96273$. We present these results in [Table 4](#). Recall that under the Dividends Hypothesis, the coefficient on p_t^T in this price-dividend spread growth regression should equal 0. The table indicates that that is indeed the case. Thus, we do not reject the property of the stricter version of the Dividends Hypothesis that the series x_t constructed from the data given β, ρ and $\bar{\phi}$ is a martingale.

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$\hat{\gamma}_{p,s}$	0.0014502	-0.0015068	0.018158	0.013945
S.E.	(0.014934)	(0.026657)	(0.03477)	(0.040873)
t-Stat	0.09711	-0.056524	0.52223	0.34118
R^2	0.000102	3.63e-05	0.00328	0.00149

Table 4: Estimates from regressions of the form in [regression 3](#) with $\beta = 80/81$ and $\rho = 0.96273$.

In these forecasting regressions we have used the valuation metric p_t^T as our explanatory variable. As shown in [equation \(15\)](#), given our model, p_t^T is comprised of a scaled version of the impact of the stationary component of dividends on price and the impact of the residual term ϕ_t on price. Under the Dividends Hypothesis, this variable is stationary.

One might be tempted to use other valuation metrics as the explanatory variable in these regressions. But under the Dividends Hypothesis, x_t is non-stationary, and thus any valuation metric that contains x_t will be non-stationary. At the same time, the realized changes in variables being forecast in [regression 1](#) and [regression 3](#) contain realized values of future changes in $x_{t+s} - x_t$. As a result, if one estimates [regression 1](#) or [regression 3](#) using a valuation metric including x_t as an explanatory variable, one will find spurious results that

appear to contradict the Dividends Hypothesis, but which are really uninformative about the truth of that hypothesis. We review this issue in Appendix D.

6 Accounting for Observed Stock Prices

Given that our forecasting regression evidence from our linear model of dividends is consistent with the Dividends Hypothesis, we now use our parameters $\beta = 80/81$, $\rho = 0.96273$, and $\bar{\phi} = -0.9$ and equation (10) with $\phi_t = \bar{\phi}$ to estimate the sequence of unobserved trend variables $\{x_t\}_{t=0}^T$ that account for the price and dividend data annually from 1929 to the present.

We show the resulting series for d_t in blue and x_t in red in Figure 3.




Figure 3: Blue Line: The ratio of dividends per share for the CRSP Value-Weighted Total Market Index to PCE (d_t), 1929-2023. Red Line: The expected long-run ratio of dividends per share for the CRSP Value-Weighted Total Market Index to PCE, x_t , that rationalizes the observed price per share of this index using equation (10), 1929-2023.

The transitory component of dividends $d_t - x_t$ is given by the difference between the blue and red lines in the figure while the permanent component of dividends x_t is given by the red line. Note that with a high value of β , only modest variation in x_t over time is needed to account for the observed volatility of p_t .

Figure 3 indicates that from the start of our sample until around 1970, investors were expecting dividends per share relative to consumption to decline over time (the red line is below the blue line). Were these expectations reasonable? Well, dividends per share relative to consumption did in fact decline steadily over this period! As described in Gârleanu and Panageas (2023), this is what one would have expected in an economy in which new firms continuously displace old ones. From 1980 onward, the model indicates that investors were expecting future ratios of dividends per share to consumption to exceed the current level (the red line is above the blue), with this gap being especially pronounced around the dot com stock boom in 2000. And in fact, from around 2000 onward, dividends per share relative to consumption have been generally rising, suggesting this optimism about future dividends per share was more than irrational exuberance.

By construction, under this Dividends Hypothesis version of the model, movements in the fundamental price p_t^* account for all the movements in p_t observed in the data. We can use equation (9) to break p_t^* into a component due to the transitory movements in dividends and a component due to the permanent component in dividends, here $\frac{\beta}{1-\beta}x_t$. In Figure 4, we show the data on p_t in blue and the path for $\frac{\beta}{1-\beta}x_t + \bar{\phi}$ in red. Given our model, movements in the difference between these two lines are accounted for by transitory fluctuations in dividends given by $\frac{\beta\rho}{1-\beta\rho}(d_t - x_t)$. We see in this figure that, under the Dividends Hypothesis, movements in the permanent component of dividends account for most of the fluctuations in price p_t .

Given the finding in Figure 4 that the large majority of fluctuation in the stock price p_t are accounted for by fluctuations in the unobserved trend component x_t of dividends under our Dividends Hypothesis, one might be tempted to ask how one can know that these fluctuations in the unobserved x_t are not in fact standing in for fluctuations in the residual component of price ϕ_t . As we argued in the previous section, this alternative hypothesis is inconsistent with the forecasting regression evidence that we have presented. Thus, we take this regression evidence as favoring the hypothesis that the fluctuations in the red line are, in fact, driven by fluctuations in expectations of the long-run ratio of dividends per share to consumption rather than by fluctuations in expected returns.

Finally, we use our parameter values and equation (5) to compute our model's implications for the dynamics of the expected return on equity one period ahead. Note that since we have normalized price p_t and dividends d_t by PCE consumption, these model implications are for the expected nominal return on equity in excess of nominal consumption growth. We present the results of this calculation of model-implied expected returns in Figure 5. The series that is plotted is net expected returns in percentage points, or $100(\mathbb{E}_t R_{t+1} - 1)$.




Figure 4: Blue Line: The ratio of price per share for the CRSP Value-Weighted Total Market Index to PCE (p_t), 1929-2023. Red Line: The component of price corresponding to fluctuations in the trend x_t for dividends given by $\frac{\beta}{1-\beta}x_t + \bar{\phi}$, 1929-2023.




Figure 5: Blue Line: model implied expected returns on equity in excess of aggregate consumption growth yearly 1929-2023. The formula for model-implied expected returns is $\mathbb{E}_t R_{t+1} = \frac{1}{\beta} \left(1 + \frac{(\beta-1)\bar{\phi}}{p_t^* + \bar{\phi}} \right)$. The series that is plotted is net expected returns in percentage points, or $100(\mathbb{E}_t R_{t+1} - 1)$.

We see in this figure that our model implies modest variation over time in the expected return on equity in excess of consumption growth. Since this variation in returns corresponds to a constant value of $\phi_t = \bar{\phi}$, it does not contribute to volatility in stock prices.

7 A Second Log Linear Model Specification

In our first model specification, we assumed that the level of dividends per share relative to PCE is integrated of order one and has normal innovations to both the trend and transitory components. We refer to this model as our linear model of dividends relative to PCE.

In this section, we consider an alternative model of the dynamics of the logarithm of the ratio of dividends per share to PCE similar to those frequently considered in the prior literature. We refer to this second model as our *log-linear model* of the dynamics of dividends. This alternative model leads to some non-linear expressions for the fundamental price p_t^* and the residual term ϕ_t that we will approximate with first order Taylor expansions. With these approximations, we can conduct forecasting exercises similar to those in [Campbell and Shiller \(1988\)](#).

We conduct forecasting regressions with this approximate model and we find results consistent with what we found with our first model. That is, we find that with a high value of β , corresponding to a low value of $\bar{\phi}$, our forecasting regressions favor the Dividends Hypothesis over the Excess Volatility Hypothesis.

Our second model of the dynamics of dividends is based on the following assumptions.

Assumption 1 Logs: Assume that $\log(d_t)$ is an ARIMA process that is integrated of order one. Let $\log(x_t)$ be the trend of the Beveridge-Nelson decomposition of this time series. Then

$$\log(x_t) \equiv \lim_{k \rightarrow \infty} \mathbb{E}_t \log(d_{t+k})$$

with

$$\log(x_{t+1}) = \log(x_t) + \epsilon_{x,t+1}. \tag{21}$$

Assumption 2 Logs: Assume that the transitory component of $\log(d_t)$ follows an AR(1) process of the form

$$(\log(d_{t+1}) - \log(x_{t+1})) = \rho (\log(d_t) - \log(x_t)) + \epsilon_{d,t+1}. \tag{22}$$

One interpretation of this model of the dynamics of log dividends is that it is simply the analog in logarithms of the first model of the dynamics of dividends that we presented above.

A second interpretation of this model is that the log of dividends relative to PCE is

subject to persistent shocks to its growth rate as in the long run risks model. In Appendix E.1, following [Morley \(2002\)](#), we show how to compute the Beveridge-Nelson trend $\log(x_t)$ and the deviation from trend $\log(d_t) - \log(x_t)$ for a model in which the dynamics of log dividends are given by

$$\log(d_{t+1}) - \log(d_t) = z_{t+1} + \epsilon_{d,t+1} \quad (23)$$

with

$$z_{t+1} = \rho_z z_t + \epsilon_{z,t+1} \quad (24)$$

Specifically, we show that the trend is given by

$$\log(x_t) \equiv \lim_{K \rightarrow \infty} \mathbb{E}_t \log(d_{t+K}) = \log(d_t) + \frac{\rho_z}{1 - \rho_z} z_t$$

with transitory component

$$(\log(d_t) - \log(x_t)) = -\frac{\rho_z}{1 - \rho_z} z_t \quad (25)$$

Note that this trend components $\log(x_t)$ is a random walk and the transitory component $(\log(d_t) - \log(x_t))$ is an AR(1).

Motivated by the work of [Gârleanu and Panageas \(2023\)](#), we favor this second interpretation of the dynamics of log dividends captured by the log versions of our Assumptions 1 and 2. This work emphasizes the role of ongoing entry by new firms via initial public offerings in driving down the share of the total stock market represented by the index of price per share as shown in [Figure 2](#). This force of initial public offerings by new firms (or seasoned issuances by incumbent firms) imparts a potentially fluctuating trend negative growth rate to the logarithm of the ratio of dividends per share to consumption. At the same time, changes in the payout policies of incumbent firms to favor share repurchases over dividends alter net share issuance by these incumbent firms. This force imparts a potentially fluctuating trend positive growth rate to the logarithm of dividends per share to consumption if share repurchases are large enough. Changes over time in the strength of these two forces appear as changes in the rate of growth of the logarithm of dividends per share to consumption and anticipated changes in this growth rate are manifest in the logarithm of the aggregate price dividend ratio.

7.1 Price and Fundamental Price in the Log-Linear Model

To compute the fundamental component of price p_t^* implied by this model of the dynamics of log dividends, we must compute the model's implications for the expected value of future

dividends in levels, not logs. It is straightforward to show that the expected value of the level of dividends implied by our second model is given by

$$\mathbb{E}_t d_{t+k} = \exp(\rho^k (\log(d_t) - \log(x_t)) + \log(x_t) + J_k),$$

where J_k is a sequence of terms given by

$$J_k = \frac{1 - (\rho^2)^k}{1 - \rho^2} \frac{\sigma_d^2}{2} + \frac{k}{2} \sigma_x^2 + \frac{1 - \rho^k}{1 - \rho} \rho_{dx} \sigma_d \sigma_x,$$

where σ_x and σ_d are the standard deviations of $\epsilon_{x,t+1}$ and $\epsilon_{d,t+1}$ respectively and ρ_{dx} is their correlation coefficient. The terms J_k are the standard correction for variance when computing the expectation of a lognormal random variable. For details of this calculation see Appendix [E.2](#).

These calculations give us the following solution for the fundamental component of price in this second model:

$$p_t^* = \sum_{k=1}^{\infty} \beta^k \exp(J_k) \exp(\rho^k (\log(d_t) - \log(x_t))) x_t. \quad (26)$$

In logs, we then have

$$\log(p_t^*) = \log \left(\sum_{k=1}^{\infty} \beta^k \exp(J_k) \exp(\rho^k (\log(d_t) - \log(x_t))) \right) + \log(x_t),$$

where the first term on the right side of this equation represents the influence of the stationary component of log dividends on log price. In what follows, we make use of the following first order approximation to this nonlinear term around the point at which $\log(d_t) = \log(x_t)$:

$$\log \left(\sum_{k=1}^{\infty} \beta^k \exp(J_k) \exp(\rho^k (\log(d_t) - \log(x_t))) \right) \approx \bar{\zeta} + \psi(\log(d_t) - \log(x_t)), \quad (27)$$

where

$$\bar{\zeta} \equiv \log \left(\sum_{k=1}^{\infty} \beta^k \exp(J_k) \right)$$

and

$$\psi \equiv \frac{1}{\sum_{k=1}^{\infty} \beta^k \exp(J_k)} \left[\sum_{k=1}^{\infty} \beta^k \exp(J_k) \rho^k \right].$$

We then work with the following first-order approximation to the fundamental component

of price

$$\log(p_t^*) \approx \bar{\zeta} + \psi(\log(d_t) - \log(x_t)) + \log(x_t). \quad (28)$$

We note from equation (1) that, by definition, $p_t^* = p_t - \phi_t$. This observation gives us the following linear approximate model for log price

$$\log(p_t - \bar{\phi}) \approx \bar{\zeta} + \psi(\log(d_t) - \log(x_t)) + \log(x_t) + \log(\xi_t) \quad (29)$$

where

$$\log(\xi_t) \equiv \log(p_t - \bar{\phi}) - \log(p_t - \phi_t).$$

Note that this approximation in equation (29) differs from that commonly used in the prior literature in that it includes the constant term $\bar{\phi} < 0$.

Under the Dividends Hypothesis, we should have $\log(\xi_t) = 0$.

7.2 Forecasting Regressions in the Log-Linear Model

We now run forecasting regressions to test the Dividends Hypothesis in the log-linear model that are directly analogous to those we applied previously to the linear model.

First, we construct a forecasting variable, $\log(p_t^T)$ that is stationary under the Dividends Hypothesis. That statistic is

$$\log(p_t^T) \equiv \log(p_t - \bar{\phi}) - \log(d_t), \quad (30)$$

which from equation (29) is given by

$$\log(p_t^T) \approx \bar{\zeta} + (\psi - 1)(\log(d_t) - \log(x_t)) + \log(\xi_t).$$

Under the Dividends Hypothesis, $\log(\xi_t) = 0$ and thus $\log(p_t^T)$ is stationary. Note that $\log(p_t^T)$ as defined in equation (30) is the standard measure in logs of the dividend-price ratio adjusted by the constant $\bar{\phi}$.

The first test of the Dividends Hypothesis that we consider in this second model is an analog to return regressions of the form in [regression 1](#) that we conducted with our first model. We develop these regressions as follows.

By definition, if we compute the returns to the fundamental component of price as

$$R_{t+1}^* \equiv \frac{p_{t+1}^* + d_{t+1}}{p_t^*},$$

we then have

$$\mathbb{E}_t R_{t+1}^* = \frac{1}{\beta}.$$

Thus, realized returns on the fundamental price should not be predictable.

Under the Dividends Hypothesis, $p_t^* = p_t - \bar{\phi}$. Thus, under the Dividends Hypothesis,

$$R_{t+1}^* = \frac{p_{t+1} - \bar{\phi} + d_{t+1}}{p_t - \bar{\phi}} \quad (31)$$

should not be predictable.

We measure realized log returns on the fundamental price over horizon s by

$$\log(R_{t,t+s}^*) = \sum_{k=0}^{s-1} \log(R_{t+k+1}^*). \quad (32)$$

We then ask whether these log returns are predictable using regressions of the form⁵

$$\log(R_{t,t+s}^*) = \alpha_{R^*,s} + \gamma_{R^*,s} \log(p_t^T) + error_{R^*,t+s}. \quad (\text{regression 4})$$

Under the Dividends Hypothesis, we expect estimates of the slope coefficient $\hat{\gamma}_{R^*,s} = 0$. Note that [regression 4](#) does not require us to take a stand on the parameters β or ψ and thus does not rely on any specific model of the transitory dynamics of log dividends.

The stricter test of the Dividends Hypothesis involves exploring whether we can reject the specified process for dividends, and in particular the assumption that $\log(x_t)$ is a martingale. The analogue of equation (17) in the log-linear model is

$$\mathbb{E}_t \log(p_{t+s} - \bar{\phi}) - \log(p_t - \bar{\phi}) = \psi (\mathbb{E}_t \log(d_{t+s}) - \log(d_t))$$

and the analogue of the expression for expected dividend growth is

$$\mathbb{E}_t \log(d_{t+s}) - \log(d_t) = \frac{(\rho^s - 1)}{\psi - 1} (\log(p_t^T) - \bar{\zeta}).$$

We therefore run the following two regressions. First, we run

$$\log(d_{t+s}) - \log(d_t) = \alpha_{d,s} + \gamma_{d,s} \log(p_t^T) + error_{d,s}, \quad (\text{regression 5})$$

⁵Here we are implicitly assuming that

$$\log(\mathbb{E}_t R_{t+1}^*) - \mathbb{E}_t \log(R_{t+1}^*) = constant$$

Given that assumption, if $\mathbb{E}_t R_{t+1}^*$ is constant over time then $\mathbb{E}_t \log(R_{t+1}^*)$ is also constant.

and compare $\gamma_{d,s}$ to the expected coefficient $\frac{(\rho^s-1)}{\psi-1}$.

Then we run

$$\log(p_{t+s}-\bar{\phi})-\log(p_t-\bar{\phi})-\psi(\log(d_{t+s})-\log(d_t)) = \alpha_{p,s}+\gamma_{p,s} \log(p_t^T)+error_{p,s}, \text{ (regression 6)}$$

and check whether $\hat{\gamma}_{p,s} = 0$. Checking whether $\hat{\gamma}_{p,s} = 0$ is equivalent to checking whether changes in $\log(p_t - \bar{\phi}) - \psi \log(d_t)$ are predictable by $\log(p_t^T)$. Recall, from equation (29), that under the Dividends Hypothesis,

$$\log(p_t - \bar{\phi}) - \psi \log(d_t) \approx \bar{\zeta} + (1 - \psi) \log(x_t).$$

Thus, finding $\hat{\gamma}_{p,s} \neq 0$ would indicate that $\log(x_t)$ is predictable, and would be evidence against the stricter version of the Dividends Hypothesis according to which $\log(x_t)$ constructed from the data with our baseline parameters is a martingale.

7.3 Forecasting Regression Results in the Log-Linear Model

To run these regressions in the log-linear model we use the same values for $\beta = 80/81$ and $\bar{\phi} = -0.9$ that we used in our first linear model. Thus, under the Dividends Hypothesis, our second model has precisely the same implications for the evolution of expected returns as shown for our first model in Figure 5.

To set the parameter ψ , observe that in the limit as the variance of shocks shrinks to zero $\psi \rightarrow (\beta\rho/(1-\beta\rho))/(\beta/(1-\beta))$. We translate ρ and β to a value for ψ using that expression. We estimate a value of $\rho = 0.96936$ using regression 5. We use $\psi = 0.2809$.

We begin with the returns regression 4. Recall that this regression does not require assumptions on β or ρ . Results are shown in Table 5. Consistent with the Dividends Hypothesis, we find estimates of $\hat{\gamma}_{R^*,s} = 0$ and regression R^2 very close to zero. We interpret this evidence as favoring the Dividends Hypothesis.

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$\hat{\gamma}_{R^*,s}$	-0.0023026	-0.012467	-0.00041287	-0.007919
S.E.	(0.018881)	(0.032463)	(0.043138)	(0.051041)
t-Stat	-0.12195	-0.38402	-0.0095709	-0.15515
R^2	0.000162	0.00167	1.1e-06	0.000309

Table 5: Estimates from regressions of the form in regression 4 with $\bar{\phi} = -0.9$.

We present regression results using the log-linear approximation to realized log returns on the fundamental price developed in Campbell and Shiller (1988) in Appendix E.3. Those re-

gressions also show no evidence of log return predictability for log returns on the fundamental price.

Table 6 presents results from estimating the dividend growth regression 5 given $\bar{\phi} = -0.9$. The third row of the table reports the model predicted values for these estimates given our values for β and ρ .

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$\hat{\gamma}_{d,s}$	0.052064	0.20036	0.35897	0.51162
$\frac{1}{1-\psi}(1-\rho^s)$	0.0426	0.2004	0.3719	0.5187
S.E.	(0.035352)	(0.05512)	(0.062982)	(0.067642)
t-Stat	1.4727	3.635	5.6996	7.5637
R^2	0.023	0.131	0.281	0.423

Table 6: Estimates from regression 5 with $\bar{\phi} = -0.9$. In row 3, we use $\rho = 0.96936$ and $\psi = 0.2809$

Finally, Table 7 reports the estimates from the price growth regression 6. Recall that the stricter test of the Dividends Hypothesis is that $\hat{\gamma}_{p,s}$ is supposed to be equal to zero. Our estimated coefficients are not significantly different from zero. Thus, this test fails to reject the stricter version of the Dividends Hypothesis.

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$\hat{\gamma}_{p,s}$	-0.0051709	-0.014316	-0.00035708	-0.010639
S.E.	(0.014785)	(0.027824)	(0.037311)	(0.044851)
t-Stat	-0.34973	-0.51452	-0.0095703	-0.23721
R^2	0.00133	0.003	1.1e-06	0.000721,

Table 7: Estimates from regression 6 with $\beta = 80/81$, $\bar{\phi} = -0.9$ and $\psi = 0.2809$.

To summarize our regression evidence to this point, we find that with parameter values $\beta = 80/81$ and $\bar{\phi} = -0.9$ our regression evidence that favors the Dividends Hypothesis over the Excess Volatility Hypothesis using both the linear and the log-linear models of the dynamics of dividends.

8 Accounting for Observed Stock Prices with our Log-Linear Model of Dividends

Given that our forecasting regression evidence from our log-linear model of dividends is consistent with the Dividends Hypothesis, we now use our parameters $\beta = 80/81$, $\bar{\phi} = -0.9$

and $\psi = 0.2809$ to estimate the sequence of unobserved trend variables $\{x_t\}_{t=0}^T$ that account for the price and dividend data annually from 1929 to the present. We estimate $\log(x_t)$ using equation (29) given the observed sequences for $\{p_t\}$ and $\{d_t\}$ using $\bar{\zeta} = \log(\beta/(1 - \beta))$ and $\log(\xi_t) = 0$.

We show the resulting series for $\log(d_t)$ in blue and $\log(x_t)$ in red in Figure 6. The transitory component of log dividends given by $\log(d_t) - \log(x_t)$ is the difference between these two lines. By comparing this figure to Figure 3, we see that results from our two specifications of the dynamics of dividends are consistent with each other. That is, with these parameter values, our model requires only modest variation over time in the trend component of dividends (here $\log x_t$) to account for the data on log stock prices $\log(p_t)$.

Given the equivalence between our log-linear model of dividends and a long run risks model of the form in equations (23) and (24), we also show the time series for the expected growth rate of the log ratio of dividends per share relative to PCE needed to account for the data on log stock prices. Specifically, we use our estimate of the sequence $\{\log(x_t)\}$ above together with a choice of $\rho = 0.96936$ and equation (25) to estimate the trend growth rates of the logarithm of the ratio dividends per share to PCE given by $\{z_t\}$.

We show results in Figure 7. For comparison purposes, we show the realized values of ten year growth rates of log dividends relative to PCE $(\log(d_{t+10}) - \log(d_t))/10$ in blue and the model-implied values of z_t in red. We argue that the model-implied trend growth rates z_t line up fairly well with the subsequently realized log dividend growth rates in the data. We see this as further confirmation of the Dividends Hypothesis in this log-linear model.

One striking feature of the results in this figure is that the trend growth rate in log dividends per share relative to consumption needed to justify the high valuation of the stock market in the year 2000 shown in red for that year is in line with the realized growth rates of log log dividends per share relative to consumption over the next twenty years (approximately 4% in blue between 2000 and 2010). Thus, in an ex-post sense, the exuberance of the stock market boom of the late 1990's regarding expected dividend growth was not overly optimistic relative to what ended up happening to dividend growth over the next twenty years.

Finally, in Figure 8, we show the implications of our log-linear model for the log price dividend ratio in the data given by $\log(p_t) - \log(d_t)$ in blue and the model-implied log price dividend ratio on the fundamental price given by $\log(p_t^T)$ in red. Note that the difference between these two series is given by

$$\log(p_t) - \log(p_t - \bar{\phi})$$

As we discuss below, it is the difference between the log price-dividend ratio using observed




Figure 6: Blue Line: The logarithm of the ratio of dividends per share for the CRSP Value-Weighted Total Market Index to PCE, $\log d_t$, 1929-2023. Red Line: The expected long-run value of the logarithm of the ratio of dividends per share to PCE, $\log x_t$, that rationalizes the observed price per share of this index, 1929-2023.




Figure 7: Blue Line: The realized annual growth rate of the logarithm of the ratio of dividends per share for the CRSP Value-Weighted Total Market Index to PCE over the 10 years after date t , $(\log d_{t+10} - \log d_t)/10$, 1929-2013. Red Line: The expected trend growth rate of the logarithm of the ratio of dividends per share to PCE, z_t , that rationalizes the observed price per share of this index, 1929-2023.

prices and the model-implied log price-dividend ratio on the fundamental price that accounts for the difference between our return and dividend growth forecasting results and prior results found with $\bar{\phi} = 0$.

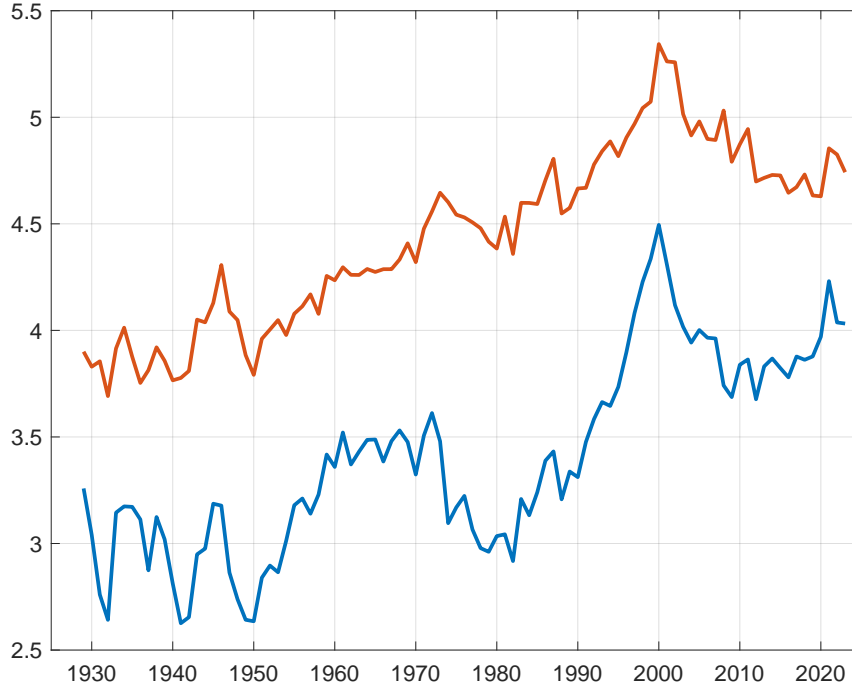


Figure 8: Blue Line: the log price-dividend ratio, $\log(p_t) - \log(d_t)$, 1929-2013. Red Line: the model implied log fundamental price-dividend ratio, $\log(p_t^*) - \log(d_t) = \log(p_t - \bar{\phi}) - \log(d_t)$, 1929-2023.

9 Regression Results With Alternative β and $\bar{\phi}$

To this point, we have examined the implications of the data for the Dividends Hypothesis that fluctuations in observed prices p_t are driven by fluctuations in the fundamental price p_t^* with a high value of $\beta = 80/81$ and a constant residual term $\bar{\phi} < 0$. We have found that, with these parameters, in our forecasting regressions, the data favor Dividends Hypothesis over the Excess Volatility Hypothesis in both our linear model and log-linear model of the dynamics of dividends.

In this section, we repeat our forecasting regressions in both models of dividends with lower values of β in the linear model and alternative values of the constant term $\bar{\phi} = 0$ in the log linear model. With low values of β in the linear model and value of $\bar{\phi}$ close to zero in the log-linear model, we find that the evidence from our quasi-return and log return forecasting regressions favors the Excess Volatility Hypothesis over the Dividends Hypothesis.

Based on these findings we argue that the question of whether stock prices are excessively volatile comes down to the question of which combination of parameters β and $\bar{\phi}$ are appropriate to use with the data to answer this question. We see this as a fruitful area for future research.

We begin with a reexamination of [regression 1](#) in the context of our linear model of the dynamics of dividends, in which the ratio of dividends to PCE is assumed to follow an ARIMA process. Recall that to construct the variables in that regression, p_t^T from equation (14) and $r_{t,t+s}$ from equation (12), we simply need the parameter β and the data on p_t and d_t itself. We do not need to specify the parameter $\bar{\phi}$ to construct these variables and run [regression 1](#).⁶ Thus, given the data, our results from the quasi-return forecasting regressions in the linear model of dividends depend only on the parameter β .

In [Figure 9](#), we show the estimated slope coefficients $\hat{\gamma}_{r,s}$ when we re-run quasi-return forecasting regressions ([regression 1](#)) with alternative values of β . The x-axis plots $\beta/(1-\beta)$ which is in units of a price-dividend ratio. We see clearly in this graph that these regression estimates point to quasi-return predictability when β is low and do not when β is high. This finding is consistent with results regarding the sensitivity of results to β in the linear model reported in [Campbell and Shiller \(1987\)](#)

We now turn to our second model in which we model the dynamics of the log of dividends. We reexamine results from our forecasting [regression 4](#) in that model using alternative values of $\bar{\phi}$ when we compute $\log(p_t^T)$. Note that, given the data, $\bar{\phi}$ is the only parameter that enters into data construction for this regression.

In [Figure 10](#), we show the estimated slope coefficients $\hat{\gamma}_{R^*,s}$ when we re-run return forecasting [regression 4](#) with alternative values of $\bar{\phi}$. The x-axis plots $\bar{\phi} \in [-1, 0]$. We see clearly in this graph that these regression estimates point to predictability of log returns on the fundamental price when $\bar{\phi}$ is closer to zero and no such predictability when $\bar{\phi}$ is close to -1 .

10 Conclusion

In this paper, we use the present value model of [Campbell and Shiller \(1987\)](#) in our equation (1) to frame the question: Do stock prices move too much to be justified by subsequent movements in dividends? This framing of the question leads to our first, linear, model of the dynamics of dividends. We then extend our analysis of this question to follow more closely the model of [Campbell and Shiller \(1988\)](#) with a second, log-linear, model of the dynamics of dividends.

⁶However, to derive our model's implications for mean expected returns and quasi-returns, we do need to specify this parameter.




Figure 9: Estimated slope coefficients $\hat{\gamma}_{r,s}$ from quasi-return forecasting regressions (regression 1) at horizons $s = 1, 5, 10$ and 15 years with alternative values of β . The x-axis plots $\beta/(1 - \beta)$ which is in units of a price-dividend ratio.




Figure 10: Estimated slope coefficients $\hat{\gamma}_{R^*,s}$ from log return forecasting regressions (regression 4) at horizons $s = 1, 5, 10$ and 15 years with alternative values of $\bar{\phi}$.

In both models, we focus on the dynamics of measures of price per share and dividends per share relative to aggregate consumption. We choose to scale the data on price and dividends per share by aggregate consumption based on a conjecture that the ratio of the price to a perpetual claim to aggregate consumption relative to current aggregate consumption is constant over time at a value of $\beta/(1 - \beta)$. This is equivalent to assuming in a Gordon growth model applied to aggregate consumption that the risk adjusted discount rate for valuing aggregate consumption less expected consumption growth has remained constant over time. In our baseline calibration we set this price-dividend ratio for a claim to aggregate consumption at a high value of 80. We reconcile that assumption with high realized returns on equity in excess of consumption growth observed in the data by assuming a negative additive constant term in asset prices, reflecting compensation for risk. This constant term depresses the observed stock price and thus raises observed rates of return on equity implied by the model.

Investors in our model receive news about dividends in the long run as captured by the Beveridge-Nelson trend in the ratio of dividends per share to aggregate consumption. When investors evaluate the price impact of this news using a high value for the discount factor β , small movements in this long run expected value generate high volatility in equilibrium stock prices.

We use a suite of forecasting regressions suggested by [Campbell and Shiller \(1987\)](#) and [Campbell and Shiller \(1988\)](#) to test the hypothesis that observed movements in stock prices relative to consumption are driven by changes in the expected present value of dividends relative to consumption against the alternative hypothesis that they reflect movements in future expected returns. We find that with the parameter assumptions described above, the Dividends Hypothesis passes these regression tests.

We then used these models to offer accounts, year-by-year from 1929 to 2023, of the specific movements in the Beveridge-Nelson trend for the ratio of dividends to aggregate consumption or the log of this ratio needed to reconcile the data on current dividends per share relative to consumption and the current ratio of price per share relative to consumption. We argue that these fluctuations look plausible. In the case of our second, log-linear model of the dynamics of dividends, we also show an interpretation of the model in terms of a simple long run risks model. One striking result from this exercise is that the expected trend growth rate for the log ratio of dividends per share relative to consumption needed to justify the high stock market valuation of year 2000 are in line with the realized growth in the log ratio of dividends per share relative to consumption observed over the next twenty years. Thus, the exuberant stock market valuations of that time period do not seem unjustified in hindsight.

Finally, we confirm that our parameter assumptions of a high value of β and a low additive

risk-adjustment term $\bar{\phi}$ are critical to our findings by repeating our analysis with alternative, lower values of β and a value of the constant $\bar{\phi} = 0$. When we do so, we find that the data favor the Excess Volatility Hypothesis over the Dividends Hypothesis.

In sum, we see our findings as indicating that the answer one gives to the question of whether the stock market moves too much to be justified by subsequent movements in dividends really depends on parameters. What guidance do we have from theory regarding appropriate parameter choices?

Much of the existing literature has focused on a log-linear model of the dynamics of dividends together with a pricing kernel that is also conditionally lognormal. In such models, risk typically impacts asset prices in a multiplicative rather than an additive way.

In contrast, in the linear model of the dynamics of dividends, dividends per share can go negative. Combining that dividend process with a pricing kernel that is conditionally lognormal delivers an additive and constant risk adjustment term in equilibrium valuations.

Which model should be preferred? When applied to data on dividends per share, the log-linear model has the advantage of being consistent with the fact that dividends per share, by definition, cannot go negative. But if one were to follow [Miller and Modigliani \(1961\)](#) and focus on valuing total cash flows to equity rather than dividends per share, and total market capitalization rather than price per share, one has to confront the fact total cash flows to equity in the data frequently do go negative. So our first linear model seems more appropriate for studying fluctuations in total market capitalization of the stock market and total flows to owners of equity from the perspective of an “equilibrium” investor who holds the entire market at every moment in time.

Ultimately, however, both models need to be consistent with the same paths of realized and expected rates of return, as we discuss in [Appendix B](#). Future research should seek to reconcile these two valuation perspectives in terms of their implications for whether or not there is a substantial constant additive risk adjustment $\bar{\phi}$ in stock market valuation. Resolving this question will further improve our understanding of the drivers of stock market volatility.

Appendices

A The Argument for Excess Volatility in Shiller (1981)

In this appendix, we review the criticisms of the interpretation of Figure 1 in Shiller (2014), which updates Shiller (1981), as evidence for excess volatility. Kleidon (1986) and Marsh and Merton (1986) criticized this interpretation on the grounds that both the levels of dividends per share and price per share are non-stationary. We focus on the criticisms of Kleidon (1986).

Let the realized data on price per share be given by $\{P_t\}$. Let the realized data on dividends per share be given by $\{D_t\}$.

Consider the following simple valuation model. In this valuation model, assume that the logarithm of dividends per share, denoted by d_t , evolves according to

$$d_{t+1} = \tilde{g} + d_t + \sigma\epsilon_{t+1},$$

where $\epsilon_{t+1} \sim N(0, 1)$ and \tilde{g} is a constant. With this assumption, we have

$$\mathbb{E}_t D_{t+1} = (1 + g)D_t$$

and, more generally,

$$\mathbb{E}_t D_{t+k} = (1 + g)^k D_t,$$

where

$$g = \exp(\tilde{g} + \frac{1}{2}\sigma^2) - 1.$$

With this model of expected dividends, create the model's implications for price per share based under constant discounting as

$$P_t = \sum_{k=1}^{\infty} (1 + r)^{-k} \mathbb{E}_t D_{t+k} = \frac{1 + g}{r - g} D_t. \quad (33)$$

The prediction for the price constructed in Shiller (1981) and Shiller (2014) under the assumption that we have an infinite realized sequence of dividends is

$$P_t^* = \sum_{k=1}^{\infty} (1 + r)^{-k} D_{t+k}.$$

That is, we use realized dividends without the expectation.

In this case, both P_t and P_t^* are non-stationary. But, theoretically, since the model's implication for the price P_t is directly proportional to the currently realized dividend D_t , we have that the standard deviation of log changes in price is given as

$$\text{Std}(\log(P_{t+1}) - \log(P_t)) = \sigma.$$

In contrast, it is straightforward to verify via a Monte Carlo simulation that the standard deviation of log changes in the predicted price constructed using the method above is

$$\text{Std}(\log(P_{t+1}^*) - \log(P_t^*))$$

which is typically at least an order of magnitude smaller than σ . [Kleidon \(1986\)](#) shows several results from such Monte Carlo simulations that lead to figures with these simulated data very similar in appearance to those in [Shiller \(1981\)](#).

The issue of why this approach to assessing stock market volatility goes wrong can be seen clearly from equation (33). If dividends are a random walk, then news that arrives between t and $t + 1$ in the form of the shock ϵ_{t+1} moves agents' expectations of future dividends out into the infinite future, since

$$\mathbb{E}_t D_{t+k} = D_t$$

and

$$\mathbb{E}_{t+1} D_{t+k} = D_{t+1}.$$

In contrast, if we follow the procedure in [Shiller \(1981\)](#) to construct P_t^* , then we are effectively assuming that agents' expectations of future dividends never move at all. That is,

$$\mathbb{E}_t^* D_{t+k} = D_{t+k}$$

and

$$\mathbb{E}_{t+1}^* D_{t+k} = D_{t+k}.$$

The only updating to P_t^* that occurs is that the first dividend is dropped and the discounting of future dividends is updated by $(1 + r)$. That is, P_t^* satisfies

$$P_t^* = \frac{1}{1 + r} [D_{t+1} + P_{t+1}^*].$$

This equation implies that

$$\log(P_{t+1}^*) - \log(P_t^*) = \log(1 + r) - \log\left(1 + \frac{D_{t+1}}{P_{t+1}^*}\right) \approx r - \frac{D_{t+1}}{P_{t+1}^*}.$$

Given that $\{\log(D_t)\}$ is assumed to be a random walk with a constant drift and our assumption of a constant discount rate r , one should not expect $\frac{D_{t+1}}{P_{t+1}^*}$ to be variable. It varies only because of random runs of positive or negative values of ϵ_t leading to positive or negative runs of realized dividend growth above or below the mean. A Monte Carlo simulation reveals this variance to be very small.

B How a Value-Weighted Stock Index is Constructed

We use data on the CRSP Value-Weighted Total Market Index 1929-2023. We are somewhat pedantic in our presentation of this data as some readers may not be familiar with its construction, and several elements of its construction are important in understanding the motivation for the key assumptions in our valuation model.

The original data we use are CRSP indices of annual returns without dividends (denoted by R_{t+1}^{nd}), returns with dividends (denoted by R_{t+1}^d), and total market capitalization (denoted by TMC_t) on the CRSP Value-Weighted Index combining stocks listed in the NYSE, AMEX, and NASDAQ exchanges for the years 1929-2023. We focus on this time period as this is the time period for which we also have NIPA data on Personal Consumption Expenditures.

As is well known, CRSP annual value-weighted returns on the total stock market are high on average and quite volatile. In our sample, the arithmetic averages of nominal and real returns with dividends (R_{t+1}^d) are 11.6% and 8.6% respectively (deflating with the PCE deflator), and these nominal and real returns have a standard deviations of 19.8% and 19.5% respectively.

The measure of price per share for the CRSP Value-Weighted Index that we use as the measure of the value of the stock market in our study is constructed from the cumulation of annual returns without dividends R_{t+1}^{nd} . Specifically, if we let P_{Dt} denote the level of price per share on the last day of year t , we construct $P_{D,t+1} = R_{t+1}^{nd} P_{Dt}$. Note that P_{Dt} is an index number in that the initial value must be normalized.

We plot the ratio of this index of price per share to Personal Consumption Expenditures (PCE) in the left panel of Figure 1. We have normalized the index of price per share so that the initial value of this ratio is equal to one. As is clearly evident in this figure, this ratio is quite volatile.

The measure of dividends per share for the CRSP Value-Weighted Index that we denote by D_t and use as our measure of cash flows to someone holding the CRSP Value-Weighted Index is constructed to solve the following equation

$$\frac{D_{t+1} + P_{D,t+1}}{P_{Dt}} = R_{t+1}^d$$

That is, given the index for price per share, the index for annual dividends per share is chosen so that returns match value-weighted returns with dividends. This equation pins down the ratio of dividends per share to price per share. The scale of dividends per share is set by the normalization of the level of price per share.

We plot the ratio of this index of dividends per share to PCE in the right panel of Figure 1.

The concepts of price per share and dividends per share for a broad stock market index are constructed to meet specific needs that are not the same as those of an academic researcher seeking to understand fluctuations in the value of the stock market. In particular, the measure of price per share represents the dynamics of the value of and payouts to the portfolio of an investor who follows a very specific trading strategy that does not correspond to equilibrium notions of “holding the market” as in [Sharpe \(1964\)](#) and [Lucas \(1978\)](#). An investor who invested to track the CRSP Value-Weighted Total Market Index, would end up holding a constantly changing share of the total market capitalization of that index, with the changes in that share of the market held engineered specifically to reduce the volatility of the cash flows to that investor, leaving that investor only with payouts from dividends. We argue, then, that it is no surprise that empirical work using these data would arrive at the conclusion that stock prices move too much to be justified by subsequent changes in payouts. This finding is hard-wired into the construction of the data.⁷

An alternative approach to assessing whether the volatility of the stock market is too high relative to the volatility of the cash flows going to someone invested in the market is to examine the cash flows to an investor who followed an “equilibrium” strategy of holding a constant fraction of the total market capitalization of the stocks in a broad stock index at every moment in time. This is the portfolio strategy that we take as the equilibrium strategy of “holding the market.” As we describe next, it is a simple exercise to construct these cash flows using data on the index returns including dividends, index returns excluding dividends, the level of the index in question, and the total market capitalization of the stocks in the index. This methodology is presented in [Dichev \(2006\)](#), who notes that it is commonly used to in the mutual fund industry to reconcile fund returns, fund flows, and fund market values.

When we do so, using the CRSP Value-Weighted Total Market Index as an illustration, we find that the cash flows associated with this “equilibrium” investment strategy are massively volatile, calling into question the conclusion that stock prices move too much to be justified by subsequent movements in dividends. It is straightforward to illustrate the same findings

⁷This concern is heightened by the recognition that in the decades following World War II, firms smoothed their dividend payouts. See [Marsh and Merton \(1986\)](#), [Chen, Da, and Priestly \(2012\)](#) and the papers cited therein for a discussion of the impact of dividend smoothing on variance bounds tests and predictive regressions.

with other broad value-weighted stock indices.

To begin, we review the basics of the construction of a broad value-weighted stock market index.

At any point in time, t , a value-weighted stock index $X(t)$ is given as a time-varying fraction of the total market capitalization of the stocks in the index. That is, if we let $\Omega(t)$ be the set of stocks in the index, and $p_i(t)$ and $s_i(t)$ be the prices and shares outstanding for those stocks, then the total market capitalization of the stocks in the index, denoted here by $TMC(t)$, is given by

$$TMC(t) = \sum_{i \in \Omega(t)} p_i(t) s_i(t). \quad (34)$$

The level of the index at t , which we denote by $X(t)$, is given by

$$X(t) = \frac{1}{\theta(t)} TMC(t), \quad (35)$$

where $\theta(t)$ is called the “divisor” for the index at t . The t in $\theta(t)$ is there to denote that this divisor changes over time. Note here that $1/\theta(t)$ represents the fraction of the total market capitalization of the stocks in the index held at t by an investor tracking the level of index rather than the total market capitalization of the stocks in the index.

The gross value-weighted return on this index between periods t and $t + 1$ not including dividends is given by

$$R_{t,t+1}^{no\ dividends} = \sum_{i \in \Omega(t)} \left(\frac{p_i(t) s_i(t)}{\sum_{j \in \Omega(t)} p_j(t) s_j(t)} \right) \frac{p_i(t+1)}{p_i(t)}. \quad (36)$$

If we denote by $d_i(t+1)$ the dividend paid by firm i at time $t+1$ to someone who owned the share at time t , then aggregate dividends paid in $t+1$ are given by

$$D(t+1) = \sum_{i \in \Omega(t)} d_i(t+1) s_i(t), \quad (37)$$

and dividends per share are given by

$$DPS(t+1) = \frac{1}{\theta(t)} D(t+1). \quad (38)$$

The gross value-weighted return on this index between periods t and $t+1$ including

dividends is given by

$$R_{t,t+1}^{w \text{ dividends}} = \sum_{i \in \Omega(t)} \left(\frac{p_i(t) s_i(t)}{\sum_{j \in \Omega(t)} p_j(t) s_j(t)} \right) \left(\frac{p_i(t+1) + d_i(t+1)}{p_i(t)} \right). \quad (39)$$

The divisor at $t + 1$, denoted by $\theta(t + 1)$ is chosen so that the change in the index level from t to $t + 1$ corresponds to the gross value-weighted return without dividends; that is,

$$\frac{X(t+1)}{X(t)} = R_{t,t+1}^{no \text{ dividends}}. \quad (40)$$

From equation (36), this implies that

$$X(t+1) = \frac{1}{\theta(t)} \sum_{i \in \Omega(t)} p_i(t+1) s_i(t).$$

With this construction, it is also the case that the gross value-weighted return including dividends corresponds in the natural manner to the returns defined in terms of price per share and dividends per share. That is,

$$\frac{X(t+1) + DPS(t+1)}{X(t)} = R_{t,t+1}^{w \text{ dividends}}. \quad (41)$$

What we have in equations (40) and (41) is that data on the *price per share* and *dividends per share* for the index can be used to reproduce the value-weighted returns on the stocks in the index without and with dividends between periods t and $t + 1$ in a natural manner consistent with what the notation would be if the entire index were a single firm.

But how is this construction achieved? In reality, the stocks in the index are not a single firm since some stocks are added and some are removed and since the incumbent firms in the index often take actions to change their number of outstanding shares. To deal with these issues, the divisor of the index is adjusted so that equation (35) is also satisfied in period $t + 1$. This approach to index construction implies that the divisor changes from period t to period $t + 1$ according to

$$\theta(t+1) = \frac{\sum_{i \in \Omega(t+1)} p_i(t+1) s_i(t+1)}{\sum_{i \in \Omega(t)} p_i(t+1) s_i(t)} \theta(t) \quad (42)$$

Here, we see that the construction of the index implies a certain trading strategy that does not correspond to holding a constant share of the market capitalization of the stocks in the index. An investor who aims to hold a portfolio that tracks this index would be required

to adjust the fraction of the total market capitalization of the stocks in the index that he or she held from $1/\theta(t)$ to $1/\theta(t+1)$, as indicated in equation (42). That is, if, at $t+1$, the shares outstanding for the firms in the index at $t+1$ have increased when evaluated at $t+1$ prices, because of either incumbent firms issuing more shares on net (raising capital), or firms being added to the index at $t+1$ being more valuable than firms leaving the index between t and $t+1$, the divisor rises, and the implied share of the total market capitalization of the stocks in the index held by an investor tracking the index falls. Likewise, if incumbent firms buy back their shares (returning capital), or if firms being added to the index at $t+1$ are less valuable than firms leaving the index between t and $t+1$, the divisor falls and the implied share of the total market capitalization of the stocks in the index held by an investor tracking the index rises.

More generally, there is a long list of circumstances that lead to changes in the number of shares outstanding for the firms in the index between t and $t+1$ that are referred to as *corporate actions*. These include initial public offerings, delistings, mergers and acquisitions, reverse mergers/takeovers, tendered shares, spin-offs, rights offerings, and certain transactions connected with warrants, options, partly paid shares, convertible bonds, contingent value rights, and so on. The staff at CRSP (and S&P Dow Jones Indices for their indices) invest considerable resources tracking all of these events and adjusting the index divisor accordingly.

What this index construction methodology implies is that an investor who aims to hold a portfolio that tracks the level of the index over time will not participate in any of these corporate actions. As a result, this investor receives only the cash flows associated with dividends paid at $t+1$ by incumbent firms in period t . This investor will receive neither the cash flows associated with new share issuance or share buybacks by these incumbent firms nor the cash flows associated with the entry and exit of firms from the index (or any of the other possible corporate actions). Instead of participating in these cash flows, an investor who aims to track the level of the index simply adjusts the fraction held of the total market capitalization of the stocks in the index instead of contributing or removing cash as indicated by these corporate actions.

How then can we use the data from the index to recover the cash flows received by an investor following the equilibrium trading strategy of “holding the market” at all times. To do this, we invoke the theorem of [Miller and Modigliani \(1961\)](#) that asserts that changes in a firm’s dividend policy to return cash to shareholders in the form of net buybacks change neither the returns nor the market capitalization of the firm. Using this principle and following [Dichev \(2006\)](#), we construct the additional cash flows accruing to an investor following the equilibrium trading strategy of “holding the market” at all times, denoted here by $CAFC_{t+1}$

for *corporate action cash flows* using the equation

$$CACF(t+1) = R_{t,t+1}^{no\ dividends} TMC(t) - TMC(t+1). \quad (43)$$

We then have the total cash flows to an equilibrium investor holding the market at $t+1$ are $D(t+1) + CACF(t+1)$.

Equation (43) is an accounting identity that follows from a reconciliation of returns on the market from t to $t+1$ and the change in market capitalization of the market as a whole. This accounting identity implies that these cash flows from corporate actions can be stated equivalently as

$$CACF(t+1) = \sum_{i \in \Omega(t)} p_i(t+1)s_i(t) - \sum_{i \in \Omega(t+1)} p_i(t+1)s_i(t+1).$$

That is, these are the cash flows that arise from all changes in the number of shares outstanding from time t to time $t+1$ when valued at prices at time $t+1$.

Now, what impact do these calculations have on the ratio of dividends per share to price per share as measured by this index? We have by definition that the ratio of dividends per share to price per share is equal to the ratio of total dividends to total market capitalization of the stocks in the index; that is,

$$\frac{DPS(t)}{X(t)} = \frac{D(t)}{TMC(t)}.$$

This then implies that

$$\frac{DPS(t)}{X(t)} = \frac{D(t) + CACF(t)}{TMC(t)} - \frac{CACF(t)}{TMC(t)}. \quad (44)$$

Consider the implications of [Miller and Modigliani \(1961\)](#) for the terms in this equation. In their analysis, they take the total cash flows to equity investors $D(t) + CACF(t)$ as given. With this assumption, they show that total market capitalization $TMC(t)$ is independent of the payout policy as determined by the split of total payouts into dividends $D(t)$ and cash flows arising from corporate actions $CACF(t)$. Thus, the first ratio on the right side of equation (44), given by $\frac{D(t)+CACF(t)}{TMC(t)}$ is fundamental. It is not impacted by changes in corporate actions. Of course, the other two ratios, $\frac{DPS(t)}{X(t)}$ and $\frac{CACF(t)}{TMC(t)}$ are impacted by corporate actions. To the extent that the ratio $\frac{CACF(t)}{TMC(t)}$ is volatile, the relative volatility of the fundamental ratio $\frac{D(t)+CACF(t)}{TMC(t)}$ and the ratio of dividends per share and price per share $\frac{DPS(t)}{X(t)}$ will be different. To the extent that there is are low frequency movements in the

ratio $\frac{CACF(t)}{TMC(t)}$ not present in the fundamental valuation ratio $\frac{D(t)+CACF(t)}{TMC(t)}$, there will be low frequency movements in the ratio of dividends per share to price per share $\frac{DPS(t)}{X(t)}$ not driven by fundamentals but instead driven by corporate actions.

In this appendix, we compare our measure of total payouts to equilibrium investors in equity as represented by the CRSP Value-Weighted Total Market Index to that constructed in Davydiuk et al. (2023) which builds on the work of Boudoukh et al. (2007) but also uses the CRSP Stock file as we do. In Figure B.1, we show in blue the ratio of total payouts on the CRSP Value-Weighted Total Market Index to total market capitalization of the stocks in that index $((D(t) + CACF(t))/TMC(t))$ from 1926-2023. In red, we show the ratio of equity cash payouts less net equity issuance to total market capitalization as measured in Davydiuk et al. (2023) for the time period 1975-2017 obtained from the *Journal of Finance* website for this article. Note that the measure constructed in Davydiuk et al. (2023) accounts for share buybacks and also for changes in entity structure due to initial public offerings (IPOs), mergers, acquisitions, and exchanges.

As is evident in this figure, these two measures are quite similar.




Figure B.1: In blue: the ratio of payouts to an equilibrium investor to total market capitalization of the stocks in the CRSP Value-Weighted Total Market Index $((D(t) + CACF(t))/TMC(t))$, where payouts are summed over the calendar year. In red: the ratio of total payouts to equity to total market capitalization of equity from Davydiuk et al. (2023).

C An Explicit Pricing Kernel

We now present a micro-foundation for the assumption that $\bar{\phi} \neq 0$. It is based on an explicit model of a pricing kernel in which innovations to the logarithm of the growth of marginal utility are normal, as are innovations to the logarithm of consumption growth, while innovations to the ratio of dividends per share to consumption d_t are normal. We show in particular that this model yields a valuation model for equity consistent with equation (10).

We begin with standard assumptions regarding the dynamics of consumption growth and of the pricing kernel used to value assets. Let the log of consumption growth between t and $t + 1$ be given by

$$g_{C,t+1} = \bar{g}_C + \sigma_{g_C} \epsilon_{C,t+1},$$

where \bar{g}_C measures trend growth, and shocks to the log growth rate $\sigma_{g_C} \epsilon_{C,t+1}$ are drawn from a Normal distribution with mean zero variance $\sigma_{g_C}^2$.

Let the log of the pricing kernel be given by

$$m_{t+1} = \bar{m} + \lambda_C \epsilon_{C,t+1} + \lambda_D \epsilon_{D,t+1} + \lambda_X \epsilon_{X,t+1},$$

where the parameters λ_C , λ_D and λ_X capture, respectively, the pricing kernel loadings on the three shocks in the model: innovations to consumption growth $\epsilon_{C,t+1}$, and the transitory and permanent innovations to the ratio of dividends per share to consumption, $\epsilon_{D,t+1}$ and $\epsilon_{X,t+1}$, that we introduced in the previous section.

These assumptions for consumption growth and the pricing kernel jointly imply that the following three variables are all constant over time: (i) the price of a claim to consumption one period ahead relative to current consumption, (ii) the expected growth of consumption, and (iii) the riskless interest rate.

In particular, the price of a claim to consumption one period ahead relative to consumption today is given by

$$\frac{P_{Ct}^{(1)}}{C_t} = \beta = \mathbb{E}_t [\exp(m_{t+1} + g_{C,t+1})] = \exp \left(\bar{m} + \bar{g}_C + \frac{1}{2} ((\lambda_C + \sigma_{g_C})^2 + \lambda_D^2 + \lambda_X^2) \right). \quad (45)$$

In equation 45, β does not depend on time so we omit the time subscript. Because this price is constant over time, Assumption 1 of our simple valuation model is satisfied.

The gross one-period risk-free interest rate implied by this pricing kernel is also constant

and given by

$$R^{RF} = \frac{1}{\mathbb{E}_t[\exp(m_{t+1})]} = \exp\left(-\bar{m} - \frac{1}{2}(\lambda_C^2 + \lambda_D^2 + \lambda_X^2)\right).$$

The expected growth rate of the level of consumption is

$$\mathbb{E}_t[\exp(g_{C,t+1})] = \exp\left(\bar{g}_C + \frac{1}{2}\sigma_{g_C}^2\right).$$

Observe that the expected return on a one-period consumption bond is

$$R^C = \frac{\mathbb{E}_t[\exp(g_{C,t+1})] C_t}{P_{Ct}^{(1)}} = \frac{\exp\left(\bar{g}_C + \frac{1}{2}\sigma_{g_C}^2\right)}{\beta}$$

Thus the expected return on a consumption bond in excess of the risk free rate is

$$R^C - R^{RF} = \exp(-\lambda_C \sigma_{g_C})$$

Thus,

$$\beta = \frac{\mathbb{E}_t[\exp(g_{C,t+1})]}{R^C} = \frac{\mathbb{E}_t[\exp(g_{C,t+1})]}{R^{RF} + \exp(-\lambda_C \sigma_{g_C})}$$

As we have noted above, in the data, the risk free interest rate appears to be below the expected growth rate of consumption. For us to have a finite value for the coefficient $\gamma_X = \beta/(1 - \beta)$, as is standard, we need to have a sufficiently large risk premium on a claim to consumption as determined by $\exp(-\lambda_C \sigma_{g_C})$.

We do not want to argue that these moments are all constant over time in the data. But the fact that they are constant in our model allows us to transparently make the point that it is possible to account for the observed volatility of stock prices based entirely on volatility of expected cash flows. We leave to future work the project of extending our valuation framework to richer models for consumption growth or for the pricing kernel under which these data moments vary over time.

C.1 Pricing Dividends

We now turn to pricing claims to dividends. We assume that the dynamics of the ratio of dividends per share to consumption are given by equations (6) and (7). Note that this model departs from standard asset pricing formulations in that innovations to the ratio of dividends to consumption are normal rather than log-normal. We now show how to compute prices of claims to dividends and equity given the dynamics of the pricing kernel and consumption

growth using Stein's Lemma.

The prices of dividends relative to consumption satisfy the recursive formula that

$$\frac{P_{D,t}^{(k)}}{C_t} = \mathbb{E}_t \left[\exp(m_{t+1} + g_{C,t+1}) \frac{P_{D,t+1}^{(k-1)}}{C_{t+1}} \right] \quad (46)$$

We guess and verify that the price of a claim to dividends k periods ahead has the following form:

$$\frac{P_{D,t}^{(k)}}{C_t} = A_k \left(\frac{D_t}{C_t} - X_t \right) + B_k X_t + H_k \quad (47)$$

We solve for the coefficients A_k , B_k , and H_k recursively using equation 46 and the method of undetermined coefficients as described in Appendix C.2. We show that the coefficients A_k, B_k satisfy the recursion

$$\begin{aligned} A_k &= \beta \rho A_{k-1} = (\beta \rho)^k \\ B_k &= \beta B_{k-1} = \beta^k \end{aligned}$$

and the coefficients H_k satisfy

$$H_k = \beta (H_{k-1} + \lambda_D A_{k-1} \sigma_D + \lambda_X (B_{k-1} - A_{k-1}) \sigma_X).$$

with $H_0 = 0$.

Note that these coefficients are independent of date t . We can then construct the value of a claim to equity as in equation (10) from

$$\frac{P_{D,t}}{C_t} = \sum_{k=1}^{\infty} \frac{P_{D,t}^{(k)}}{C_t} = \gamma^D \left(\frac{D_t}{C_t} - X_t \right) + \gamma^X X_t + \phi,$$

where $\gamma^D = \frac{\beta \rho}{1 - \beta \rho}$, $\gamma^X = \frac{\beta}{1 - \beta}$, and $\phi = \sum_{k=1}^{\infty} H_k$. Note that H_k and thus ϕ are constant over time, so this affine model satisfies Assumption 2 of our valuation framework.

C.2 Solving for A_k , B_k , and H_k

The price for a claim to dividends in the current period is given by

$$\frac{P_{D,t}^{(0)}}{C_t} = \frac{D_t}{C_t}$$

so $H_0 = 0$, $A_0 = B_0 = 1$.

We have the following recursion for all other horizons k :

$$\frac{P_{Dt}^{(k)}}{C_t} = \mathbb{E}_t \left[\exp(m_{t+1} + g_{C,t+1}) \frac{P_{D,t+1}^{(k-1)}}{C_{t+1}} \right]. \quad (48)$$

We use this to solve for A_k , B_k and H_k as follows.

This recursive equation implies that⁸

$$A_k \left(\frac{D_t}{C_t} - X_t \right) + B_k X_t + H_k = \mathbb{E}_t [\exp(m_{t+1} + g_{C,t+1})] \left[\rho A_{k-1} \left(\frac{D_t}{C_t} - X_t \right) + B_{k-1} X_t + H_{k-1} \right] + \exp(\bar{m} + \bar{g}_C) \mathbb{E}_t [\exp((\lambda_C + \sigma_{g_C})\epsilon_{C,t+1} + \lambda_D \epsilon_{D,t+1} + \lambda_X \epsilon_{X,t+1}) [A_{k-1} \sigma_D \epsilon_{D,t+1} + (B_{k-1} - A_{k-1}) \sigma_X \epsilon_{X,t+1}]]$$

Matching coefficients on $(\frac{D_t}{C_t} - X_t)$ and X_t gives us that the coefficients A_k , B_k satisfy the recursion

$$A_k = \beta \rho A_{k-1} = (\beta \rho)^k$$

$$B_k = \beta B_{k-1} = \beta^k$$

To solve for the recursion for the constant term H_k , we need to solve for the term

$$\exp(\bar{m} + \bar{g}_C) \mathbb{E}_t [\exp((\lambda_C + \sigma_{g_C})\epsilon_{C,t+1} + \lambda_D \epsilon_{D,t+1} + \lambda_X \epsilon_{X,t+1}) [A_{k-1} \sigma_D \epsilon_{D,t+1} + (B_{k-1} - A_{k-1}) \sigma_X \epsilon_{X,t+1}]]$$

To do so, we use the result that if x and y and z are independent standard normal random variables and a, b, c, d are scalar constants, then

$$\mathbb{E} [\exp(ax + by)(cx + dz)] = ca \exp((a^2 + b^2)/2) \quad (49)$$

This formula is an application of Stein's Lemma. We prove it in Appendix C.2.

This gives us that

$$\begin{aligned} \exp(\bar{m} + \bar{g}_C) \mathbb{E}_t [\exp((\lambda_C + \sigma_{g_C})\epsilon_{C,t+1} + \lambda_D \epsilon_{D,t+1} + \lambda_X \epsilon_{X,t+1}) [A_{k-1} \sigma_D \epsilon_{D,t+1} + (B_{k-1} - A_{k-1}) \sigma_X \epsilon_{X,t+1}]] = \\ \exp(\bar{m} + \bar{g}_C) \exp \left(\frac{1}{2} (\lambda_C + \sigma_{g_C})^2 + \frac{1}{2} (\lambda_D^2 + \lambda_X^2) \right) (\lambda_D A_{k-1} \sigma_D + \lambda_X (B_{k-1} - A_{k-1}) \sigma_X) = \\ \beta (\lambda_D A_{k-1} \sigma_D + \lambda_X (B_{k-1} - A_{k-1}) \sigma_X). \end{aligned}$$

⁸This equation is derived by substituting eq. (47) into the left-hand side of eq. (48) and again into the right-hand side (this time evaluated at $t+1$ and $k-1$). Next we used eqs. (7) and (6) to express $(\frac{D_{t+1}}{C_{t+1}} - X_{t+1})$ and X_{t+1} in terms of $(\frac{D_t}{C_t} - X_t)$ and X_t plus innovation terms.

This result implies that we can solve for the coefficients H_k recursively from

$$H_k = \beta (H_{k-1} + \lambda_D A_{k-1} \sigma_D + \lambda_X (B_{k-1} - A_{k-1}) \sigma_X).$$

C.3 Proof of Formula (49)

One can prove this formula by using the moment generating function for normal random variables. In particular, we start by computing for a normal random variable

$$\mathbb{E} \exp(atx) = \exp(at\mu + \frac{1}{2}a^2t^2\sigma^2).$$

We then have

$$\mathbb{E} ax \exp(atx) = \mathbb{E} \frac{d}{dt} \exp(atx) = \exp(at\mu + \frac{1}{2}a^2t^2\sigma^2)(a\mu + ta^2\sigma^2).$$

If we evaluate this expression at $t = 1$ with $\mu = 0$ and $\sigma = 1$ for a standard normal distribution, we have

$$\mathbb{E} ax \exp(ax) = \exp(\frac{1}{2}a^2)a^2.$$

We multiply by c/a to obtain

$$\mathbb{E} cx \exp(ax) = \exp(\frac{1}{2}a^2)ca.$$

We then have

$$\mathbb{E} \exp(ax + by)(cx + dz) = \mathbb{E} \exp(by) \mathbb{E} cx \exp(ax) + \mathbb{E} \exp(by) \mathbb{E} \exp(ax) \mathbb{E} dz$$

by the independence of x, y and z . Finally, since $\mathbb{E}z = 0$ and $\mathbb{E} \exp(by) = \exp(\frac{1}{2}b^2)$, we get equation (49).

D Spurious Regressions

Define

$$p_t^L = p_t - \frac{\beta\rho}{1 - \beta\rho} d_t.$$

From the equilibrium pricing equation (10),

$$p_t^L = \phi_t + \Gamma x_t.$$

Thus, this price-dividend spread isolates the combined price impact of the trend component of dividends x_t and the residual component ϕ_t . Under the Dividends Hypothesis, $\phi_t = \bar{\phi}$ and x_t is a martingale. Thus, under the Dividends Hypothesis, changes in p_t^L should not be predictable.

We here explore the exercise of forecasting changes in p_t^L with the level of p_t^L itself. Specifically, we consider regressions of the form

$$p_{t+s}^L - p_t^L = \alpha_{p^L,s} + \gamma_{p^L,s} p_t^L + error_{p^L,t+s}. \quad (\text{regression 7})$$

One might be tempted to interpret results from these regressions as a means of evaluating the Dividends Hypothesis. Observe, however, that under the Dividends Hypothesis, x_t is not constant and non-stationary. Thus, [regression 7](#) effectively explores whether non-stationary x_t can predict changes in $x_{t+s} - x_t$. As is well known, regressions of changes in a random walk on its level can yield spurious regression estimates, where changes appear to be predictable in sample even though the truth is that they are not.

For the sake of exploration, we run [regression 7](#) despite the fact that the regression is spurious under the Dividends Hypothesis. We find the results in [Table D.1](#).

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$\hat{\gamma}_{p^L,s}$	-0.089419	-0.29457	-0.54201	-0.80954
S.E.	(0.048319)	(0.092106)	(0.11771)	(0.124)
t-Stat	-1.8506	-3.1981	-4.6047	-6.5285
R^2	0.0359	0.104	0.203	0.353

Table D.1: Results from regressions of the form in [regression 7](#).

A naive interpretation of these regression results would be that they indicate the presence of a large transitory component in the valuation statistic p_t^L , in contradiction to the Dividends Hypothesis.

We instead interpret these findings as an artifact of running spurious regressions. To demonstrate this observation, we Monte Carlo artificial data for x_t with 95 observations constructed as a random walk with standard normal innovations and then run regressions of the form in [regression 7](#) replacing p_t^L with our artificial series for x_t . In [Figures D.1](#) and [D.2](#) we show the simulated histograms of parameter estimates $\hat{\gamma}_{p^L,s}$ from these regressions at horizons $s = 1, 5, 10,$ and 15 . These figures illustrate that the point estimates $\hat{\gamma}_{p^L,s}$ in [Table D.1](#) are well within the range of what one would expect from such regressions under the Dividends Hypothesis. Thus, we do not see these regression results as having any bearing on the validity of that hypothesis.




Figure D.1: Histograms of estimated slope coefficients $\hat{\gamma}_{p^L,s}$ from [regression 7](#) under the assumption that p_t^L is a random walk with standard normal innovations with a sample of length 95. Left panel: Estimates $\hat{\gamma}_{p^L,s}$ for $s = 1$. Right panel: Estimates $\hat{\gamma}_{p^L,s}$ for $s = 5$.




Figure D.2: Histograms of estimated slope coefficients $\hat{\gamma}_{p^L,s}$ from [regression 7](#) under the assumption that p_t^L is a random walk with standard normal innovations with a sample of length 95. Left panel: Estimates $\hat{\gamma}_{p^L,s}$ for $s = 10$. Right panel: Estimates $\hat{\gamma}_{p^L,s}$ for $s = 15$.

E Calculations for the Second Model of $\log(d_t)$

In this section of the Appendix we present details of calculations referred to in the text [Section 7](#).

E.1 Connection to Simple Long Run Risks Model

We now show how to map a simple Long-Run Risks Model with persistent shocks to the expected growth rate of log dividends into the framework of our second model. To do so, we construct the Beveridge-Nelson decomposition of the dynamics of log dividends in such a model.

To start, define

$$\Delta \log(d_{t+1}) = \log(d_{t+1}) - \log(d_t)$$

Let the process for the stationary dynamics of dividend growth be driven by a latent variable z_t that follows an AR1. That is, let

$$\Delta \log(d_{t+1}) = z_{t+1} + \epsilon_{d,t+1}$$

with

$$z_{t+1} = \rho_z z_t + \epsilon_{z,t+1}$$

Recall that we define the Beveridge-Nelson trend of log dividends as $x_t = \lim_{K \rightarrow \infty} \mathbb{E}_t d_{t+K}$ with perhaps some adjustment for a deterministic trend. For simplicity, we do not include such a trend here.

We then have

$$\mathbb{E}_t \Delta \log(d_{t+s}) = \rho_z^s z_t$$

and

$$\mathbb{E}_t \log(d_{t+K}) = \log(d_t) + \sum_{s=1}^K \rho_z^s z_t$$

This calculation gives us a Beveridge-Nelson trend for log dividends

$$\log(x_t) \equiv \lim_{K \rightarrow \infty} \mathbb{E}_t \log(d_{t+K}) = \log(d_t) + \frac{\rho_z}{1 - \rho_z} z_t$$

and a transitory component

$$(\log(d_t) - \log(x_t)) = -\frac{\rho_z}{1 - \rho_z} z_t$$

Note that we can confirm that $\mathbb{E}_t \log(x_{t+1}) = x_t$ by computing

$$\mathbb{E}_t \log(x_{t+1}) - \log(x_t) = \mathbb{E}_t \log(d_{t+1}) - \log(d_t) + \left(\frac{\rho_z}{1 - \rho_z} \right) (\mathbb{E}_t z_{t+1} - z_t) =$$

$$\rho_z z_t + \frac{\rho_z}{1 - \rho_z} (\rho_z - 1) z_t = 0$$

Note that this model has the property that the transitory component of log dividends $\log(d_t) - \log(x_t)$ follows an AR1 process and the trend component $\log(x_t)$ follows a random walk.

E.2 Solving for J_k

We have that for $k = 0$, $J_0 = 0$ by definition and the J_k satisfies the recursion implied by the law of iterated expectations for $k \geq 1$

$$\mathbb{E}_t d_{t+k} = \mathbb{E}_t \mathbb{E}_{t+1} d_{t+k}.$$

Using the form of our guess for $\mathbb{E}_t d_{t+k}$ above, we have

$$\begin{aligned} \exp(\rho^k \log(y_t) + \log(x_t) + J_k) &= \mathbb{E}_t \exp(\rho^{k-1} \log(y_{t+1}) + \log(x_{t+1}) + J_{k-1}) = \\ &= \mathbb{E}_t \exp(\rho^k \log(y_t) + \rho^{k-1} \epsilon_{d,t+1} + \log(x_t) + \epsilon_{x,t+1} + J_{k-1}). \end{aligned}$$

Cancelling the terms $\exp(\rho^k \log(y_t) + \log(x_t))$ gives

$$\begin{aligned} \exp(J_k) &= \exp(J_{k-1}) \mathbb{E}_t \exp(\rho^{k-1} \epsilon_{d,t+1} + \epsilon_{x,t+1}) = \\ &= \exp\left(J_{k-1} + \frac{1}{2}(\rho^2)^{k-1} \sigma_d^2 + \frac{1}{2} \sigma_x^2 + \rho^{k-1} \rho_{dx} \sigma_d \sigma_x\right). \end{aligned}$$

E.3 Derivation of Equation (32)

With strictly positive dividends d_t , we can write realized returns on the fundamental price as

$$R_{t+1}^* \equiv \frac{p_{t+1}^*/d_{t+1} + 1}{p_t^*/d_t} \frac{d_{t+1}}{d_t}$$

Taking logs gives

$$\log(R_{t+1}^*) = \log(p_{t+1}^*/d_{t+1} + 1) - \log(p_t^*/d_t) + \log(d_{t+1}) - \log(d_t)$$

As is standard, we have the first order approximation to realized log returns on the fundamental price taken around a point at which $d_t = x_t$ that

$$\begin{aligned} \log(R_{t+1}^*) &\approx -\log(\beta) + \beta [(\log(p_{t+1}^*) - \log(d_{t+1})) - \log(\beta) + \log(1 - \beta)] - \\ &= (\log(p_t^*) - \log(d_t)) + \log(d_{t+1}) - \log(d_t) \end{aligned}$$

Define a constant

$$\log(\bar{R}) = -\log(\beta) + \beta (\log(1 - \beta) - \log(\beta))$$

We then have the approximation holding for all future realizations of the data

$$\log(p_t^*) - \log(d_t) = \log(d_{t+1}) - \log(d_t) - (\log(R_{t+1}^*) - \log(\bar{R})) + \beta (\log(p_{t+1}^*) - \log(d_{t+1}))$$

We can iterate on this formula to get

$$\begin{aligned} \log(p_t^*) - \log(d_t) &= \sum_{k=0}^{s-1} \beta^k (\log(d_{t+k+1}) - \log(d_{t+k})) - \\ &\sum_{k=0}^{s-1} \beta^k (\log(R_{t+k+1}^*) - \log(\bar{R})) + \beta^s (\log(p_{t+s}^*) - \log(d_{t+s})) \end{aligned}$$

We directly define approximate long horizon log returns on the fundamental price as

$$\log(R_{t,t+s}^*) \equiv \sum_{k=0}^{s-1} \beta^k (\log(R_{t+k+1}^*) - \log(\bar{R}))$$

We report return forecasting regressions using this log-linear approximation to returns as the dependent variable and $\log(p_t^T)$ as the independent variable in Table E.1

horizon	$s = 1$	$s = 5$	$s = 10$	$s = 15$
$\hat{\gamma}_{R^*,s}$	-0.0023671	-0.013212	-0.004707	-0.016221
S.E.	(0.018924)	(0.031822)	(0.041022)	(0.046794)
t-Stat	-0.12509	-0.41518	-0.11474	-0.34664
R^2	0.00017	0.00195	0.000159	0.00154

Table E.1: Estimates from regressions of the form in [regression 4](#) using the log-linear approximation to log returns on the fundamental price as the dependent variable with $\beta = 80/81$ and $\bar{\phi} = -0.9$.

References

- Atkeson, Andrew, Jonathan Heathcote, and Fabrizio Perri. 2024. “Reconciling Macroeconomics and Finance for the US Corporate Sector: 1929 - Present.”
- Bansal, Ravi and Christian Lundblad. 2002. “Market Efficiency, Asset Returns, and the Size of the Risk Premium in Global Equity Markets.” *Journal of Econometrics* 109 (2):195–237.
- Bansal, Ravi and Amir Yaron. 2004. “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles.” *Journal of Finance* 59 (4):1481–1509.
- Barsky, Robert B. and J. Bradford De Long. 1993. “Why Does the Stock Market Fluctuate?” *Quarterly Journal of Economics* 108 (2):291–311.
- Beveridge, Stephen and Charles R. Nelson. 1981. “A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the ‘Business Cycle’.” *Journal of Monetary Economics* 7 (2):151–174.
- Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael R. Roberts. 2007. “On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing.” *Journal of Finance* 62 (2):877–915.
- Campbell, John Y. and John Cochrane. 1999. “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior.” *Journal of Political Economy* 107 (2):205–251.
- Campbell, John Y. and Robert J. Shiller. 1987. “Cointegration and Tests of Present Value Models.” *Journal of Political Economy* 95 (5):1062–1088.
- . 1988. “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors.” *The Review of Financial Studies* 1 (3):195–228.
- Chen, Long, Zhi Da, and Richard Priestly. 2012. “Dividend Smoothing and Predictability.” *Management Science* 58 (10):1834–1853.
- Cochrane, John. 2011. “Presidential Address: Discount Rates.” *Journal of Finance* 66 (4):1047–1108.
- Davydiuk, Tetiana, Scott Richard, Ivan Shaliastovich, and Amir Yaron. 2023. “How Risky Are U.S. Corporate Assets?” *Journal of Finance* 78 (1):141–208.
- Dichev, Ilia D. 2006. “What Are Stock Investors’ Actual Historical Returns? Evidence from Dollar-Weighted Returns.” *American Economic Review* 97 (1):386–401.
- Greenwald, Daniel, Martin Lettau, and Sydney Ludvigson. 2023. “How the Wealth Was Won: Factors Shares as Market Fundamentals.” Working Paper 25769, NBER.
- Gârleanu, Nicolae and Stavros Panageas. 2023. “Heterogeneity and Asset Prices: An Intergenerational Approach.” *Journal of Political Economy* 131 (4):839–876.

- Kleidon, Allan W. 1986. "Variance Bounds Tests and Stock Price Valuation Models." *Journal of Political Economy* 94 (5):953–1001.
- Larrain, Borja and Motohiro Yogo. 2008. "Does Firm Value Move Too Much to Be Justified by Subsequent Changes in Cash Flow?" *Journal of Financial Economics* 87 (1):200–226.
- Leroy, Stephen and Richard Porter. 1981. "The Present-Value Relation: Tests Based on Implied Variance Bounds." *Econometrica* 49 (3):555–574.
- Lucas, Jr., Robert E. 1978. "Asset Prices in an Exchange Economy." *Econometrica* 46 (6):1429–1445.
- Lustig, Hanno N., Stijn Van Nieuwerburgh, and Adrien Verdelhan. 2013. "The Wealth-Consumption Ratio." *The Review of Asset Pricing Studies* 3 (1):38–94.
- Marsh, Terry A. and Robert C. Merton. 1986. "Dividend Variability and Variance Bounds Tests for the Rationality of Stock Market Prices." *American Economic Review* 76 (3):483–498.
- Miller, Merton H. and Franco Modigliani. 1961. "Dividend Policy, Growth, and the Valuation of Shares." *Journal of Business* 34 (4):411–433.
- Morley, James C. 2002. "A State-Space Approach to Calculating the Beveridge-Nelson Decomposition." *Economics Letters* 75 (1):123–127.
- Sharpe, William F. 1964. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *Journal of Finance* 19 (3):425–442.
- Shiller, Robert J. 1981. "Do Stock Prices Move Too Much to be Justified by Subsequent Movements in Dividends?" *American Economic Review* 71 (3):421–436.
- . 2014. "Speculative Asset Prices." *American Economic Review* 104 (6):1486–1517.