Explaining the Variance of Price-Dividend Ratios

John H. Cochrane

Explaining the Variance of Price–Dividend Ratios

John H. Cochrane
University of Chicago

I report a bound on the variance of price–dividend ratios and a decomposition of their variance into terms that reflect changes in dividend growth and discount rates. The specification is not restrictive. The test statistics do not require construction of ex post present values; instead, they are restrictions on means, variances, and covariances of price–dividend ratios, dividend growth, and discount rates. I consider implications for the mean price–dividend ratio, and I evaluate whether a low mean discount rate can rationalize the mean and variance of price–dividend ratios. The results do not indicate any striking rejections of present-value models. However, the bulk of the variance of price–dividend ratios must be accounted for by changing forecasts of discount rates, and discount rates must possess some unusual characteristics.

Stock prices are volatile, yet dividends and the usual measures of the discount rate are relatively smooth. Are prices too volatile? In this article, I reexamine two volatility tests that address this question, a *variance bound* and a *variance decomposition.*

Why reexamine volatility tests? In the view of many

financial economists, price-based volatility tests seem to give qualitatively different information from return-based Euler equation tests. As a result, volatility-test rejections are often interpreted to reject rational asset pricing in general, where Euler-equation rejections simply suggest adjustments to existing asset-pricing models [see Shiller (1989), Cochrane (1991b)]. However, volatility tests have been plagued by controversies over their specification. This situation motivates another look at volatility tests, taking care to overcome as many of the specification problems as possible.

The basic idea of the tests is most easily explained in the context of a constant discount factor present value model:[1]

$$P_t = E_t \sum_{j=1}^{\infty} \rho^j D_{t+j},$$ (1)

where P_t is the stock price at time t, D_t is the dividend at time t, ρ is the discount factor, and E_t is the expectation operator conditional on information at time t. Since $\text{var}(E_t(x)) \leq \text{var}(x)$ for any random variable x, (1) implies the variance bound

$$\text{var}(P_t) \leq \text{var}\left(\sum_{j=1}^{\infty} \rho^j D_{t+j} \right).$$ (2)

Taking the variance inside the sum, (2) can be expressed in terms of the autocovariances of dividends rather than the ex post present value,

$$\text{var}(P_t) \leq \frac{\rho^2}{1 - \rho^2} \sum_{j=-\infty}^{\infty} \rho^{|j|} \text{cov}(D_t, D_{t-j}).$$ (3)

To derive the basic idea of the variance decomposition, multiply both sides of the present-value relation (1) by $P_t - E(P_t)$ and take expectations, yielding

$$\text{var}(P_t) = \text{cov}\left(P_t, \sum_{j=1}^{\infty} \rho^j D_{t+j} \right)$$ (4)

or

$$\text{var}(P_t) = \sum_{j=1}^{\infty} \rho^j \text{cov}(P_t, D_{t+j}).$$ (5)

Equations (4) and (5) reflect the fact that changes in the price must reflect news about future dividends. Below, I derive versions of (4)

---

[1] The first tests by LeRoy and Porter (1981) and Shiller (1981) were roughly in the form of Equations (2) and (4).

and (5) that add a similar discount-rate term similar to the dividend term in (4) and (5), and hence the name "variance decomposition."

1. Specification Issues and Discount Rate Models

A variety of specification issues have plagued volatility tests. In this section, I discuss them, and my approach to their resolution.

1.1. Units

Since price and dividend levels are not stationary, there is no reason to expect sample counterparts of (2)–(5) to hold. Volatility tests are, in fact, quite sensitive to transformations that attempt to achieve stationarity, and many criticisms of the early tests centered on the appropriate transformation of variables. [See, among others, Kleidon (1986), Marsh and Merton (1986), and West (1987, 1988). Shiller (1989) and Gilles and LeRoy (1991) give excellent reviews of this literature.] For this reason, I derive a variance bound and decomposition that use the price–dividend *ratio* and dividend *growth* and discount *rates,* which are more likely to be stationary. This specification follows the practice in most current articles, including Campbell and Shiller (1988), LeRoy and Parke (1990), and Mankiw, Romer, and Shapiro (1991).

1.2 Terminal prices

Many volatility tests are implementations of equations analogous to (2) and (4). [A recent example is Mankiw, Romer, and Shapiro (1991).] However, in a finite sample one must truncate the sums inside the large parentheses of (2) and (4) with a terminal price. Also, one must correct inferences for the severe serial correlation of these sums [see, among others, Flavin (1983), Flood and Hodrick (1990), and Kleidon (1986)].

For this reason, I use statistics that, like (3) and (5), are restrictions on the covariance functions of price–dividend ratios, dividend growth rates, and discount rates, and do not rely on the calculation of ex post present values. Of course, one must also truncate sums of covariances in expressions like (3) and (5), but it is reasonable to hope that the covariances, or *forecasts* of dividend growth and discount rates, die out more quickly than their ex post levels, alleviating the terminal price and truncation problems.

1.3 Time-series restrictions

Many volatility tests restrict the time-series structure of the variables they consider. For example, LeRoy and Parke (1990) assume that dividends follow a pure random walk. However, as Shiller (1989)

notes, negative serial correlation in dividend growth could overturn their variance bound. The bound derived here allows an arbitrary time-series structure for dividend growth, so it overcomes Shiller's objection. Similarly, Campbell and Shiller (1988) infer sums of covariances in expressions similar to (4) from a VAR. Instead, I estimate the covariances in analogues to (3) and (5) directly.

1.4 Mean price–dividend ratio and discount rates

If the mean discount rate is low enough, the present-value model is consistent with an *arbitrarily* high variance of the price–dividend ratio. To see this point, suppose expected dividend growth at $t + 1$ rises, so dividends from $t + 1$ on rise by ΔD. The price then rises by $\Delta P = \Sigma_{j=1}^{\infty} \rho^j \Delta D = \rho \Delta D / (1 - \rho)$. Thus, a small change in dividend growth can have an arbitrarily large effect on the price, if the discount factor ρ is close enough to 1. However, if the discount factor is close to 1, the *mean* price–dividend ratio will also be high. Hence, the mean price–dividend ratio can be used to restrict the range of mean discount rates one can invoke to explain the *variance* of the price–dividend ratio.[2]

For this reason, I examine the model's predictions for the mean price–dividend ratio as well as its variance, and I examine the effect of a range of assumed mean discount rates. Most volatility tests only examine the present-value model's implications for second moments, and simply assume a "reasonable" value for the mean (usually constant) discount rate.

1.5 Time-varying discount rates, bubbles, and fads

Prices can vary, with no change in dividends, if discount rates vary. However, and despite the volume of evidence for time-varying expected returns, most volatility tests are still content to reject the constant discount rate present-value model. [Campbell and Shiller (1988) is one of the few exceptions.] I adopt a specification that allows for a time-varying discount rate, and varying risk premiums across assets.

Statistical rejections of present-value models have been interpreted as support for three rough categories of alternatives. Each alternative amounts to a different statement about discount rates. I treat discount rates in three ways, to address each alternative. The alternatives are as follows.

---

[2] At first glance, one might think that the mean discount rate is easy to determine, but this is not the case. For example, a common candidate is the mean return. But the discount rate is a *single* object that prices *many* assets, and its mean cannot therefore be equal to the mean return on every asset simultaneously.

(i) *Bubbles,* that is, there is *no* discount-rate process that can rationalize the variance of price–dividend ratios. The present-value relation is derived from an Euler equation and a transversality condition. Bubbles occur when the transversality condition fails. Since there always exist discount rates that satisfy the Euler equation, the only way that *no* discount rate can rationalize the present value model is if there is a bubble.

To address this alternative, I calculate the variance bound and decomposition with no assumptions on (unobserved) discount rates beyond stationarity. I also calculate the variance bound and decomposition using returns in the place of discount rates. Since $1 = R_{t+1}^{-1}R_{t+1}$ $= E_t(R_{t+1}^{-1}R_{t+1})$, the inverse return is a "discount factor" that satisfies the Euler equation by construction. A rejection in this case must be attributed to bubbles rather than an incorrect discount factor model.[3]

(ii) *Fads,* that is, there is no *reasonable* discount-rate process that can explain the variance of price–dividend ratios. Here it is admitted that there *exist* unobserved discount-rate processes that can explain the variance of price–dividend ratios. However, it is claimed that any such discount-rate process must have extreme statistical properties that discount-rate processes based on fundamentals are unlikely to have. For example, Poterba and Summers (1988) calculate that discount rates must have a standard deviation of 5.8 percent. They find it "difficult to think of risk factors that could account for such variation in required returns" (p. 51). Grossman and Shiller (1981), Hansen and Jagannathan (1991), Mehra and Prescott (1985), and West (1988) also present calculations that suggest that the standard deviation of discount rates required to reconcile a variety of observations with the present-value model is "too high."

To address this alternative, I calculate mean–standard deviation frontiers for the unobserved discount-rate processes one can invoke to satisfy the volatility tests. Then, the reader can assess whether the required discount rates are "reasonable," according to his or her priors, or the implications of discount-rate models not considered here. This methodology is close to that of Hansen and Jagannathan (1991) and, more loosely, to Watson (1991).

(iii) *Rejection of particular asset-pricing models.* Asset-pricing models tie discount rates to observables. For example, the consump-

---

[3] However, the transversality condition is not testable in a finite sample, so this "test" should be interpreted with caution [see Diba and Grossman (1988), Hamilton and Whiteman (1985), and reviews in Cochrane (1991b) and especially Flood and Hodrick (1990)]. One way to state the problem is that we can never be sure that prices are not responding to changing forecasts of dividends and discount rates in the distant future. Also, the tests are derived under the null hypothesis that no bubbles are present, and many test statistics do not have moments under the alternative. Therefore, the power of the tests is an open question.

tion-based model ties the discount rate to consumption growth, and the CAPM ties it to the market return. Volatility tests can be constructed using specific asset-pricing or discount-rate models, and a rejection can then be interpreted as a rejection of the asset-pricing model. Most of the variance-bound literature tests constant discount rates; Grossman and Shiller (1981) use consumption-based discount rates; Campbell and Shiller (1988) also try a variety of interest rate–based models.

1.6 Approximate present-value model

The exact present-value model expressed in terms of price–dividend ratios, dividend growth, and discount rates is not linear, as (1) is. Therefore, one must either construct tests based on ex post present values using completely specified discount-rate models, or use an approximately present-value model for which analogues to (3) and (5) can be constructed. I follow the latter course.

The approximate present-value model is independently interesting, since it endogenously generates cross-sectional variation in risk premiums, unlike Campbell and Shiller's (1988) similar linearization. The mean price–dividend ratio varies across assets according to their long-run correlation with the discount rate, generating a price version of standard expected return models. In principle, this feature allows one to conduct present-value tests on many assets simultaneously, as Euler equation tests are routinely conducted. However, I do not exploit cross-sectional predictions in this article.

2. Present-Value Model, Variance Bound, and Variance Decomposition

This section derives and interprets analogues to the present value (1), variance bound (3), and variance decomposition (5) that incorporate the specification issues mentioned above.

Start with a general one-period Euler equation:

$$1 = E_t(\gamma_{t+1}R_{t+1}), \tag{6}$$

where $R_{t+1} = (P_{t+1} + D_{t+1})/P_t$ and γ_{t+1} is the discount factor or intertemporal marginal rate of substitution.[4] (Appendix A contains a list of all the symbols.) Iterating (6), together with the transversality condition

---

[4] This equation is "general" since a discount factor that satisfies (6) exists under weak no-arbitrage conditions. See Hansen and Richard (1987). Hence, every asset-pricing model implies a representation of the form (6).

$$\lim_{j \to \infty} E_t \left(\prod_{k=1}^{j} \gamma_{t+k} \right) P_{t+j} = 0,$$

we obtain a similar general present-value model,

$$P_t = E_t \left(\sum_{j=1}^{\infty} \left(\prod_{k=1}^{j} \gamma_{t+k} \right) D_{t+j} \right).$$

However, prices and dividends are not stationary, so moments calculated from this present-value model cannot be estimated by sample counterparts. To reexpress the present-value model as a relation between stationary variables, divide both sides by dividends,

$$\frac{P_t}{D_t} = E_t \sum_{j=1}^{\infty} \left(\prod_{k=1}^{j} \gamma_{t+k} \right) \frac{D_{t+j}}{D_t} = E_t \sum_{j=1}^{\infty} \prod_{k=1}^{j} (\gamma_{t+k} \eta_{t+k}), \qquad (7)$$

where $\eta_t = D_t/D_{t-1}$. Next, express the result in terms of log discount rates, $g_t = -\ln(\gamma_t)$, and dividend growth rates, $n_t = \ln(\eta_t)$:

$$\frac{P_t}{D_t} = E_t \sum_{j=1}^{\infty} \exp\left(\sum_{k=1}^{j} (n_{t+k} - g_{t+k}) \right). \qquad (8)$$

Under the assumption that price–dividend ratios, dividend growth, and discount rates or factors are stationary, we can form tests based on the implications of Equations (7) and (8) for moments of these variables.[5]

Equations (7) and (8) are not linear in γ_t and η_t or g_t and n_t as (1) is linear in D_t, so we cannot take the expectation and variance operators inside the present value as we did to get rid of the ex post present values in (3) and (5). To achieve the same result, I use an approximation to the present-value model (8) for which analogues to (3) and (5) can be derived. The approximate model is

$$\frac{P_t}{D_t} \cong \frac{\Omega}{1 - \Omega} + \frac{\Omega}{2(1 - \Omega)^2} \sum_{j=-\infty}^{\infty} \Omega^{|j|} \operatorname{cov}(n_t - g_t, n_{t-j} - g_{t-j})$$

$$+ \frac{1}{1 - \Omega} E_t \left(\sum_{j=1}^{\infty} \Omega^j (\tilde{n}_{t+j} - \tilde{g}_{t+j}) \right), \qquad (9)$$

where $\tilde{n}_t = n_t - E(n)$, $\tilde{g}_t = g_t - E(g)$, and $\Omega = e^{E(n) - E(g)}$.

Equation (9) is derived in Appendix B. Basically, I take a second-order Taylor expansion of the term inside the expectation operator in (8) with respect to n_{t+k} and g_{t+k}, about their unconditional means

---

[5] Equations (7) and (8) are *consistent* with the assumption that price–dividend ratios, dividend growth rates, and discount rates are stationary. In Appendix B, I discuss conditions under which stationary dividend growth and discount rates *imply* a stationary price–dividend ratio.

$E(n)$ and $E(g)$. This expression is linear and quadratic in $n_{t+k} - E(n)$ and $g_{t+k} - E(g)$. Taking expectations, I obtain an expression in means, variances, and covariances of n and g, Equation (9). In Appendix B, I also argue that the approximate model is reasonably accurate, and compare it to the similar approximate model derived by Campbell and Shiller (1988).

Taking the unconditional expected value of (9), we obtain the mean price–dividend ratio,

$$E\left(\frac{P}{D}\right) = \frac{\Omega}{1 - \Omega} + \frac{\Omega}{2(1 - \Omega)^2} \sum_{j=-\infty}^{\infty} \Omega^{|j|} \operatorname{cov}(n_t - g_t, n_{t-j} - g_{t-j}). \quad (10)$$

The approximate model (9) consists of three terms, and the mean price–dividend ratio (10) contains the first two of those terms. The first term is the price–dividend ratio in a certainty world with constant dividend growth $E(n)$ and constant discount rate $E(g)$. As the mean discount rate $E(g)$ approaches the mean dividend growth rate $E(n)$, Ω approaches 1, so this constant term (and the other terms as well) diverges to infinity. This behavior is not an artifact of the approximation: the exact present-value model [Equations (7) and (8)] also diverges to infinity when $E(g) \rightarrow E(n)$.[6]

The second term in (9) and (10) adjusts the mean price–dividend ratio for the covariance of dividend growth with discount rates. This term generates cross-sectional risk premiums, just as the covariance with the market drives cross-sectional risk premiums in the CAPM. If the dividend growth of one of two otherwise identical assets has greater covariance with the discount rate, that asset will have a lower price–dividend ratio and hence a higher average return. More precisely, the weighted sum of covariances, or the covariance between long-run or low-frequency movements in dividend growth and discount rates, drives cross-sectional risk premiums in price–dividend ratios.[7]

The third term on the right-hand side of (9) captures variation in price–dividend ratios over time due to changing forecasts of dividend growth and discount rates. The price–dividend ratio rises above its

---

[6] See Appendix C. This observation motivated the Taylor expansion about $E(g)$ and $E(n)$, rather than some other value of g and n. With any of the latter, the exact and approximate present-value models would have diverged to infinity at different values.

[7] We have

$$\frac{1}{1 - \Omega^2} \sum_{j=-\infty}^{\infty} \Omega^{|j|} \operatorname{cov}(n_t - g_t, n_{t-j} - g_{t-j})$$

$$= \operatorname{var}\left(\sum_{j=0}^{\infty} \Omega^j n_{t-j}\right) + \operatorname{var}\left(\sum_{j=0}^{\infty} \Omega^j g_{t-j}\right) - 2\operatorname{cov}\left(\sum_{j=0}^{\infty} \Omega^j n_{t-j}, \sum_{j=0}^{\infty} \Omega^j g_{t-j}\right).$$

These terms are also the spectral and cross-spectral densities in a window near frequency zero.

mean if dividend growth rates are forecast to rise, or if discount rates are forecast to fall. Since the sums start at $j = 1$, only *predictable* changes in the discount rate and dividend growth rate can explain variation in the price–dividend ratio: if the discount rate and dividend growth rate are variable but unforecastable, the price–dividend ratio should be a constant. Also, it has seemed puzzling to some authors that stock and bond prices do not move together, but (9) shows that changing expectations of *future* discount rates can induce changes in the price–dividend ratio, with no change in current bond prices, interest rates, or expected returns.

The approximate model (9) does not capture changes in the *conditional* covariance of dividend growth and discount rates, or *combinations* of time-series and cross-sectional variation. This desirable feature requires a third-order expansion, which proved algebraically intractable. As a result, the model has only a limited ability to capture time-varying risk premiums.

Since for any X, $\mathrm{var}(E_t(X)) \leq \mathrm{var}(X)$, (9) implies the variance bound

$$\mathrm{var}\!\left(\frac{P}{D}\right) \leq \frac{1}{(1-\Omega)^2}\,\mathrm{var}\!\left(\sum_{j=1}^{\infty} \Omega^j(\tilde{n}_{t+j} - \tilde{g}_{t+j})\right), \tag{11}$$

in analogy to (2). Taking the variance inside the sum, we obtain the variance bound in terms of second moments:

$$\mathrm{var}\!\left(\frac{P}{D}\right) \leq \frac{\Omega^2}{(1-\Omega^2)(1-\Omega)^2}\sum_{j=-\infty}^{\infty} \Omega^{|j|}\,\mathrm{cov}(n_t - g_t,\, n_{t-j} - g_{t-j}), \tag{12}$$

in analogy to (3).

The variance bound (12) is a function of a weighted sum of covariances. This observation explains why variance bounds have been so sensitive to detrending methods. If discount rates are constant and log dividend *levels* are stationary, then the variance of long-moving averages of log dividend *growth* (their spectral density at low frequencies) tends to 0. In this case, the weighted moving average on the right-hand side of the variance bound (12) is close to 0, practically by construction.

The empirical work checks the implications of the mean and variance equations *together*. Since the variance and covariance term is the same in the mean price–dividend ratio (10) and variance bound (12), we can substitute out that term, so the content of the pair (10) and (12) can be summarized by the mean price–dividend ratio (10) together with

$$\mathrm{var}\!\left(\frac{P}{D}\right) \leq \frac{2\Omega}{1-\Omega^2}\left(E\!\left(\frac{P}{D}\right) - \frac{\Omega}{1-\Omega}\right). \tag{13}$$

251

This is a variance bound that holds for *arbitrary* discount rates.

Premultiplying (9) by $P_t/D_t - E(P/D)$ and taking expectations yield the variance decomposition, in analogy to (4) and (5),

$$\text{var}\left(\frac{P}{D}\right) = \frac{1}{1 - \Omega} \sum_{j=1}^{\infty} \Omega^j \text{cov}\left(\frac{P_t}{D_t}, n_{t+j}\right)$$

$$+ \frac{1}{1 - \Omega} \sum_{j=1}^{\infty} \Omega^j \text{cov}\left(\frac{P_t}{D_t}, -g_{t+j}\right). \tag{14}$$

This equation decomposes the variance of the price–dividend ratio into terms that reflect changing forecasts of future dividend growth and changing forecasts of future discount rates, using the price–dividend ratio as the forecasting variable. It is not an orthogonal decomposition, so terms less than 0 percent and greater than 100 percent are possible. High price–dividend ratios may be associated with low future dividend growth if they are also associated with *much* lower future discount rates.

Note that we can write the sums of covariances in (14) as

$$\sum_{j=1}^{\infty} \Omega^j \text{cov}\left(\frac{P_t}{D_t}, x_{t+j}\right) = \text{cov}\left(\frac{P_t}{D_t}, \sum_{j=1}^{\infty} \Omega^j x_{t+j}\right), \qquad x = n \text{ or } g.$$

The right-hand expression is the numerator of a regression coefficient of long-horizon movements in x_t on price–dividend ratios. Since $\Omega \cong 1$, the facts we document about price–dividend ratios' ability to forecast a geometrically weighted sum of x's can generally be interpreted in terms of their ability to forecast the usual truncated but unweighted long-horizon sums, as in Fama and French (1988). Thus, we can understand (14) as a restriction between the variance of price–dividend ratios and the strength of price–dividend ratio–based forecasts of dividend growth and discount rates.[8]

3. Discount Rates

3.1 Discount-rate models

I consider the following models for the discount rate g_t.

(i) *Discount rate = a constant.* I include this model because it is widely tested in the volatility test literature, and it helps to understand

---

[8] An alternate way of expressing (14) is $1 = \beta_n - \beta_g$, where β_n is the regression coefficient in

$$\sum_{j=1}^{\infty} \Omega^j n_{t+j} = \alpha + \beta_n \frac{P_t}{D_t} + \text{error},$$

and similarly for β_g.

how the more plausible models work. The constant discount rate model predicts no risk premiums—expected returns on all assets are the same.

(ii) *Discount rate = reference return plus risk premium.* There is a long tradition in the volatility test and investment or capital-budgeting literature that measures time-varying discount rates from interest rates plus risk premiums that are constant over time. To capture this idea, we can model discount rates as a reference return r_t^0 (a constant value, or data on an interest rate) plus an unobserved i.i.d. (not necessarily mean 0) random variable ϵ_t,

$$g_t = r_t^0 + \epsilon_t, \tag{15}$$

where $\text{cov}(z_{t-j}, \epsilon_t) = 0$ for any z_{t-j}.

With this model, discount rates in the variance decomposition (14) can be measured, since $\text{cov}(P_t/D_t, g_{t+j}) = \text{cov}(P_t/D_t, r_{t+j}^0)$. The ϵ_t term in (15) generates cross-sectional risk premiums: the mean price–dividend ratio of otherwise identical assets varies as the correlation of their dividend growth with ϵ_t varies.

(iii) *Consumption-based discount rates.* With constant relative risk aversion, time-separable utility,

$$E\left(\sum_{t=0}^{\infty} \rho^t u(c_t)\right) = E\left(\sum_{t=0}^{\infty} \rho^t \frac{c_t^{1-\alpha} - 1}{1 - \alpha}\right),$$

the discount rate is

$$g_t = -\ln[\rho u'(c_t)/u'(c_{t-1})] = -\ln(\rho) + \alpha \ln(c_t/c_{t-1}). \tag{16}$$

(iv) *Discount rate = return.* As discussed in Section 1, we can test for bubbles by using returns in the place of discount rates, since $1 = E_t(R_{t+1}^{-1}R_{t+1})$. Taking logs, this discount rate model is just $g_t = r_t$. Imposing the transversality condition, the model "discount rate = return" is vacuous, so it can also be interpreted as a measure of the accuracy of the approximate present-value equation (9).

3.2 Bounds on the mean discount rate

Even if we treat discount rates as unobservable, the mean discount rate is not a totally free parameter. It must satisfy three lower bounds. These bounds are important, since each of the statistics (10), (12) and (13), and (14) are sensitive functions of the assumed mean discount rate, through Ω. They are reported in the graphs of the results that follow.[9]

---

[9] The bounds are estimated with relatively large standard errors. Few of the point estimates are more than two standard errors above zero, and all are less than two standard errors above 4 percent. Thus, the data are consistent with a wide variety of mean discount rates.

First, the mean discount rate must be greater than the mean dividend growth rate of every asset or portfolio, or the present value is infinite. Thus, $E(g) > E(n)$. Second, the variance terms in the mean price–dividend ratio equation (10) must be positive. Hence, the mean price–dividend ratio can be no lower than $\Omega/(1 - \Omega)$. Solving for $E(g)$ gives

$$E(g) \geq E(g)_{\min} \equiv E(n) - \ln\left(\frac{E(P/D)}{1 + E(P/D)}\right).$$

Third, the mean discount rate must be greater than the mean log return on any asset or portfolio of assets. To see this, start with the Euler equation (6). With $r = \ln(R)$ and $g = \ln(\gamma)$,

$$1 = E(\gamma R) = E(e^{r-g}) \geq e^{E(r-g)}.$$

Taking logs, $E(g) \geq E(r)$.

4. Variance Bounds

The data are based on the value-weighted and equally weighted NYSE portfolios maintained by the Center for Research in Security Prices (CRSP). The sample consists of annual data from 1926 to 1988. The data are described in detail in Appendix F.

4.1 Constant discount rate
The mean price–dividend ratio (10) and variance bound (12) calculated with constant discount rates are presented in Figure 1. Both decline as the discount rate rises, mostly because of the leading term in Ω.

Variance bounds tests typically pick a value for the (constant) discount rate and check whether the variance bound is satisfied at that value. Figure 1 shows that this procedure will not lead to a violation of the variance bound in the value-weighted portfolio so long as the chosen discount rate is less than about 6 percent.

To test the mean price–dividend ratio equation (10) and the variance bound (12) *together* (still assuming constant discount rates), we can solve (10) for the mean discount rate. This value is marked "var(g) = 0" in Figure 1 and the following graphs. Then, we verify that the variance bound is satisfied at that value. (The same result for the equally weighted portfolio is discussed below.)

4.2 Unobserved discount rates
The variance bound with no restrictions on discount rates [Equation (13)] is presented in Figure 2. The variance bound (13) of Figure 2 *rises* as a function of the mean discount rate, where the variance bound
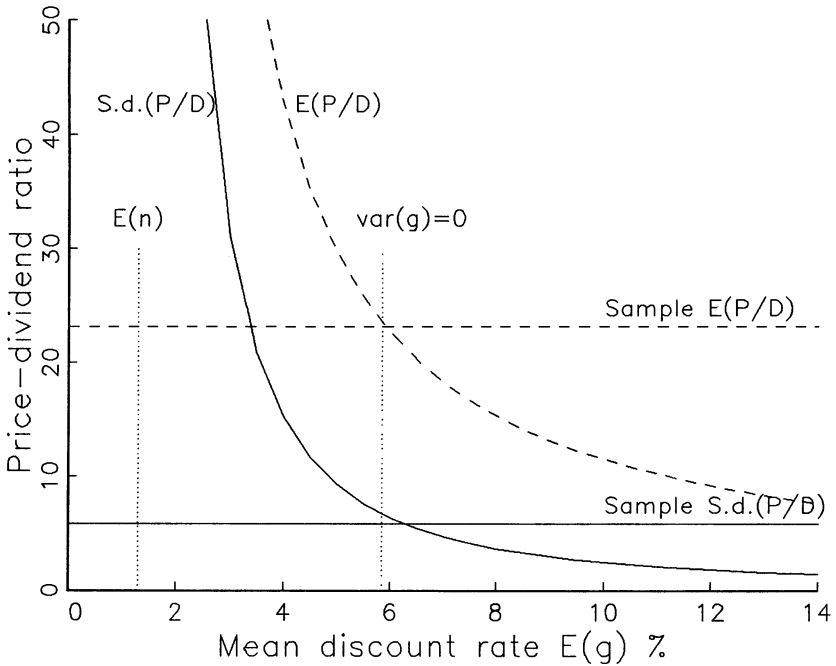
Figure 1
Mean price–dividend ratio and variance bound with constant discount rates, value-weighted NYSE
The dashed curve marked "$E(P/D)$" reports the predicted mean price–dividend ratio with constant discount rates,

$$E\left(\frac{P}{D}\right) = \frac{\Omega}{1 - \Omega} + \frac{\Omega}{2(1 - \Omega)^2} \sum_{j=-15}^{15} \Omega^{|j|} \, \text{cov}(n_t, \, n_{t+j}), \qquad (10)$$

where $\Omega = e^{E(n) - E(g)}$, n is the log of real dividend growth, and $E(g)$ is the assumed constant discount rate given on the horizontal axis. For comparison, the dashed horizontal line marked "Sample $E(P/D)$" reports the sample mean price–dividend ratio.

The solid curve marked "S.d.(P/D)" reports the variance bound with constant discount rates,

$$\text{SD}\left(\frac{P}{D}\right) \leq \left[\frac{\Omega^2}{(1 - \Omega^2)(1 - \Omega)^2} \sum_{j=-15}^{15} \Omega^{|j|} \, \text{cov}(n_t, \, n_{t+j}) \right]^{1/2} \qquad (12)$$

For comparison, the solid horizontal line marked "Sample S.d.(P/D)" gives the sample standard deviation of the price–dividend ratio.

The vertical line labeled "$E(n)$" marks the discount rate equal to mean dividend growth. The vertical line labeled "var$(g) = 0$" marks the constant discount rate at which the predicted mean price–dividend ratio equals its sample value. Since the variance bound is greater than the sample variance at this mean discount rate, the variance bound and mean price–dividend ratio together do not reject the constant-discount-rate model. Data are annual, 1926–1988.

with constant discount rates in Figure 1 *fell*. This feature is due to the fact that the mean price–dividend ratio equation (10) is always implicitly satisfied in using the bound (13). For example, consider the leftmost point of Figure 2. Here, the mean discount rate is $E(g)_{\text{min}}$, at which $E(P/d) = \Omega/(1 - \Omega)$. For the mean price–dividend ratio equation (10) to be satisfied here, we must implicitly assume that the discount rate equals the dividend growth rate, so that the term in

Value weighted NYSE

Equally weighted NYSE
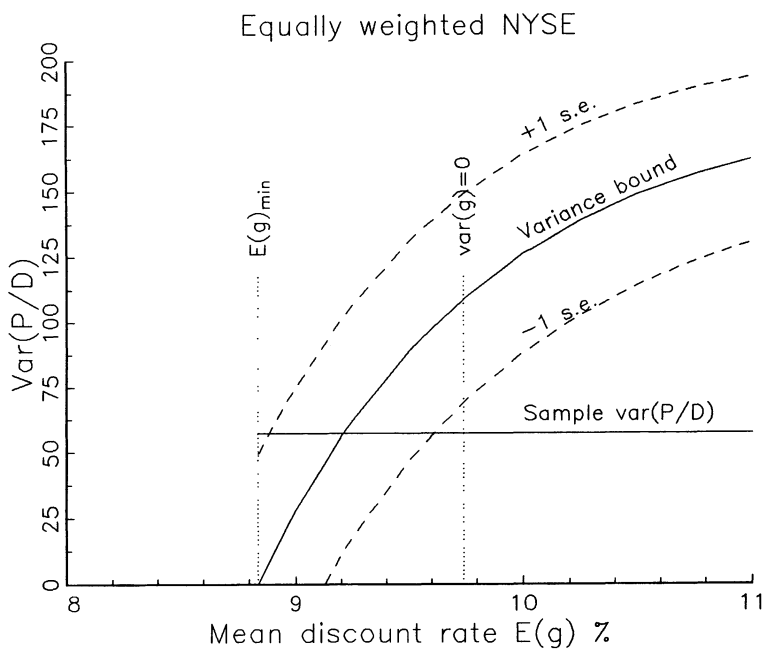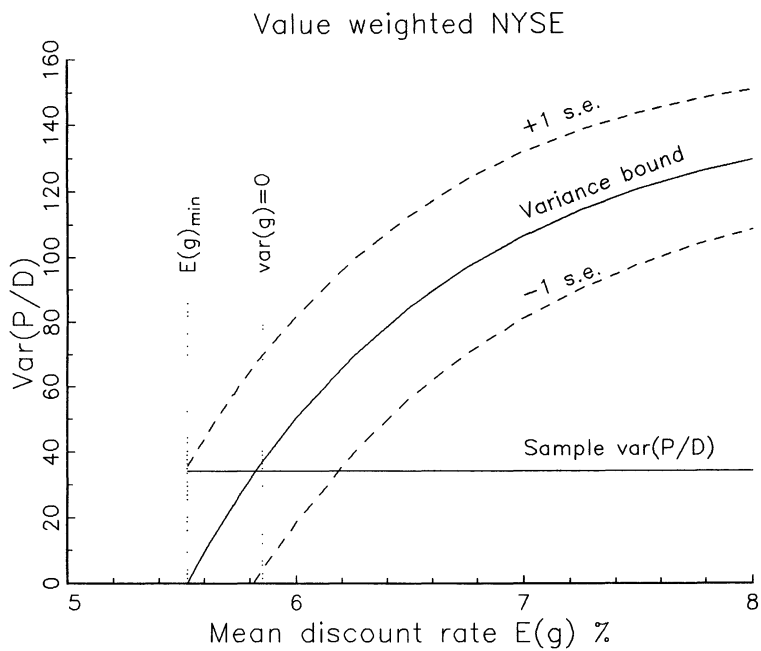
covariances on the right-hand side of (10) is zero. But this means the covariance term on the right-hand side of the variance bound (12) is also zero, which is why the variance bound (13) graphed in Figure 2 is zero. Similarly, to satisfy (10) at high mean discount rates, one must assume that unobserved discount rates generate a large covariance term, and this raises the variance bounds (12) and (13).

Thus, there is always a region of mean discount rates near $E(g)_{min}$ in which the point estimate of the variance bound (13) is violated. However, even where the point estimate of (13) is below the point estimate of the variance of the price–dividend ratio, it is never even one standard error below, so we do not reject the present-value model for any value of the mean discount rate at which it can be constructed [above $E(g)_{min}$].

The bound with no restrictions on discount rates is the same as the bound with constant discount rate at the mean discount rate marked "$var(g) = 0$," since that discount rate solves (10). The sample variance is below the bound at "$var(g) = 0$" in Figure 2, so the variance bound does not reject the constant-discount-rate model for either portfolio.

4.3 Discount rate models
One might suppose that the variance bound with any time-varying discount-rate model is automatically satisfied because the variance bound with constant discount rates is satisfied. This is not true. The variance bound with time-varying discount rates can be lower than the variance bound with constant discount rates. If the discount-rate process is positively correlated with dividend growth, it lowers the covariance terms in (12) relative to their value with constant discount rates. If discount rates *equal* dividend growth rates (log utility, consumption = dividend), then the variance bound with time-varying discount rates is *zero*.

However, a discount-rate model can only restrict the variance bound and mean price–dividend ratio together if it restricts the mean dis-

---

←

Figure 2
Variance bound with no restrictions on discount rates
The "variance bound" gives the right-hand side of

$$\text{var}\left(\frac{P}{D}\right) \leq \frac{2\Omega}{1 - \Omega^2}\left(E\left(\frac{P}{D}\right) - \frac{\Omega}{1 - \Omega}\right), \tag{13}$$

where $\Omega = e^{E(n) - E(g)}$, n is the log of dividend growth, and $E(g)$ is the assumed mean discount rate given on the horizontal axis. The ± 1 standard error lines show standard errors for ($\text{var}(P/D) -$ bound). The "Sample $\text{var}(P/D)$" line gives the sample variance of the price–dividend ratio.

The vertical line labeled "$E(g)_{min}$" marks the minimum value of the mean discount rate consistent with the mean price–dividend ratio [i.e., where $E(P/D) = \Omega/(1 - \Omega)$]. The vertical line labeled "$var(g) = 0$" marks the discount rate at which the predicted mean price–dividend ratio equals its sample value with constant discount rates. Data are annual, 1926–1988.

Table 1
Variance bound and mean price–dividend ratio with discount rate = return

	Value-weighted NYSE			Equally weighted NYSE		
	Mean	Vari-ance	Vari-ance bound	Mean	Vari-ance	Vari-ance bound
Sample	23.09	34.02		25.96	57.01	
Standard error	1.71	7.75		1.96	15.80	
Model	22.56		49.73	27.81		106.76
Standard error	1.52		31.42	4.03		96.06
Standard error (sampleless model)	2.13			4.77		

"Sample" gives the sample mean and variance of the price–dividend ratio, annual data 1926–1988. The "model," "mean," and "bound" give

$$E\left(\frac{P}{D}\right) = \frac{\Omega}{1 - \Omega} + \frac{\Omega}{2(1 - \Omega)^2} \sum_{j=-15}^{15} \Omega^{|j|} \operatorname{cov}(n_t - g_t, n_{t-j} - g_{t-j}), \tag{10}$$

$$\operatorname{var}\left(\frac{P}{D}\right) \leq \frac{\Omega^2}{(1 - \Omega^2)(1 - \Omega^2)} \sum_{j=-15}^{15} \Omega^{|j|} \operatorname{cov}(n_t - g_t, n_{t-j} - g_{t-j}), \tag{12}$$

respectively, where g is the log of return, n is the log of dividend growth, and $\Omega = \exp(E(n) - E(g))$. The estimated covariances also include a Bartlett weighting to ensure that their sum is positive. All standard errors contain a Hansen (1982)/Newey–West (1987) correction, as detailed in Appendix G. Data are annual, 1926–1988.

count rate, since only the mean discount rate enters (13). For example, the constant-discount-rate model had one free parameter, $E(g)$, which was estimated by (10) as the value marked "var$(g) = 0$." The constant-risk-premium and interest-rate-plus-constant-risk-premium discount-rate models leave more than one free parameter, so these discount-rate models do not restrict the variance bound beyond the values presented above with no assumptions on discount rates.

Calculations of the variance bound (12) and the mean price–dividend ratio (10) with discount rate = return are presented in Table 1. The predicted mean price–dividend ratios [the right-hand side of (10)] differ from sample means by about one standard error of the sample mean. The variance bounds [the right-hand side of (12)] are well above the sample variance in both cases. Thus, there is no indication of either bubbles or serious defects of the approximate model.

In Table 2, variance bounds are presented using consumption-based discount rates. Since there are two free parameters, the variance bound is presented for a variety of assumed subjective discount factors ρ. For each ρ, I calculated a risk-aversion parameter α to solve the mean price–dividend equation (10). Then, I estimated the variance bound (12) or (13) at the mean discount rate corresponding to the assumed ρ and the estimated α. In all cases, the point estimate of the variance bound is substantially greater than the sample variance of the price–dividend ratio.

The feature that drives this result is the low correlation of long-

Table 2
Variance bound using consumption-based discount rates

	Value-weighted NYSE					Equally weighted NYSE				
	Utility parameters		Mean discount rate	Variance bound			Utility parameters		Mean discount rate	Variance bound	
ρ	α	$E(g)$ (%)	Bound	SE		ρ	α	$E(g)$ (%)	Bound	SE
.99	4.28	6.50	84.76	10.55		.99	8.71	12.21	174.22	5.49
.97	2.32	6.03	53.13	10.84		.97	6.43	11.32	167.30	7.58
.95	0.60	5.90	41.50	10.95		.95	4.20	10.54	149.71	9.71
	[Sample var(P/D) = 34.02]						[Sample var(P/D) = 57.01]			

Discount rates are generated by the consumption-based model:

$$g_t = -\ln(\rho) + \alpha \ln(c_t/c_{t-1}).$$

For a given choice of ρ, α is calculated to satisfy the mean price–dividend ratio

$$E\left(\frac{P}{D}\right) = \frac{\Omega}{1-\Omega} + \frac{\Omega}{2(1-\Omega)^2} \sum_{j=-15}^{15} \Omega^{|j|} \mathrm{cov}(n_t - g_t, n_{t-j} - g_{t-j}), \tag{10}$$

where n is the log dividend growth and $\Omega = e^{E(n)-E(g)}$. The variance bound is then calculated by

$$\mathrm{var}\left(\frac{P}{D}\right) \le \frac{\Omega^2}{(1-\Omega^2)(1-\Omega^2)} \sum_{j=-15}^{15} \Omega^{|j|} \mathrm{cov}(n_t - g_t, n_{t-j} - g_{t-j}). \tag{12}$$

[Equivalently, the mean discount rate $E(g)$ is calculated and used in the variance bound (13).] All standard errors contain a Hansen (1982)/Newey–West (1987) correction, as detailed in Appendix G. Data are annual, 1926–1988.

run (low-frequency) movements in consumption growth and dividend growth. If they were highly and positively correlated, the covariance term on the right-hand side of (10) would have been small, so I would have estimated an α that implied a low mean discount rate to satisfy the mean price–dividend ratio (10). In turn, that low mean discount rate would have implied a violation of the variance bound (13).

5. Variance Decomposition

The fraction of the variance of price–dividend ratios accounted for by dividend growth, interest rates, and returns, as a function of the assumed mean discount rate, $E(g)$, is presented in Figure 3.

The stylized fact that drives the results is the size of price–dividend ratio forecasts of long-horizon dividend growth and discount-rate measures. For example, if a rise in price–dividend ratios forecast a large rise in future dividend growth, then the "fraction due to dividend growth" will be high.

5.1 Unobserved discount rates
With no assumptions about discount rates, all we can do is calculate the fraction of the variance of the price–dividend ratio due to dividend

Value weighted NYSE

Equally weighted NYSE
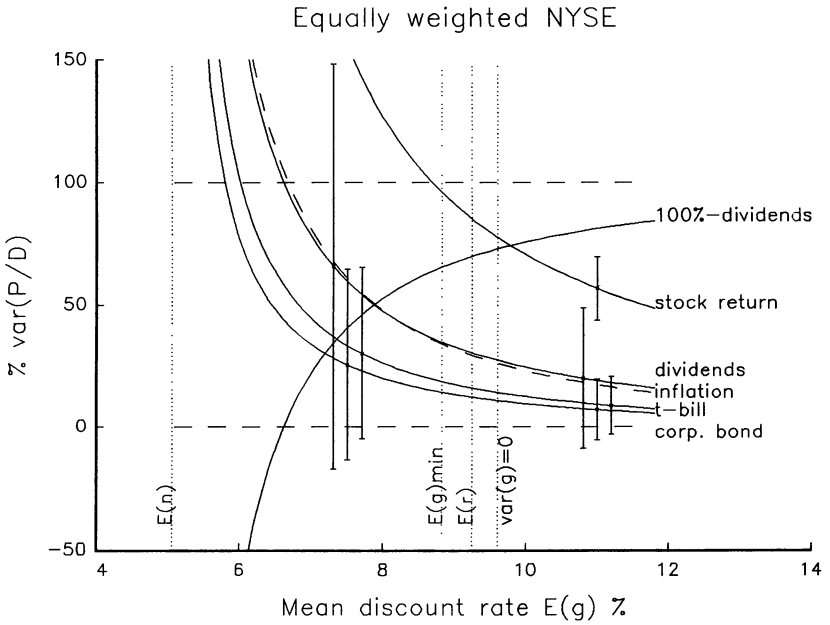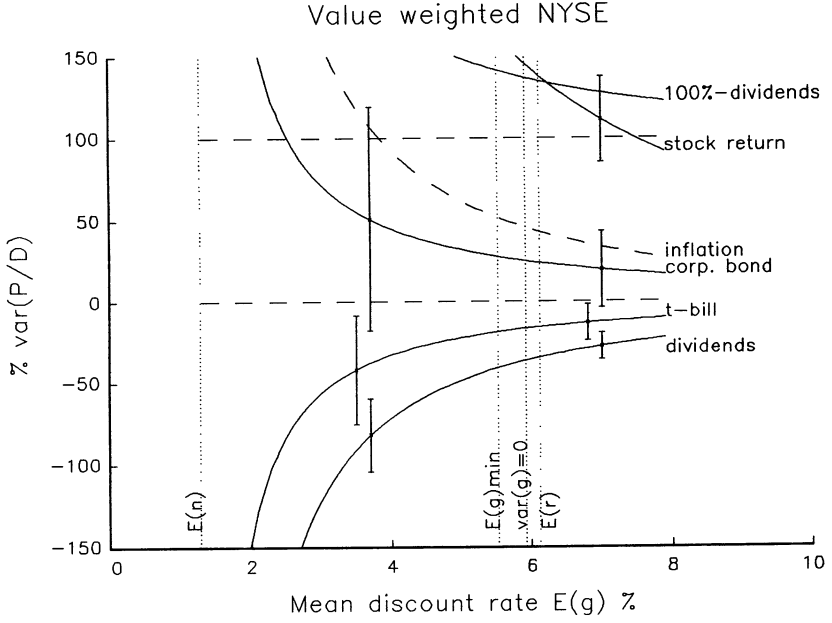
growth, and infer the fraction that must be attributed to unobserved discount rates. Thus, examine the lines marked "dividends" and "100% − dividends" in Figure 3.

For the equally weighted (EW) price–dividend ratio, a rise in the price–dividend ratio forecasts a rise in dividend growth, as we expect. However, the value-weighted (VW) price–dividend ratio actually forecasts a *decline* in real dividend growth. Thus, real dividend growth accounts for a *negative* fraction of the variance of the VW price–dividend ratio, and discount-rate changes must account for more than 100 percent.

As the mean discount rate $E(g)$ declines, Ω approaches 1, so the terms of the variance decomposition diverge to infinity. Thus, so long as a rise in the price–dividend ratio forecasts a rise in dividend growth, there is *some* mean discount rate at which dividend growth accounts for *all* the variance of the price–dividend ratio (there is a point where the "dividends" line crosses the 100 percent reference line, and the "100% − dividends" line crosses zero). For the EW portfolio, this is at about $E(g) = 6.5\%$. For larger values of the mean discount rate, increasing fractions of the variance of price–dividend ratios must be attributed to changing discount rates—the "100% − dividends" line rises. In particular, the largest lower bound on mean discount rates is the mean return $E(r)$; at this mean discount rate only 30 percent of the variance of the price–dividend ratio is accounted for by dividend growth. The remaining 70 percent must be due to changing forecasts of discount rates.

The relatively large standard errors in Figure 3 are also noteworthy. Price–dividend ratios only forecast small movements in dividend

---

←

Figure 3
Variance decomposition
Each curve shows the percent of the variance of the price–dividend ratio attributed to dividend growth or a discount rate series. For a given mean discount rate $E(g)$, these are calculated by

$$\% \, \mathrm{var}\left(\frac{P}{D}\right) = 100 \times \frac{1}{1 - \Omega}\left[\sum_{j=1}^{15} \Omega^j \mathrm{cov}\left(\frac{P_t}{D_t}, x_{t+j}\right)\right]\left[\mathrm{var}\left(\frac{P}{D}\right)\right]^{-1}$$

[see Equation (14)], where $\Omega = e^{(E(n) - E(g))}$, n is the log of dividend growth, and x is the dividend growth, Treasury-bill returns, corporate bond returns, or stock returns, as indicated. The dashed curve reports the same calculation with $x =$ inflation.

The curve labeled "100% − dividends" reports 100% − the percent of the variance of the price–dividend ratio attributed to dividend growth. The discount rate curves should intersect the "100% − dividends" curve at the true mean discount rate.

The horizontal dashed lines indicate 0 percent and 100 percent for reference. The error bars report selected one standard error intervals.

"$E(n)$" marks the mean discount rate equal to the mean dividend growth rate. "$E(g)_{\min}$" marks the minimum value of the mean discount rate consistent with the mean price–dividend ratio [the value of $E(g)$ at which $E(P/D) = \Omega/(1 - \Omega)$]. "var$(g) = 0$" marks the constant discount rate at which the predicted mean price–dividend ratio equals its sample value with constant discount rates. "$E(r)$" marks the mean log return. Data are annual, 1926–1988.

growth rates. On the other hand, only small forecasts are required: dividend growth and discount rate forecasts on the order of one standard error higher than the point estimates in more than 60 years of data would be sufficient to explain the variance of price–dividend ratios at relatively high mean discount rates.[10]

5.2 Discount-rate models

Constant discount rate. With a constant discount rate, the mean price–dividend ratio equation (10) can be solved for a unique value of the discount rate, marked "var(g) = 0" in Figure 3. If the constant-discount-rate model is correct, dividend growth should account for 100 percent of the variance of the price–dividend ratio at this discount rate—the "dividends" line should intersect the 100 percent reference line and the "100% − dividends" line should intersect zero at the mean discount rate "var(g) = 0."

Dividend growth accounts for a negative fraction of the variance of the VW price–dividend ratio at the "var(g) = 0" mean discount rate. Thus, the variance decomposition rejects the constant-discount-rate model for the VW portfolio.[11] Dividend growth accounts for about 30 percent of the variance of the EW price–dividend ratio at the "var(g) = 0" mean discount rate. This fraction is less than two standard errors from 0 percent and 100 percent. Thus, the constant-discount-rate model is not statistically rejected for the EW portfolio, but neither is the view that dividend growth accounts for 0 percent of the variance of the EW price–dividend ratio.[12]

---

[10] In deriving Figure 3, I used real dividend growth and discount rates. However, price–dividend ratios should not forecast inflation, so it should not matter whether we use real or nominal dividend growth and discount rates. However, price–dividend ratios *do* forecast inflation, so there is a significant difference between nominal and real results.

Included in Figure 3 is the weighted sum of covariances of the price–dividend ratio with subsequent inflation. To obtain the fraction of the variance of the price–dividend ratio explained by *nominal* dividend growth, add the inflation line to the real dividend growth line in Figure 3; to obtain the fraction explained by a nominal interest rate, subtract the inflation line from the real interest rate lines. (The sum of the two fractions is unchanged; the use of nominal data just shifts the covariance with inflation from the discount rate to the dividend growth term.)

Since both price–dividend ratios forecast rises in inflation, the fraction of the variance accounted for by nominal dividend growth is higher than for real dividend growth, and is positive in both cases. The mean discount rates at which 100 percent of the variance is accounted for by dividends is then higher, $E(g)$ = 2 percent and 7.5 percent for the VW and EW portfolios, respectively.

[11] This calculation assumes that the mean discount rate inferred from the mean price–dividend ratio is perfectly measured. Standard errors calculated including the estimation of mean discount rates are not much different.

[12] Nominal dividend growth accounts for less than 10 percent of the variance of the VW price–dividend ratio at "var(g) = 0," which is positive but more than two standard errors from 100 percent. It accounts for 50 percent of the variance of the EW price–dividend ratio, which is still less than two standard errors from 0 and 100 percent. Thus, the results are the same with real and nominal dividend growth.

Discount rate equals constant plus risk premium. With this model, the mean price–dividend ratio equation (10) only provides a lower bound $E(g) \geq E(g)_{min}$, rather than identifying the mean discount rate. Lower mean discount rates can increase the fraction of the variance of price–dividend ratios attributed to dividend growth. However, the bounds of $E(g)_{min}$ are not enough lower than the "var(g) = 0" points (see Figure 3) to substantially change the estimates and standard errors. Furthermore, the lower bound $E(g) > E(r)$ is higher than $E(g)_{min}$, and dividend growth accounts for an even smaller fraction of the variance of the price–dividend ratio at mean discount rates $E(g) > E(r)$.

Thus, the constant-risk-premium model does little better than the constant-discount-rate model: some time variation in discount rates is required to explain the variance of price–dividend ratios.

Discount rate equals interest rate plus risk premium. With these discount-rate models, the covariances of the price–dividend ratio with future discount rates are equal to the covariances of the price–dividend ratio with future interest rates. However, these models do not identify the mean discount rate. Thus, if these models work, the "t-bill" or "corp. bond" lines should intersect the "100% − dividends" line in Figure 3, at some mean discount rate above the lower bounds $E(g)_{min}$ and $E(r)$.

All the interest rates account for a positive fraction of the variance for the EW portfolio; the corporate bond rate does so for the VW portfolio (Figure 3). However, the fraction of the variance of the price–dividend ratio accounted for by these interest rates is well below the required amount (the interest rate lines are well below the "100% − dividends" lines) at mean discount rates above the discount rate bounds. A rise in the price–dividend ratio forecasts a decline in these interest rates, but the decline is too small. The forecasts are also statistically insignificant: the interest rate lines are less than one standard error above 0 in Figure 3.[13]

Discount rate equals return. This model works beautifully. At the mean discount rate equal to mean return, marked "$E(r)$" in Figure 3, dividend growth and discount rates together account for almost exactly 100 percent of the variance of the price–dividend ratio (the "100% − dividends" line intersects the "stock return" line).

---

[13] The contribution of nominal interest rates is the real interest rate line *less* the inflation line in Figure 3. Since all the interest rate lines lie below the inflation lines, all nominal interest rates contribute *negative* fractions of the variance of price–dividend ratios. A rise in the price–dividend ratio forecasts a rise in inflation that is larger than the decline in real interest rates, so it forecasts a *rise* in nominal interest rates.

Table 3
Variance decomposition with discount rate = return

Portfolio		Real divi-dends	Real return	Nominal divi-dends	Nominal return	Total
		Percent of variance explained by				
Value-weighted NYSE	% var (P/D)	−34.62	137.79	8.38	94.79	103.17
	Standard error	10.00	31.72	18.91	19.77	
Equally weighted NYSE	% var (P/D)	30.41	84.93	59.74	55.60	115.33
	Standard error	41.34	18.89	41.04	16.85	

The table entries are terms of (14),

$$100\frac{1}{1-\Omega}\left[\sum_{j=1}^{15}\Omega^{j}\text{cov}\left(\frac{P_{t}}{D_{t}}, x_{t+j}\right)\right]\left[\text{var}\left(\frac{P}{D}\right)\right]^{-1},$$

where x is the dividend growth (n) and x is minus the return (r), respectively, and $\Omega = e^{E(n)-E(r)}$. All standard errors contain a Hansen (1982)/Newey–West (1987) correction, as detailed in Appendix G. Data are annual, 1926–1988.

The variance decomposition with discount rate = return is also presented in Table 3. It shows that, in fact, slightly *more* than 100 percent of the variance of the price–dividend ratios is accounted for, verifying the approximation and rejecting bubbles. However, the variance of price–dividend ratios is largely due to changing forecasts of returns rather than to changing forecasts of dividend growth. Since the mean discount rate must be at least as large as the mean return, the bulk of the variance of price–dividend ratios will have to be explained by discount-rate movement with any discount-rate model.

The stylized fact behind this result is that price–dividend ratio forecasts of real and nominal returns are large and negative, confirming Fama and French (1988). Even when price–dividend ratios forecast that interest rates and dividend growth rates move in the right direction, the magnitudes of the forecasts are much smaller than the magnitude of the forecast change in returns.

Consumption-based discount rates. The variance decomposition using consumption-based discount rates is presented in Table 4.

As before, the mean price–dividend ratio equation (10) can be used to infer one parameter of the utility function, so I estimated the risk-aversion coefficient α by imposing the mean price–dividend ratio equation (10) for various assumed values of the subjective discount factor ρ. All the consumption-based discount rate contributions to the variance of price–dividend ratios are negative, and many are more than two standard deviations below 0.

A high price–dividend ratio forecasts *higher* long-run consumption growth. It should forecast *lower* consumption growth: any wealth effects of a stock price rise should be incorporated into consumption

Table 4
Variance decomposition using consumption-based discount rates

Utility parameters		Mean discount rate	Variance decomposition				
			% var(P/D) due to			SE	
ρ	α	$E(g)$ (%)	Dividend growth	Discount rates	Total	Dividend growth	Discount rates
1. Value-weighted NYSE							
.99	4.28	6.50	−31.34	−26.08	−57.42	9.17	26.02
.97	2.32	6.03	−35.56	−15.75	−51.31	10.23	15.97
.95	0.60	5.90	−36.92	−4.19	−41.11	10.58	4.27
2. Equally weighted NYSE							
.99	8.71	12.21	14.62	−12.63	1.99	22.08	22.67
.97	6.43	11.32	17.70	−9.87	7.83	25.92	19.92
.95	4.20	10.54	21.30	−6.78	14.52	30.34	15.44

Discount rates are generated by the consumption-based model:

$$g_t = -\ln(\rho) + \alpha \ln(c_t/c_{t-1}).$$

For a given choice of ρ, α is calculated to satisfy the mean price–dividend ratio equation

$$E\left(\frac{P}{D}\right) = \frac{\Omega}{1-\Omega} + \frac{\Omega}{2(1-\Omega)^2} \sum_{j=-15}^{15} \Omega^{|j|} \mathrm{cov}(n_t - g_t, n_{t-j} - g_{t-j}), \tag{10}$$

where n is the log of real dividend growth and $\Omega = e^{E(n-g)}$. The variance decomposition gives the percent of the variance of the price–dividend ratio due to dividends and discount rates, from (14):

$$100 \frac{1}{1-\Omega} \left[\sum_{j=1}^{15} \Omega^j \mathrm{cov}\left(\frac{P_t}{D_t}, x_{t+j}\right) \right] \left[\mathrm{var}\left(\frac{P}{D}\right) \right]^{-1},$$

where x is the dividend growth n and discount rates g as indicated. All standard errors contain a Hansen (1982)/Newey–West (1987) correction, as detailed in Appendix G. Data are annual, 1926–1988.

immediately, and then consumption *growth* should be lower, as discount rates should be lower.

6. Bounds on the Mean and Standard Deviation of Discount Rates

So far, we have found that there *exist* unobserved discount-rate processes that explain the variance of price–dividend ratios, but the observable discount-rate proxies (other than the trivial case of the return itself) do not satisfy the variance decomposition. Are the unobserved discount rates that one must invoke to satisfy the mean price–dividend ratio, variance bound, and variance decomposition "reasonable," or must they have unusual time-series processes suggestive of "fads"? To address this equation, this section computes the minimum standard deviation of discount rates required to satisfy the mean price–dividend ratio, the variance bound, and the variance decomposition.[14]

---

[14] The choice of mean and variance rather than other moments more appropriate to nonnormal variables is arbitrary. It reflects a tradition in the discount-rate-puzzle literature rather than any deep observations about the importance of variances over other moments.

Value weighted NYSE

Equally weighted NYSE
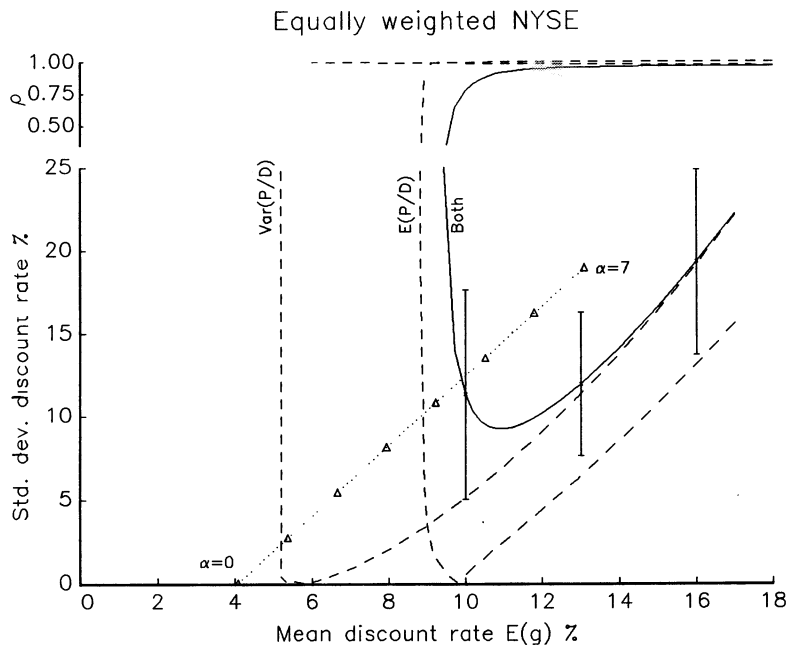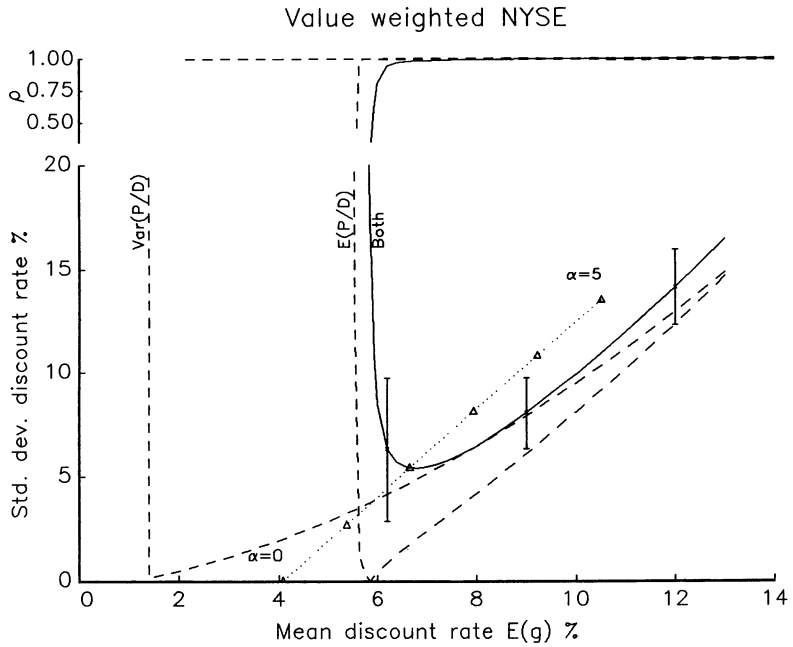
The problem is to find the discount-rate process with minimum variance for a given mean discount rate, subject to the constraints that the mean price–dividend ratio (10), the variance bound (12) or (13), and/or variance decomposition (14) are satisfied. The solution to this problem is a discount-rate process that is a singular function of dividend growth and the price–dividend ratio:

$$g_t^* = \alpha(L)n_t + \beta(L)(P_t/D_t).$$ (17)

[$\alpha(L)$ and $\beta(L)$ may be two-sided.] Adding a noise term to (17) just adds variance without helping to satisfy the constraints. The problem is then to find the forms of $\alpha(L)$ and $\beta(L)$ that minimize the variance of g_t subject to the constraints. The solution is straightforward but algebraically unpleasant, and so is presented in Appendix E. The results are presented in Figure 4.

First, examine the bound on the standard deviation of discount rates that satisfy the mean price–dividend ratio (10), marked "$E(P/D)$" in Figure 4. This bound has a global minimum of 0 at the mean discount rate that solves (10) [marked "var$(g) = 0$" in previous graphs]. As explained in conjunction with the variance bound in Figure 2, the discount rate must have higher variance for both larger and smaller mean discount rates to keep the mean price–dividend ratio prediction (10) satisfied. This bound rises to infinity at the mean discount rate $E(g)_{min}$, since no lower discount rate is consistent with the mean price–dividend ratio equation.

Next, examine the bound on the standard deviation of discount rates that satisfy the variance decomposition (14). This bound basically rises with the mean discount rate, as the fraction of the variance of the price–dividend ratio attributed to dividend growth declines (see Figure 3). For the EW portfolio, there is a global minimum standard deviation of 0, at the mean discount rate where real dividend growth accounts for 100 percent of the variance of the price–dividend ratio.

←

Figure 4
Bounds on the standard deviation of discount rates that satisfy the mean price–dividend ratio and variance decomposition, autocorrelation of discount rates that attain the bounds, and mean and standard deviation of consumption-based discount rates
The bounds marked "$E(P/D)$," "Var(P/D)," and "Both" give the minimum standard deviation of discount rates that satisfy the mean price–dividend ratio equation (10), the variance decomposition (14), and (10) and (14) together, respectively, using real dividend growth. The error bars report selected standard errors of the "Both" bound. The top part of each graph reports the autocorrelation of the discount rate processes that attain each bound. The line marked "$\alpha = 0$," ..., "$\alpha = 5$" reports the mean and standard deviation of consumption-based discount rates,

$$g_t = -\ln(\rho) + \alpha \ln(c_t/c_{t-1}),$$

using subjective discount factor $\rho = .96$ and the indicated risk aversion α. Data are annual, 1926–1988.

The variance bound (12) is an inequality constraint and so only adds the information that the mean discount rate must be greater than the value at which the variance bound is just satisfied (see Figure 2). This value is very close to $E(g)_{min}$ at which the mean price–dividend bound ["$E(P/D)$" in Figure 4] diverges to infinity, and so is not separately reported in Figure 4.

The minimum values of the bounds on the standard deviation of discount rates implied by the mean price–dividend ratio, variance bound, and variance decomposition *separately* (as they are usually analyzed) are all equal to—or very near—0. However, the minimum value of the bound ("Both") that imposes all these conditions simultaneously is positive. For the value-weighted portfolio, a standard deviation of roughly 5 percent is required at a mean discount rate of about 7 percent, while for the equally weighted portfolio, a standard deviation of roughly 10 percent is required at a mean discount rate of 11 percent. The minimum standard deviations are about two standard errors above zero.

Are these discount rates "reasonable"? To get some idea, the mean and standard deviation of consumption-based discount rates are also presented in Figure 4. The utility parameters are a subjective discount factor $\rho = .96$ and a variety of risk-aversion coefficients below 7. Changing the subjective discount factor ρ simply shifts the curve to the left or right. As the figure shows, the consumption-based discount rates have means and standard deviations in the required regions with a range of "reasonable" parameters: risk-aversion coefficients of 5 are sufficient. (The comparison with consumption-based discount rates is *not* a test of the consumption-based model, as that model fails to generate moments other than mean and variance correctly. Figure 3 and Table 4 showed that forecasts of consumption growth from price–dividend ratios had the wrong sign.)

By contrast, the simplest similar bounds constructed by Hansen and Jagannathan (1991) for discount factors that generate the unconditional equity premium require risk-aversion coefficients over 40, and their bounds on discount factors that generate the unconditional term premium require risk-aversion coefficients in the hundreds. Thus, the minimum standard deviation of discount rates required to explain the variance of price–dividend ratios is an order of magnitude lower than that required to explain unconditional return premiums.

The required discount rates have some other unusual characteristics. Figure 4 shows the first-order autocorrelation coefficients of the discount-rate processes on the standard-deviation bound. For most mean discount rates, this autocorrelation is very near 1. The variance-minimizing discount rates are also highly predictable from past price–dividend ratios and discount rates. Consumption growth is nearly

unpredictable both from its own past and past price–dividend ratios or dividend growth. (However, the mean, standard deviation, auto-correlation, and predictability are consistent with marginal rates of *transformation* implied by standard intertemporal production functions [see Cochrane (1991a)].)

In summary, the unobserved discount rates we must invoke to explain the variance of price–dividend ratios are "reasonable" from the usual criterion of their standard deviation, but are quite unlike consumption-based discount rates in their autocorrelation and predictability.

7. Concluding Remarks

In this article, I examine a variance bound and a variance decomposition for price–dividend ratios. I try to answer three questions with each test: (1) Is there *any* discount rate process that is consistent with the variance of price–dividend ratios, or is its variance an indication of "bubbles"? (2) Are there *reasonable* unobserved discount rate processes that explain the variance of price–dividend ratios? (3) Do particular discount rate *models* account for the variance of price–dividend ratios? The tests are derived with an approximate present-value model that extends the specification of similar models in the literature. In particular, it captures the importance of the mean discount rate to the volatility of stock prices, it allows us to use information about the mean as well as the variance of the price–dividend ratio, and it can generate cross-sectional risk premiums because of varying low-frequency correlations of dividend growth with the discount rate.

The variance bound is easily satisfied for all discount-rate models. The central ingredient in this result is the choice of units: price–dividend ratios and dividend growth rates. Most variance bounds in the literature that use these units do not reject, even with more restrictive specifications. The extensions pursued here—serial correlation in dividend growth, time-varying discount rates, varying mean discount rates, and checking the implications of the mean price–dividend ratio—could have big effects, but turn out not to.

The variance decomposition yields mixed results. The variance of the price–dividend ratio is entirely accounted for by forecasts of dividend growth together with *returns,* so there is no indication of bubbles or serious shortcomings of the modeling approximations. However, changing return forecasts account for the bulk of the variance of price–dividend ratios, so most of that variance must be chalked up to discount-rate movements. But the discount-rate models—interest rate plus risk premium, and consumption based—fail.

269

Two observations are crucial to this result. First, a rise in price–dividend ratios forecasts a large decline in returns, but its forecasts of dividend growth, interest rates, and consumption growth rates are typically smaller, less statistically significant, and not always of the "right" sign.[15] This fact by itself need not be a problem for the variance decomposition, as we might just invoke a low mean discount rate. However, the second observation is that the *mean* price–dividend ratio equation must also be satisfied. This observation rules out the mean discount rates for which dividend growth, alone or together with interest rate proxies for discount rates, can account for the variance of the price–dividend ratio.

The consumption-based model also fails miserably. A rise in the price–dividend ratio forecasts a long-term *rise* in consumption growth, rather than a decline as it should. This is a different problem than the low contemporaneous correlation of returns and consumption growth documented by Euler equation rejections.

Mean–standard deviation frontiers for unobserved discount rates that satisfy the variance bound, variance decomposition, and mean price–dividend ratio are calculated, to see if these discount rates are "reasonable" or suggestive of "fads." Here again, it is important to use the information in the *mean* price–dividend ratio as well as the variance bound or decomposition. Constant discount rates could satisfy the mean, variance bound, and variance decomposition *separately,* but not *together.* The discount rates required to explain the mean price–dividend ratio and the bounds have standard deviations that seem "reasonable," but they are more predictable and autocorrelated than consumption growth rates.

In summary, discount-rate processes *exist* that explain the variance of price–dividend ratios. But none of the discount-rate models considered here (except the trivial model "discount rate = return") are forecastable enough from price–dividend ratios to explain the variance of price–dividend ratios.

In the end, these tests of the present-value relation look much like tests of Euler equations: there are no striking rejections as in the original articles by LeRoy and Porter (1981) and Shiller (1981), and there is no requirement for outrageous discount-rate behavior. However, existing discount-rate models are rejected. The remaining difficulty is the same—namely, to find successful and nontrivial measures of discount rates, such as measures of marginal rates of substitution or transformation.

---

[15] It might be possible to raise the dividend forecast by using earnings measures. If firms smooth earnings to dividends, then a rise in price–dividend ratios may forecast a large change in earnings a few years out, but the corresponding rise in dividends may not occur inside the 15-year limit I used for dividend growth forecasts.

Appendix A: List of Symbols

The symbols used in this article are

P_t = stock price (claim to dividends from $t + 1$ forward);
D_t = dividends in period t;
η_t = dividend growth, $\quad \eta_t = D_t/D_{t-1}$;
n_t = log dividend growth, $\quad n_t = \ln(\eta_t)$;
R_t = gross rate of return, $\quad R_t = (P_t + D_t)/P_{t-1}$;
r_t = $\ln(R_t)$;
γ_t = discount factor, e.g., $1 = E_{t-1}(\gamma_t R_t)$;
g_t = discount rate, $\quad g_t = -\ln(\gamma_t)$;
ρ = subjective discount factor in utility;
α = coefficient of risk aversion, $\quad u'(c) = c^{-\alpha}$;
Ω = $e^{E(n)-E(g)}$.

$E(g)_{\min}$ denotes the minimum value of $E(g)$ consistent with the mean price–dividend ratio equation (10):

$$E(g)_{\min} \equiv E(n) - \ln[E(P/D)/(1 + E(P/D))].$$

$\text{var}(g) = 0$ denotes the $E(g)$ at which the mean price–dividend ratio equation (10) holds with a constant discount rate.

Appendix B: Stationarity of the Price–Dividend Ratio in the Exact Present-Value Model

Equation (8) is *consistent* with the assumption that price–dividend ratios, dividend growth, and discount rates are all stationary. It would be nice to show that stationary dividend growth and discount rates *imply* a stationary price–dividend ratio.

A strongly stationary dividend growth and discount rate *imply* a strongly stationary price–dividend ratio, if the expected value of the sum converges almost surely. $E(g) > E(n)$ is necessary and sufficient for the *sum* to converge a.s. Precisely, if γ_t and η_t are strongly stationary and positive a.s., then one can show that the weak law of large numbers implies

$$\lim_{T\to\infty} \sum_{j=1}^{T} \prod_{k=1}^{j} (\gamma_{t+k}\eta_{t+k}) < \infty \text{ a.s.}$$

if and only if $E(g) > E(n)$. However, $E(g) > E(n)$ is not sufficient for the *expected value* of the sum to converge a.s. For example, suppose $w_t = n_t - g_t$ is lognormally distributed. Then, $E(\exp(w)) = \exp(E(w) + \frac{1}{2}\text{var}(w))$, so

$$E\left(\frac{P_t}{D_t}\right) = E\left(\sum_{j=1}^{\infty} \exp \sum_{k=1}^{j} w_{t+k}\right) = \sum_{j=1}^{\infty} \exp\left(jE(w) + \frac{1}{2}\text{var}\left(\sum_{k=1}^{j} w_{t+k}\right)\right).$$

Now,

$$\lim_{j \to \infty} \frac{1}{j} \operatorname{var}\left(\sum_{k=1}^{j} w_{t+k}\right) = s_w(0),$$

where $S_w(0)$ denotes the spectral density of w at frequency 0. Thus, the expected price–dividend ratio (or the price–dividend ratio itself if the moments above are conditional on time t information) is finite if and only if $E(w) + S_w(0) < 0$, which is more stringent than $E(w) < 0$. Furthermore, establishing strong stationarity does not establish the existence of second moments for weak stationarity, which is the basis of the usual linear time-series methods.

Lacking both necessary and sufficient auxiliary conditions [such as $E(g) > E(n)$] for dividend growth and discount-rate stationarity to *imply* a stationary price–dividend ratio, the logic of the text is to *assume* that the price–dividend ratio, dividend growth, and discount rates are all stationary with first and second moments, and to point out that this is *consistent* with the present-value formula [Equations (7) and (8)].

Appendix C: Derivation of the Approximate Present-Value Model

Start from the exact present-value model (8). Multiply both sides by any variable Z_t observed at time t and take expectations, yielding

$$E\left(Z_t \frac{P_t}{D_t}\right) = E\left[Z_t \sum_{j=1}^{\infty} \left(\exp \sum_{k=1}^{j} w_{t+k}\right)\right],$$

where $w_t = n_t - g_t$. Construct a second-order Taylor expansion of the expression in the brackets, with respect to w_{t+j}, $j = 1, 2, \ldots$, and Z_t about their means $E(w)$ and $E(Z)$. With $\tilde{w}_t = w_t - E(w)$, that Taylor expansion is

$$
Z_t \sum_{j=1}^{\infty} \left[\exp\left(\sum_{k=1}^{j} w_{t+k}\right)\right]
$$

$$
\cong \frac{Z_t \Omega}{1 - \Omega} + \frac{Z_t}{1 - \Omega} \sum_{j=1}^{\infty} \left(\Omega^j \tilde{w}_{t+j}\right)
$$

$$
+ \frac{1}{2} \frac{E(Z)}{1 - \Omega} \sum_{j=1}^{\infty} \left[\Omega^j \left(\tilde{w}_{t+j}^2 + 2 \sum_{k=1}^{\infty} \Omega^k \tilde{w}_{t+j} \tilde{w}_{t+j+k}\right)\right],
$$

where $\Omega = e^{E(w)} = e^{E(n) - E(g)}$. Taking expectations,

$$
E\left(Z_t \frac{P_t}{D_t}\right) = E(Z_t)\left(\frac{\Omega}{1 - \Omega} + \frac{\Omega}{2(1 - \Omega)^2} \sum_{j=-\infty}^{\infty} \Omega^{|j|} \operatorname{cov}(w_t, w_{t-j})\right)
$$

$$
+ \frac{1}{1 - \Omega} E\left(Z_t \sum_{j=1}^{\infty} \Omega^j \tilde{w}_{t+j}\right).
$$

Since this equation holds for all variables Z_t known at time t, it is equivalent to

$$\frac{P_t}{D_t} = \frac{\Omega}{1 - \Omega} + \frac{\Omega}{2(1 - \Omega)^2} \sum_{j=-\infty}^{\infty} [\Omega^{|j|} \operatorname{cov}(w_t, w_{t-j})]$$

$$+ \frac{1}{1 - \Omega} E_t \left(\sum_{j=1}^{\infty} \Omega^j \tilde{w}_{t+j} \right).$$

Substituting $w_t = n_t - g_t$ yields Equation (9) in the text.

Appendix D: Accuracy of Approximation

The mean price–dividend ratio, variance bound, and total of the variance decomposition with discount rate equal to return, presented in the text, provide measures of the accuracy of approximation. As another measure, this appendix examines the accuracy of the Taylor approximation as applied to the ex post price–dividend ratio with discount rates equal to returns, that is, the quality of the approximation:

$$\frac{P_t}{D_t} \cong E\left(\frac{P}{D}\right) + \frac{1}{1 - \Omega} \sum_{j=1}^{\infty} \Omega^j (\tilde{n}_{t+j} - \tilde{r}_{t+j}), \qquad (D1)$$

where

$$E\left(\frac{P}{D}\right) = \frac{\Omega}{1 - \Omega} + \frac{\Omega}{2(1 - \Omega)^2} \sum_{j=-\infty}^{\infty} \Omega^{|j|} \operatorname{cov}(n_t - g_t, n_{t-j} - g_{t-j}).$$

The model (7) and (8) holds exactly, ex post (i.e., dropping the E_t) with returns in the place of discount rates. Thus, the only reason the calculation of (D1) differs from the exact price–dividend ratio is the accuracy of the approximation, together with the fact that the sums are only taken to the terminal date (1988) rather than to infinity. A calculation of (D1), together with the actual price–dividend ratio, is presented in Figure 5.

For comparison, the Campbell–Shiller (1988) approximation to the price–dividend ratio, again using ex post returns to discount dividends, is also presented in Figure 5. Campbell and Shiller start with the identity
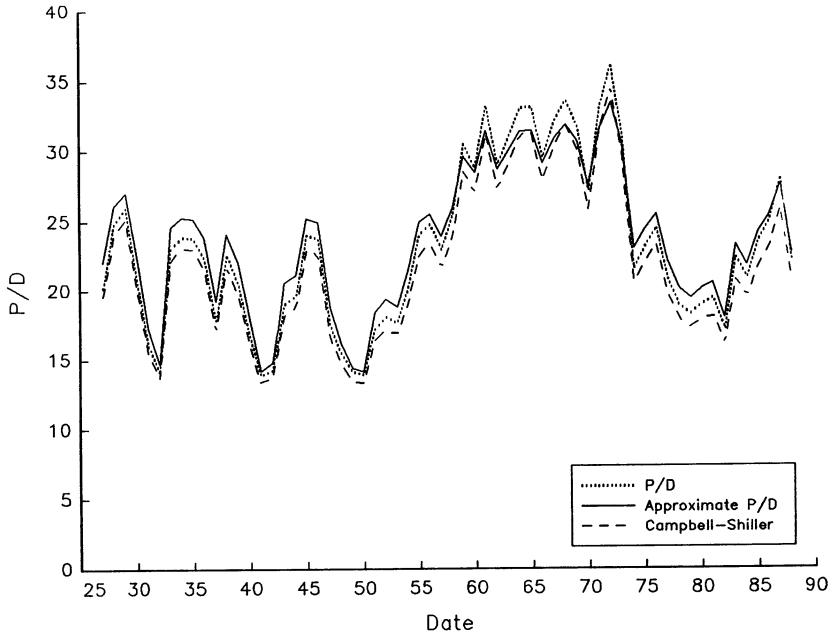
$$1 = R_{t+1}^{-1} R_{t+1}.$$

This identity implies

$$\frac{P_t}{D_t} = R_{t+1}^{-1} \frac{D_{t+1}}{D_t} \left(\frac{P_{t+1}}{D_{t+1}} + 1 \right).$$

Take logs of the last equation, Taylor approximate the log of the term

Figure 5
Accuracy of approximation
The line marked "P/D" gives the price–dividend ratio of the value-weighted NYSE. The line marked "Approximate P/D" gives the approximate ex post present value:

$$\frac{P_t}{D_t}(\text{approximate}) = E\!\left(\frac{P}{D}\right) + \frac{1}{1-\Omega} \sum_{j=1}^{T-t} \Omega^j (\tilde{n}_{t+j} - \tilde{r}_{t+j}),$$

$$E\!\left(\frac{P}{D}\right) = \frac{\Omega}{1-\Omega} + \frac{\Omega}{2(1-\Omega)^2} \sum_{j=15}^{15} \Omega^{|j|} \operatorname{cov}(n_t - r_t, n_{t-j} - r_{t-j}),$$

where $T = 1988$. (The graph is essentially the same if one uses the sample mean price–dividend ratio in place of the predicted value from the covariances of dividend growth less return.) The line marked "Campbell–Shiller" gives the equivalent Campbell–Shiller approximation:

$$\frac{P_t}{D_t}(\text{Campbell–Shiller}) = \exp\!\left(-\left[\sum_{j=1}^{T-t} \Omega^{j-1}(n_{t+j} - r_{t+j}) + \frac{\Omega^{T-t}}{1-\Omega}\!\left(E(n) - E(r)\right) - \frac{k}{1-\Omega}\right]\right),$$

$$k = (1-\Omega)\ln(1-\Omega) + \Omega \ln(\Omega).$$

in parentheses, and iterate forward on the log price–dividend ratio to obtain

$$\ln\!\left(\frac{D_t}{P_t}\right) \cong \sum_{j=1}^{\infty} \rho^{j-1}(r_{t+j} - \Delta d_{t+j}) - \frac{k}{1-\rho},$$

where

$$k = -[(1-\rho)\ln(1-\rho) + \rho \ln(\rho)]$$

and ρ is an arbitrary constant determined by the point at which the Taylor approximation is taken. In arriving at Figure 5, I used

$$\rho = \Omega = e^{E(r) - E(n)},$$

following Campbell and Shiller. As Figure 5 shows, the approximation used in this article and the Campbell–Shiller approximation perform about equally in this test. Of course, the real test is how well the two models perform with the "true" discount-rate model rather than ex post returns. Also, my approximation was designed to endogenously generate risk premiums, which have to be added to the Campbell–Shiller approximation through an expected-return model.

Appendix E: Finding the Mean–Standard Deviation Frontier for Discount Rates

The problem is to minimize var(g_t), subject to the mean (10) and variance decomposition (14). Using $Z_t = P_t/D_t$ in the right-hand side of (14), the constraints (10) and (14) can be rewritten as

$$\left(E\left(\frac{P}{D}\right) - \frac{\Omega}{1 - \Omega} \right) \frac{2\Omega(1 - \Omega)^2}{1 - \Omega^2} = X_1 = \text{var}\left(\sum_{j=1}^{\infty} \Omega^j (n_{t+j} - g_{t+j}) \right),$$

$$\text{cov}\left(Z_t, \sum_{j=1}^{\infty} \Omega^j n_{t+j} \right) - (1 - \Omega) \, \text{cov}\left(Z, \frac{P}{D} \right) = X_2 = \text{cov}\left(Z_t, \sum_{j=1}^{\infty} \Omega^j g_{t+j} \right).$$

The problem is most easily solved in the frequency domain,

$$\min \frac{1}{\pi} \int_0^{\pi} S_g(\omega) \, d\omega$$

subject to

$$X_1 = \frac{1}{\pi} \int_0^{\pi} |b(e^{-i\omega})|^2 S_{n-g}(\omega) \, d\omega$$

and

$$X_2 = \frac{1}{2\pi} \int_0^{\pi} S_{z,b(L)g}(\omega) \, d\omega,$$

where

$$b(L) = \sum_{j=1}^{\infty} \Omega^j L^{-j}$$

so

$$b(e^{-i\omega}) = \frac{\Omega e^{i\omega}}{1 - \Omega e^{i\omega}}.$$

$S_x(\omega)$, and $S_{x,y}(\omega)$ denote spectral and cross-spectral densities, respectively.

The variance-minimizing process is singular, that is, there are an α and β such that $g_t = \alpha(L)n_t + \beta(L)z_t$. With this form for the g_t process, and using $S_{x,b(L)y} = b(e^{i\omega})S_{x,y}$, the problem becomes

$$\min_{\{\alpha,\beta\}} \frac{1}{\pi} \int_0^\pi d\omega \, |\alpha|^2 S_n + |\beta|^2 S_z + \alpha\beta^* S_{nz} + \alpha^*\beta S_{nz}^* \qquad \text{(E1)}$$

subject to

$$X_1 = \frac{1}{\pi} \int_0^\pi d\omega \, |b|^2 (|1 - \alpha|^2 S_n + |\beta|^2 S_z - (1 - \alpha)\beta^* S_{nz}$$
$$- (1 - \alpha)^* \beta S_{nz}^*),$$

$$X_2 = \frac{1}{2\pi} \int_0^\pi d\omega \, (b\alpha S_{nz} + b^*\alpha^* S_{nz}^*) + (b\beta + b^*\beta^*) S_z,$$

where the $(e^{-i\omega})$ notation is suppressed following α, β, b, and the spectral densities.

This is a frequency-by-frequency Langrangian maximization. The first-order conditions yield

$$\alpha = \frac{\delta_1 |b|^2}{1 + \delta_1 |b|^2}, \qquad \beta = \frac{\delta_2 b^*}{2(1 + \delta_1 |b|^2)},$$

where δ_1 and δ_2 are Lagrange multipliers on the two constraints. Substituting into (E1), the minimum variance of discount rates is

$$\text{var}(g_t) = \frac{1}{\pi} \int_0^\pi d\omega \, \frac{|b|^2}{(1 + \delta_1 |b|^2)^2} \left(\delta_1^2 |b|^2 S_n + \frac{\delta_2^2}{4} S_z \right.$$
$$\left. - \frac{\delta_1 \delta_2}{2} (b S_{nz} + b^* S_{nz}^*) \right), \qquad \text{(E2)}$$

where the constraints δ_1 and δ_2 are found from the constraint equations

$$X_1 = \frac{1}{\pi} \int_0^\pi d\omega \, \frac{|b|^2}{(1 + \delta_1 |b|^2)^2} \left(S_n + \frac{\delta_2^2}{4} |b|^2 S_z + \frac{\delta_2}{2} (b S_{nz} + b^* S_{nz}^*) \right), \qquad \text{(E3)}$$

$$X_2 = \frac{1}{2\pi} \int_0^\pi d\omega \, \frac{|b|^2}{1 + \delta_1 |b|^2} \left(\delta_1 (b S_{nz} + b^* S_{nz}^*) - \delta_2 S_z \right). \qquad \text{(E4)}$$

The second constraint can be solved for δ_2 as a function of δ_1:

$$\delta_2 = \left(2\pi X_2 - \int_0^\pi d\omega \, \frac{|b|^2}{1 + \delta_1 |b|^2} \delta_1 (b S_{nz} + b^* S_{nz}^*) \right)$$
$$\times \left(\int_0^\pi d\omega \, \frac{|b|^2}{1 + \delta_1 |b|^2} S_z \right)^{-1}. \qquad \text{(E5)}$$

To calculate the minimum standard deviation of discount rates for a given mean discount rate, I followed this procedure: (1) I constructed the spectral densities S_n, S_z, S_{nz} using the first 15 covariances, weighted as in Newey and West (1987) to ensure that the spectral density matrix is positive definite; (2) I used (E5) to substitute for δ_2 in (E3), and then searched for a value of δ_1 that satisfies (E3), performing the integrals numerically; (3) with the resulting values of δ_1 and δ_2, I evaluated (E2) to give the minimum variance of discount rates.

The case $\delta_1 = 0$, which corresponds to the same minimization imposing only the variance decomposition, yields a natural interpretation as a regression of discount rates on a variable observed at time t. In this case, the minimum variance reduces to

$$\text{var}(g_t) = \frac{1}{\pi} \int_0^\pi d\omega \, \frac{\delta_2^2}{4} \, |h|^2 S_z,$$

with

$$X_2 = -\frac{1}{2\pi} \int_0^\pi d\omega \, \delta_2 |h|^2 S_z.$$

Solving the constraint for δ_2 and substituting in the variance, we obtain

$\text{var}(g_t)$

$$= \frac{1}{\pi} \int_0^\pi d\omega \, |h|^2 S_z \left(-X_2 \left[\frac{1}{\pi} \int_0^\pi d\omega \, |h|^2 S_z \right]^{-1} \right)^2$$

$$= \left(\text{cov}\left(Z_t, \sum_{j=1}^\infty \Omega^j n_{t+j} \right) - (1 - \Omega) \, \text{cov}\left(Z, \frac{P}{D} \right) \right)^2 \left[\text{var}\left(\sum_{j=1}^\infty \Omega^j Z_{t-j} \right) \right]^{-1}$$

$$= \left(\text{cov}\left(g_t, \sum_{j=1}^\infty \Omega^j Z_{t-j} \right) \left[\text{var}\left(\sum_{j=1}^\infty \Omega^j Z_{t-j} \right) \right]^{-1} \right)^2 \text{var}\left(\sum_{j=1}^\infty \Omega^j Z_{t-j} \right).$$

This is the variance of the fitted value of the regression of g_t on $\sum_{j=1}^\infty \Omega^j Z_{t-j}$. Thus, the minimum variance of discount rates compatible with the variance decomposition only can also be found by inferring the OLS regression coefficient of discount rates g_t on the variable $\sum_{j=1}^\infty \Omega^j Z_{t-j}$. Also, this calculation shows that the variance-minimizing discount rates are thus one-step-ahead predictable from past Z_t.

Appendix F: Data Description

The Treasury bill, government bond index, corporate bond index, and CPI are from the Ibbotson–Sinquefield database. The real per capita consumption series is the same as in Campbell and Shiller

(1988), extended to 1988 with NIPA real per capita nondurable and services consumption growth.

Price–dividend ratio and dividend growth data are based on the CRSP value-weighted and equally weighted NYSE portfolio returns, with and without dividends, converted to annual frequencies to avoid the seasonal in dividends. (The model can apply to monthly data, but the approximation would be worse with big seasonals in dividend growth.) Precisely, the dividend–price ratio and dividend growth are formed from the CRSP annualized total (R_t^T) and price-only (R_t^C) annualized gross returns as follows:

$$\frac{D_t}{P_t} = \frac{(R_t^T - R_t^C)}{R_t^C} = \left(\frac{P_t + D_t}{P_{t-1}} - \frac{P_t}{P_{t-1}}\right)\frac{P_{t-1}}{P_t},$$

$$\frac{D_t}{D_{t-1}} = \frac{D_t}{P_t}\frac{P_{t-1}}{D_{t-1}}\frac{P_t}{P_{t-1}},$$

(F1)

where P_t is the December 31, year $t - 1$ closing price. One can show that this construction means that annual dividends in year t are the monthly dividends, brought forward from the end of the month in which they are paid to December 31 at the total monthly return R^T.

Appendix G: Standard Errors

Organize the data into a vector y_t, whose mean gives the means, variances, and covariances of price–dividend ratios, dividend growth, and discount rates:

$$y_t = [n_t \quad g_t \quad P_t/D_t \quad \tilde{n}_t^2 \quad \tilde{g}_t^2 \quad (P_t/D_t)^2 \quad \tilde{n}_t\tilde{n}_{t-1} \quad \tilde{n}_t\tilde{n}_{t-2} \quad \cdots],$$

$$E(y_t) = [E(n_t) \quad E(g_t) \quad E(P_t/D_t) \quad \text{var}(n_t) \quad \text{var}(g_t) \quad \text{var}(P_t/D_t) \quad \text{cov}(n_t, n_{t-1}) \quad \cdots]$$

$$= \mu.$$

All the statistics of the article can be expressed as differentiable functions $f(\mu)$, which are estimated by using sample moments μ_T in place of the population moment μ. Assuming the variables in y_t are stationary (and other regularity conditions), Hansen (1982, 1985) shows that

$$T^{1/2}[f(\mu_T) - f(\mu)]$$

converges in distribution to a normally distributed random vector with mean 0 and covariance matrix $\nabla f(\mu)' V_0 \nabla f(\mu)$, where

$$V_0 = \lim_{J \to \infty} \sum_{j=-J}^{J} \left(1 - \frac{|j|}{J}\right) E\left((y_t - \mu)(y_{t-j} - \mu)'\right).$$

I calculated ∇f analytically where possible, and otherwise as the

numerical derivative of the procedures that calculate test statistics as a function of sample moments. $J = 5$ is used throughout the tables. (I did some experiments and found little difference in standard errors with more than $J = 5$ covariances.)

References

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

Cochrane, J. H., 1991a, "Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations," *Journal of Finance,* 46, 207–234.

Cochrane, J. H., 1991b, "Volatility Tests and Market Efficiency: A Review Essay," *Journal of Monetary Economics,* 27, 463–485.

Diba, B. T., and H. I. Grossman, 1988, "Rational Inflationary Bubbles," *Journal of Monetary Economics,* 21, 35–46.

Fama, E. F., and K. R. French, 1988, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics,* 22, 3–25.

Flavin, M., 1983, "Excess Volatility in the Financial Markets: A Reassessment of the Empirical Evidence," *Journal of Political Economy,* 96, 246–273.

Flood, R. P., and R. J. Hodrick, 1990, "On Testing for Speculative Indeterminacies in Models with Rational Implications," *Journal of Economic Perspectives,* 4, 85–101.

Gilles, C., and S. F. LeRoy, 1991, "Econometric Aspects of the Variance-Bounds Tests: A Survey," *Review of Financial Studies,* 4, 753–791.

Grossman, S. J., and R. J. Shiller, 1981, "The Determinants of the Variability of Stock Market Prices," *American Economic Review,* 71, 222–227.

Hamilton, J. D., and C. H. Whiteman, 1985, "The Observable Implications of Self-Fulfilling Expectations," *Journal of Monetary Economics,* 16, 353–373.

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

Hansen, L. P., 1985, "Using Martingale Difference Approximations to Obtain Covariance Matrix Bounds for Generalized Method of Moments Estimators," NORC Discussion Paper 85-16.

Hansen, L. P., and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy,* 99, 225–262.

Hansen, L. P., and S. F. Richard, 1987, "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica,* 55, 587–613.

Kleidon, A. W., 1986, "Variance Bounds Tests and Stock Price Valuation Models," *Journal of Political Economy,* 94, 953–1001.

LeRoy, S. F., and W. R. Parke, 1990, "Stock Price Volatility: Tests Based on the Geometric Random Walk," working paper, University of California at Santa Barbara.

LeRoy, S. F., and R. D. Porter, 1981, "The Present Value Relation: Tests Based on Implied Variance Bounds," *Econometrica,* 49, 555–574.

Mankiw, N. G., D. Romer, and M. D. Shapiro, 1991, "Stock Market Forecastability and Volatility: A Statistical Reappraisal," *Review of Economic Studies,* 58, 455–477.

Marsh, T. A., and R. C. Merton, 1986, "Dividend Variability and Variance Bounds Tests for the Rationality of Stock Market Prices," *American Economic Review,* 76, 483–498.

Mehra, R., and E. C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics,* 15, 154–161.

Newey, W. K., and K. D. West, 1987, "A Simple, Positive-Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica,* 55, 703–708.

Poterba, J. M., and L. H. Summers, 1988, "Mean Reversion in Stock Prices: Evidence and Implications," *Journal of Financial Economics,* 22, 27–59.

Shiller, R. J., 1981, "Do Stock Prices Move too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review,* 71, 421–426.

Shiller, R. J., 1989, *Market Volatility,* MIT Press, Cambridge.

Watson, M. W., 1991, "Measures of Fit for Calibrated Models," Technical Working Paper 102, National Bureau of Economic Research.

West, K. D., 1987, "A Specification Test for Speculative Bubbles," *Quarterly Journal of Economics,* 102, 553–580.

West, K. D., 1988, "Dividend Innovations and Stock Price Volatility," *Econometrica,* 56, 37–61.