Two Trees

John H. Cochrane[*]
University of Chicago

Francis A. Longstaff[**]
The UCLA Anderson School and NBER

Pedro Santa-Clara[**]
The UCLA Anderson School and NBER

We solve a model with two i.i.d. Lucas trees. Although the corresponding one-tree model produces a constant price-dividend ratio and i.i.d. returns, the two-tree model produces interesting asset-pricing dynamics. Investors want to rebalance their portfolios after any change in value. Because the size of the trees is fixed, prices must adjust to offset this desire. As a result, expected returns, excess returns, and return volatility all vary through time. Returns display serial correlation and are predictable from price-dividend ratios. Return volatility differs from cash-flow volatility, and return shocks can occur without news about cash flows.

Returns that are independent over time are the standard benchmark for theory and empirical work in asset pricing. Yet, on reflection, i.i.d. returns seem impossible with multiple positive net supply assets. If a stock or a sector rises in value, investors will try to rebalance away from it. But we cannot all rebalance, as the average investor must hold the market portfolio. It seems that the successful asset's expected returns must rise, or some other return moment must change, in order to induce investors to hold more of the successful securities.

We characterize the asset price and return dynamics that result from this market-clearing mechanism in a simple context. We solve an asset-pricing model with two Lucas (1978) trees. Each tree's dividend stream follows a geometric Brownian motion. The representative investor has log utility and consumes the sum of the two trees' dividends. Prices adjust so that investors are happy to consume the dividends. We obtain closed-form solutions for prices, expected returns, volatilities, correlations, and so forth. Despite its

simple ingredients, and although the corresponding one-tree model produces a constant price-dividend ratio and i.i.d. returns, the two-tree model displays interesting dynamics.

The two trees can represent industries or other characteristic-based groupings. The two trees can represent broad asset classes such as stocks versus bonds, or stocks and bonds versus human capital and real estate. The two trees can represent two countries' asset markets, providing a natural benchmark for asset market dynamics in international finance.[1] Valuation ratios and market values of the two trees vary over time, so portfolio strategies that hold assets based on value/growth, small/large, momentum, or related characteristics give different average returns. Many interesting results continue to hold as one tree becomes vanishingly small relative to the other, so the model has implications for expected returns and dynamics of one asset relative to a much larger market.

Underlying the dynamics, we find that expected returns typically rise with a tree's share of dividends, to attract investors to hold that larger share. As a result, a positive dividend shock, which increases current prices and returns, also typically raises subsequent expected returns. Thus, returns tend to display positive autocorrelation or "momentum," prices typically seem to "underreact" or not to "fully adjust" to dividend news, and to "drift" upward for some time after that news. However, there are also parameters and regions of the state space in which expected returns decline as functions of the dividend share, leading to "mean reversion," price "overreaction," and "downward drift," with corresponding "excess volatility" of prices and returns.

When one asset has a positive dividend shock, this shock lowers the share of the other asset, so the expected return of the other asset typically declines. We see negative cross-serial correlation. We see movements in the other asset's price even with no news about that asset's dividends, a "discount rate effect," and another source and form of apparent "excess volatility." Finally, we see that asset returns can be positively contemporaneously correlated with each other even when their underlying dividends are independent. The lower expected return raises the price of the other asset. A "common factor" or "contagion" emerges in asset returns even though there is no common factor in cash flows.

Because price-dividend ratios vary despite i.i.d. dividend growth, price-dividend ratios forecast returns in the time series and in the cross section. Thus, we see "value" and "growth" effects: high price-dividend ratio assets have low expected returns and vice versa. We see that times in which a given asset has a high price-dividend ratio are *times* when that asset has a low expected return. Thinking of the "two trees" as stocks versus all other assets (bonds, real estate, human capital), price-dividend ratios forecast stock index returns. These effects coexist with the positive short-run autocorrelation of returns described above.

---

[1] Hau and Rey (2004), Guibaud and Coeurdacier (2006), and Pavlova and Rigobon (2007) are some recent articles that use multiple-tree frameworks to model countries.

However, although these dynamics are superficially reminiscent of those in the empirical asset pricing literature, and although we have used the terminology of that literature to help describe the model's dynamics, we do not claim that our model quantitatively matches the empirical literature. Our ingredients—the log utility function, the pure-endowment production structure, and especially the geometric Brownian motion dividend processes—are simple but empirically unrealistic.

Our model is simple because our goal is purely theoretical: to understand the dynamics induced by the market-clearing mechanism. Other papers in the emerging literature that price multiple cash-flow processes, including Bansal et al. (2002), Menzly et al. (2004), and Santos and Veronesi (2006), include non-i.i.d. dividend processes and temporally nonseparable preferences in order to better fit some aspects of the data. Other papers in the emerging general-equilibrium literature such as Gomes et al. (2003) and Gala (2006) add an interesting investment and production side to endogenize cash-flow dynamics.

For our aim, less is more. If we were to add these kinds of ingredients, we would no longer see what dynamics result from market clearing alone versus what dynamics come from the temporal structure of preferences, cash flows, or technology. We also would be forced to more complex and less transparent solution methods. As the standard one-tree model, though unrealistic, delivers useful insights into key asset-pricing issues, this simple two-tree model can isolate market-clearing effects that will be part of the story in more complex models, and allows us a clearer economic intuition for those effects.

Higher risk aversion and greater numbers of trees are desirable extensions that would not by themselves introduce dynamics. The first might raise the magnitudes of market-clearing dynamics, providing a better fit to the data, and the second is clearly important in its own right. However, our solution method works only for log utility and two trees.

Was it a mistake to believe in the logical possibility if not the reality of i.i.d. returns for all these years, given that our world does contain multiple nonzero net supply assets? The answer is no; it is possible to construct multiple nonzero net supply asset models with i.i.d. returns. For example, this can occur if the *supply* of assets changes instantly to accommodate changes in *demand*. If a price rise is instantly matched by a share repurchase, and a price decline is instantly matched by a share issuance, then the market value of each security can remain constant despite any variation in prices. No change in return moments is necessary, because no change in the market portfolio occurs. In this situation, investors can and do collectively rebalance.

In economics language, this situation is equivalent to the assumption of linear technologies. Output is a linear function of capital with no adjustment costs, diminishing returns, or irreversibilities, in contrast to our assumption of fixed endowments (trees, share supply). Investors can then instantly and costlessly transfer physical capital from successful projects to unsuccessful ones or to consumption, keeping all market weights constant. Such a linear technology

assumption is explicit, for example, in Cox et al. (1985). Most simply put, we can just write the i.i.d. rate of return processes or "technologies," and ask investors how much they want to hold, allowing them collectively to hold as much or as little as they want at any time. The resulting model will of course have i.i.d. returns.

Although such models are possible, these clearly unrealistic supply or technological underpinnings of i.i.d. returns are perhaps not often appreciated. It *is* a mistake to think that i.i.d. returns emerge from a multiperiod version of the usual market-equilibrium derivation of the CAPM, in which demand adjusts to a fixed supply of shares. And as the last example makes clear, we should really think of models that deliver i.i.d. returns in this way as "asset quantity" models, not "asset pricing" models. They are models of the composition of the market portfolio, because the asset prices and expected returns are given exogenously.

Which is the right assumption? In reality, market portfolio weights do change over time. Thus, a realistic model should have at least some short-run adjustment costs, irreversibilities, and other impediments to aggregate rebalancing. It will therefore contain some market-clearing dynamics of the sort we isolate and study in our simple two-tree exchange economy. On the other hand, new investment is made, new shares are issued, and some old capital is allowed to depreciate or is reallocated to new uses. Thus, a realistic model cannot specify a pure endowment structure. It needs some mechanism that allows aggregate rebalancing in the long run. The dynamics of the sort we study will apply less and less at longer horizons. In any case, this discussion and the examples of this paper make clear that the *technological* underpinnings of asset pricing models are more important to asset-price dynamics than is commonly recognized. Dynamics do not depend on preferences alone.

1. Model and Results

1.1 Model setup
The representative investor has log utility,

$$U_t = E_t \left[\int_0^\infty e^{-\delta\tau} \ln(C_{t+\tau}) \, d\tau \right]. \tag{1}$$

There are two trees. Each tree generates a dividend stream $D_i \, dt$. The dividends follow geometric Brownian motions with identical parameters,

$$\frac{dD_i}{D_i} = \mu \, dt + \sigma \, dZ_i, \tag{2}$$

where $i = 1, 2$, and dZ_i are standard Brownian motions, uncorrelated with each other.

To make the notation more transparent, we suppress time indices, for example, $dD_i/D_i \equiv dD_{it}/D_{it}$, etc., unless needed for clarity. Also, we focus on the first asset and suppress its index, for example, $dD/D = dD_1/D_1$. Finally,

since instantaneous moments are of order dt, we will typically omit the dt term in expressions for moments.

This is an endowment economy, so prices adjust until consumption equals the sum of the dividends, $C = D_1 + D_2$.

This economy is a straightforward generalization of the well-known one-tree model, and the one-tree model is the limit of our two-tree model as either tree becomes dominant. In the one-tree model, the price-dividend ratio is a constant, $P/D = 1/\delta$, and returns are i.i.d.

1.2 Dividend share dynamics

The relative sizes of the two trees give a single state variable for this economy. We find it convenient to capture that state by the dividend share,

$$s = \frac{D_1}{D_1 + D_2}. \tag{3}$$

Expected returns and other variables are functions of state, so we can understand their dynamics by understanding those functions of state and understanding how the state variable s evolves.

Applying Itô's Lemma to the definition in Equation (3), we obtain the dynamics for the dividend share process,

$$ds = -2\sigma^2 \, s(1 - s)\,(s - 1/2)\,dt + \sigma s(1 - s)(dZ_1 - dZ_2). \tag{4}$$

The drift of the dividend share process is an S-shaped function. The drift is zero when s equals 0, 1/2, or 1. The drift is positive for s between 0 and 1/2 and is negative for s between 1/2 and 1. Thus, there is a tendency for the dividend share to mean revert toward a value of 1/2. The two dividend-growth rates are independent, and there is no force raising an asset's dividend-growth rate if its share becomes small. Mean reversion in the share results from the nonlinear nature of the share definition in Equation (3) through second-order Itô's Lemma effects—the drift is zero when $\sigma^2 = 0$. In general, however, the drift is small so the dividend share is a highly persistent state variable—the path of its conditional means tends only slowly back to 1/2. For example, at $s = 1/4$, the drift is $3/32 \times \sigma^2$, and with $\sigma = 0.20$ that implies a drift of only 0.375 percentage points per year.

Share volatility is a quadratic function of the share, largest when the trees are of equal size, $s = 1/2$, and declining to zero at $s = 0$ or $s = 1$. Share volatility is substantial. For example, at $s = 1/2$, and with $\sigma = 0.20$, share volatility is five percentage points per year. In turn, this means that if expected returns are functions of the share, then they have the potential to vary significantly over time.

The dispersing effects of volatility overwhelm the mean-reverting effects of the drift, so this share process does not have a stationary distribution. A share process with a stationary distribution might be more appealing, but that would require putting dynamics in the dividend processes, so that small assets catch up. We want to make it clear that all dynamics in this model come from

market-clearing, not from dynamics of the inputs. We return to the issue of the nonstationary share and its effects on asset pricing below, after we see how asset prices behave in the model.

1.3 Consumption dynamics

In the one-tree model, consumption equals the dividend, so consumption growth is i.i.d. with constant mean $E[dC/C] = \mu$ and variance $\text{Var}[dC/C] = \sigma^2$.

In the two-tree model, aggregate consumption $C = D_1 + D_2$ follows

$$\frac{dC}{C} = \mu \, dt + \sigma \, s \, dZ_1 + \sigma \, (1 - s) \, dZ_2. \tag{5}$$

Consumption growth is no longer i.i.d. Mean consumption growth is still constant, but consumption volatility is a convex quadratic function of the share,

$$\text{Var}_t \left[\frac{dC}{C} \right] = \sigma^2 \left[s^2 + (1 - s)^2 \right]. \tag{6}$$

Consumption-growth volatility is still σ^2 at the limits $s = 0$ and $s = 1$, but declines to one-half that value at $s = 1/2$. Volatility is lower for intermediate values of the dividend share as consumption is then diversified between the two dividends.

1.4 The riskless rate

The investor's first-order conditions imply that marginal utility is a discount factor that prices assets, i.e.,

$$M_t = \frac{e^{-\delta t}}{C_t}. \tag{7}$$

The instantaneous interest rate is given from the discount factor by

$$r \, dt = -E_t \left[\frac{dM_t}{M_t} \right] = \delta \, dt + E_t \left[\frac{dC}{C} \right] - \text{Var}_t \left[\frac{dC}{C} \right]. \tag{8}$$

In the one-tree model, consumption equals the dividend so we have

$$r = \delta + \mu - \sigma^2. \tag{9}$$

We see the standard discount rate (δ), consumption growth (μ), and precautionary savings (σ^2) effects. Because the riskless rate is constant, the entire term structure is constant and flat. (We can compute a riskless rate, even though the riskless asset is in zero net supply.)

Substituting the moments of the consumption dynamics, Equation (5), into Equation (8), the interest rate in the two-tree model is

$$r = \delta + \mu - \sigma^2 \left[s^2 + (1 - s)^2 \right]. \tag{10}$$

Thus, the riskless rate varies over time as a quadratic function of the dividend share. The riskless rate is higher for intermediate values of the dividend share because dividend diversification lowers consumption volatility, which lowers the precautionary savings motive. Because the interest rate is not constant, the term structure is not flat.

1.5 Market portfolio price and return

As is usual in log utility models, the price-dividend ratio V_M for the market portfolio or claim to consumption stream is a constant

$$V_{Mt} \equiv \frac{P_{Mt}}{C_t} = \frac{1}{C_t} E_t \left[\int_0^\infty \frac{M_{t+\tau}}{M_t} C_{t+\tau} \, d\tau \right]$$

$$= E_t \left[\int_0^\infty e^{-\delta\tau} \frac{C_{t+\tau}}{C_{t+\tau}} \, d\tau \right] = \frac{1}{\delta}. \tag{11}$$

This calculation is the same for the one-tree and two-tree models, and it is valid for all consumption dynamics.

The total instantaneous return R_M on the market equals price appreciation plus the dividend yield,

$$R_M = \frac{dP_M}{P_M} + \frac{C}{P_M} \, dt = \frac{dC}{C} + \delta \, dt. \tag{12}$$

[In the second equality, we use the fact from Equation (11) that $P_M = C/\delta$.] Substituting in consumption dynamics, we find for the one-tree model that the market return is i.i.d.

$$R_M = (\mu + \delta) \, dt + \sigma \, dZ, \tag{13}$$

with constant expected return and variance:

$$E_t[R_M] = (\mu + \delta), \tag{14}$$

$$\mathrm{Var}_t[R_M] = \sigma^2. \tag{15}$$

For the two-tree model, the consumption dynamics in Equation (5) imply

$$R_M = (\mu + \delta) \, dt + \sigma \, s \, dZ_1 + \sigma \, (1 - s) \, dZ_2. \tag{16}$$

The expected market return and variance are now

$$E_t[R_M] = (\mu + \delta), \tag{17}$$

$$\mathrm{Var}_t[R_M] = \sigma^2[s^2 + (1 - s)^2]. \tag{18}$$

The expected return on the market is the same as in the one-tree case, but the variance of the market return equals the variance of consumption growth, which is now a quadratic function of the dividend share, declining for intermediate shares.

Subtracting the riskless rate in Equation (10) from the expected market return in Equation (17) shows that the equity premium equals the variance of the market return:

$$E_t[R_M] - r = \text{Var}_t[R_M], \tag{19}$$

as is usual in log utility models. From Equation (18), the variance of the market is a convex quadratic function of the dividend share. This fact means that the equity premium and market Sharpe ratios are also time varying, and increase as the market becomes more polarized.

1.6 The price-dividend ratio

The price P of the first asset is given by

$$P_t = E_t\left[\int_0^\infty e^{-\delta\tau} \frac{C_t}{C_{t+\tau}} D_{t+\tau}\, d\tau\right]. \tag{20}$$

Again, we suppress asset and time subscripts unless necessary for clarity, and we focus on the first asset because the second follows by symmetry.

For the remainder of this section, we impose the parameter restriction $\delta = \sigma^2$. This restriction, in conjunction with the symmetry of the assets, gives much simpler formulas and more transparent intuition than the general case. This restriction is not unreasonable: with $\sigma = 0.20$, $\delta = 0.04$. In the following section, we treat the general case that breaks this restriction, allows the trees to have different values of μ and σ, and allows correlated shocks. We present formulas using whichever of δ or σ^2 gives a more intuitive appearance.

Using the definition of the dividend share in Equation (3), Equation (20) can be rewritten to give the price-consumption ratio for the first asset:

$$\frac{P}{C} = E_t\left[\int_0^\infty e^{-\delta\tau} s_{t+\tau}\, d\tau\right]. \tag{21}$$

Valuing the asset is formally identical to risk-neutral pricing (using the discount rate δ) of an asset that pays a cash flow equal to the dividend share. The dividend share plays a similar role in many tractable models of long-lived cash flows, including Bansal et al. (2002), Menzly et al. (2004), Longstaff and Piazzesi (2004), and Santos and Veronesi (2006).

Solving Equation (21) with share dynamics from Equation (4), the price-dividend ratio of the first asset is

$$V \equiv \frac{P}{D} = \frac{1}{s}\frac{P}{C} = \frac{1}{2\delta s}\left[1 + \left(\frac{1-s}{s}\right)\ln(1-s) - \left(\frac{s}{1-s}\right)\ln(s)\right]. \tag{22}$$

The Appendix gives a short proof, along the following lines. First, because s follows a Markov process, Equation (21) implies that P/C is a function of s. Using Itô's Lemma and the dynamics for s from Equation (4), we obtain an

expression for $E[d(P/C)]$ in terms of the first and second derivatives of P/C. Second, from Equation (21), the price-consumption ratio satisfies

$$E\left[d\left(\frac{P}{C}\right)\right] = \delta\frac{P}{C} - s.$$
(23)

Equating these two expressions for $E[d(P/C)]$, we obtain a differential equation for P/C. We then verify that Equation (22) solves this differential equation.

The formula in Equation (22) and subsequent ones are all given in terms of elementary functions, and thus they can be characterized straightforwardly. However, it is easier and makes for better reading simply to plot the functions, so we follow that route in our analysis. For comparability, we use the parameter values $\delta = 0.04$, $\mu = 0.02$, and $\sigma = 0.20$ throughout all the plots presented in this section.

Figure 1 plots the price-consumption and price-dividend ratios for $\delta = 0.04$. The price-consumption ratio lies close to a linear function of the dividend share. This behavior is largely a scale effect; larger dividends command higher prices. The slightly S-shaped deviations from linearity are thus the most interesting features of this graph. The price-consumption ratio initially is higher than the share and then falls below the share. We anticipate that the first asset will have a "high" price at low shares and a "low" price at high shares.

The price-dividend ratio varies greatly from the constant value $P/D = 1/\delta = 25$ of the one-tree model. This variation of the price-dividend ratio as a function of the dividend share drives the return dynamics that follow. The price-dividend ratio is equal to the price-consumption ratio divided by share s, so its behavior is driven by the small deviations from linearity in the price-consumption ratio function.

The price-dividend ratio in Figure 1 is largest for small shares, declines until its value is less than the price-dividend ratio of the market portfolio starting at $s = 0.5$, and then increases to equal the market price-dividend ratio of 25 when $s = 1$. This initially surprising nonmonotonic behavior is necessary, because the constant price-dividend ratio of the market portfolio must equal the share-weighted mean of the price-dividend ratios of the individual assets. If the price-dividend ratio of the first asset is greater than that of the market at a share of, say, 0.25, then by symmetry, it must be less than that of the market at a share of 0.75. It must then recover to equal the market price-dividend ratio when it *is* the market at $s = 1$. The S-shape of the price-consumption ratio about the 45-degree line is exactly symmetric in this way, reflecting the symmetry of Equation (22).

The price-dividend ratio increases rapidly as the dividend share decreases toward zero. The basic mechanism for this behavior is a decline in risk premium. As the first asset share declines, its dividends become less correlated with aggregate consumption. This fact lowers their risk premium and discount rate,
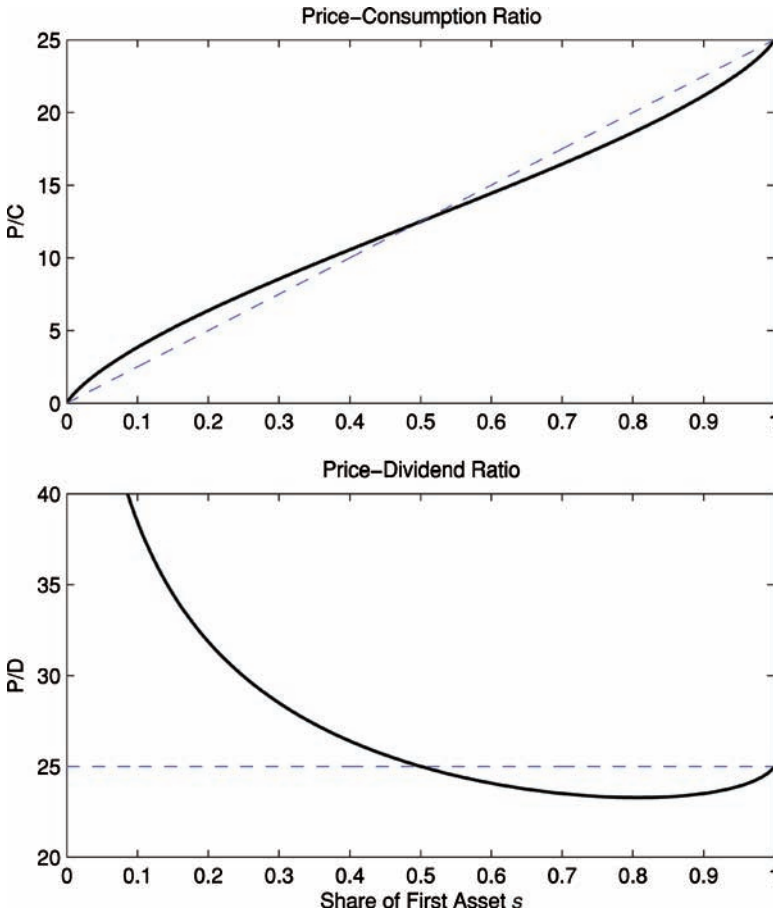
Figure 1
Price-consumption ratio and price-dividend ratio
The top panel presents the price-consumption ratio of the first asset, in the simple case, using parameters $\mu = 0.02$, $\delta = \sigma^2 = 0.04$. The bottom panel presents the price-dividend ratio. The dashed line in the top panel gives the 45-degree line; the dashed line in the bottom panel gives the price-dividend ratio of the market portfolio and of the one-tree model, $1/\delta = 25$.

raising their valuations. An asset with a small share is more valuable from a diversification perspective.

As $s \to 0$, the price-dividend ratio rises to infinity. As $s \to 0$, the first tree's dividend becomes completely uncorrelated with consumption, because consumption consists entirely of the second tree's dividend. As a result, the first tree is valued as a risk-free security. In this parameterization, the $s = 0$ limit of the interest rate equals the dividend-growth rate $r = \mu$, so the price-dividend ratio of the growing dividend stream explodes. Although the rise in the price-dividend ratio as $s \to 0$ is generic, because the asset becomes less risky, the

left limit is finite for other parameterizations in which the interest rate exceeds the mean dividend-growth rate.

1.7 Asset returns and moments

Returns follow now that we know prices. Again suppressing asset and time subscripts, the instantaneous return R of the first asset is

$$R \equiv \frac{D}{P} \, dt + \frac{dP}{P}. \tag{24}$$

For both calculations and intuition, it is convenient to express this return in terms of the price-dividend or valuation ratio,

$$R = \frac{1}{V} \, dt + \frac{dD}{D} + \frac{dV}{V} + \frac{dD}{D} \frac{dV}{V}. \tag{25}$$

The return equals the dividend yield, the dividend growth, the change in valuation or growth in the price-dividend ratio, and an Itô term.

Although the price-dividend ratio is constant in the one-tree model, the last two terms in Equation (25) are zero and the asset return is just the dividend yield plus dividend (consumption) growth. With two assets, as we have seen, the price-dividend ratio is no longer constant and varies through time as the relative weights of the assets evolve, so the latter two terms matter. Both the mean and variance of returns vary over time as functions of the state variable s.

Now we can find return moments. Taking the expectation of Equation (25), the expected return of the first asset is

$$E_t[R] = \frac{1}{V} \, dt \, + \, E_t \left[\frac{dD}{D} \right] \, + \, E_t \left[\frac{dV}{V} \right] \, + \, \mathrm{Cov}_t \left[\frac{dD}{D}, \frac{dV}{V} \right]. \tag{26}$$

The instantaneous variance of the first asset's return is

$$\begin{aligned}
\mathrm{Var}_t[R] &= \mathrm{Var}_t \left[\frac{dD}{D} + \frac{dV}{V} \right] \\
&= \mathrm{Var}_t \left[\frac{dD}{D} \right] + \mathrm{Var}_t \left[\frac{dV}{V} \right] + 2 \, \mathrm{Cov}_t \left[\frac{dD}{D}, \frac{dV}{V} \right].
\end{aligned} \tag{27}$$

In the single-asset model with a constant valuation ratio, the variance of returns is equal to the variance of dividend growth, the first term in Equation (27). In the two-asset model, variation in the price-dividend ratio can provide additional return volatility through the second term in Equation (27). However, if the price-dividend ratio is strongly negatively correlated with dividend growth—if an increase in dividends and thus share strongly reduces the price-dividend ratio—then return volatility can be less than dividend-growth volatility through the third term in Equation (27).

Applying Itô's Lemma to the price-dividend ratio V in Equation (22) and using the share process in Equation (4) gives the following expressions for

moments of dV/V in Equations (26) and (27):

$$E_t\left[\frac{dV}{V}\right] = \delta\left(1 + 3(1 - s)\right) - \frac{1}{V}\left(1 - 2\ln(s)\right), \qquad (28)$$

$$\text{Var}_t\left[\frac{dV}{V}\right] = \frac{2}{\delta}\left[\delta(1 + (1 - s)) + \frac{1}{1 - s}\frac{\ln(s)}{V}\right]^2, \qquad (29)$$

$$\text{Cov}_t\left[\frac{dD}{D}, \frac{dV}{V}\right] = -\left[\delta(1 + (1 - s)) + \frac{1}{1 - s}\frac{\ln(s)}{V}\right]. \qquad (30)$$

Substituting the expected dividend-growth rate $E_t[dD/D] = \mu$ and Equations (28) and (30) into Equation (26) and rearranging gives us the expected return as a function of state,

$$E_t[R] = \mu + 2\delta(1 - s) + \left(1 - \frac{s}{1 - s}\right)\frac{\ln(s)}{V}. \qquad (31)$$

Subtracting the riskless rate in Equation (10) from the expected return in Equation (31) gives the expected excess return

$$E_t[R] - r = 2\delta(1 - s)^2 + \left(1 - \frac{s}{1 - s}\right)\frac{\ln(s)}{V}. \qquad (32)$$

Similarly, substituting from Equations (29) and (30) into Equation (27) gives the return variance

$$\text{Var}_t[R] = \frac{\delta}{2} + 2\delta\left[\frac{1}{2} + (1 - s) + \frac{1}{\delta}\frac{1}{1 - s}\frac{\ln(s)}{V}\right]^2. \qquad (33)$$

Figure 2 plots the expected return given in Equation (31) as a function of the dividend share. Expected returns rise with the share, reach an interior maximum, and then decline slightly. This behavior of expected returns in Figure 2 mirrors the behavior of the price-dividend ratio in Figure 1. With constant expected dividend growth, expected returns are the only reason price-dividend ratios vary at all, so low expected returns must correspond to high price-dividend ratios. One way to understand the nonmonotonic behavior of expected returns, then, is as the mirror image of the nonmonotonic behavior of the price-dividend ratio studied above. Expected returns must be higher than the market expected return for high shares (near $s = 0.8$) so that expected returns can be lower than the market expected return for small shares.

The behavior of expected returns as a function of state in Figure 2 drives asset return dynamics. A positive shock to the first asset's dividends increases the dividend share. For share values below about 0.80, this event increases expected returns. In this range, then, a positive dividend shock leads to a string of expected price increases. Prices will seem to "underreact" and "slowly"
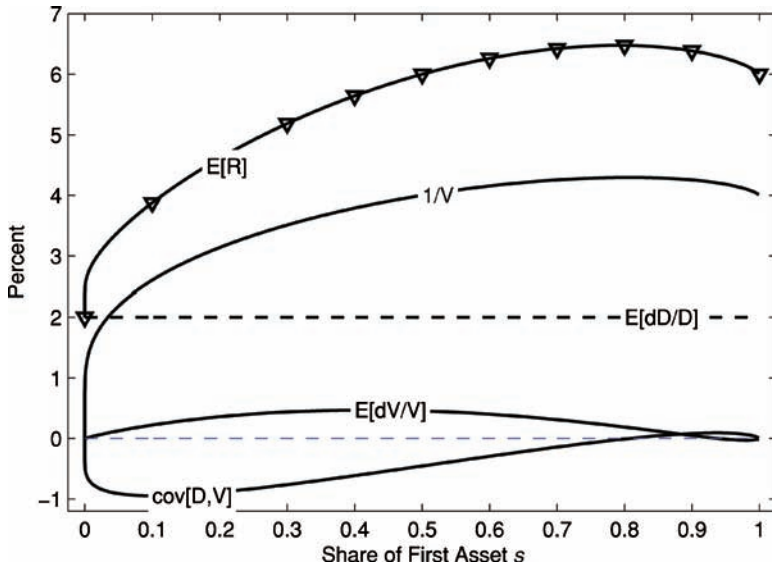
Figure 2
Expected return and components
$E[R]$ gives the expected return of the first asset as a function of the share of the first asset. The remaining lines give components of this expected return. $1/V$ gives the dividend-price ratio. $E[dD/D]$ gives the expected dividend-growth rate. $E[dV/V]$ gives the expected change in the price-dividend ratio. $Cov[D,V]$ gives $Cov[dV/V,dD/D]$, the covariance of dividend growth with price-dividend ratio shocks.

incorporate dividend news. To the extent that own-dividend shocks dominate other-dividend shocks as a source of price movement, we expect to see here positive autocorrelation and "momentum" of returns.

For share values above about 0.80, however, expected returns decline in the dividend share. Here, a positive own-dividend shock leads to lower subsequent expected returns; we see "overreaction" to or "mean reversion" after the dividend shock, and we expect to see negative autocorrelation, mean reversion, and "excess volatility" of returns.

Equation (26) expresses the expected return as a sum of four components. The individual components are also plotted in Figure 2. As illustrated, the dividend yield $1/V$ is generally the largest component of expected returns, followed closely by the expected growth rate of dividends $E[dD/D]$. For small dividend shares, the negative covariance between dividends and the valuation ratio reduces the expected return substantially. The negative covariance appears because a positive shock to dividends increases the share, and this has a strong negative effect on the valuation ratio as per Figure 1. The expected change in the valuation ratio is small, generally positive, and has its largest effect on the expected return for small to intermediate values of the dividend ratio. All three of these components change sign or slope at the point where the price-dividend ratio starts rising as a function of share, for high share values.

Expected excess returns represent risk premia, reflecting the covariance of returns with consumption growth,

$$E_t[R] - r = \text{Cov}_t\left[R, \frac{dC}{C}\right]. \tag{34}$$

We can express this risk premium as the sum of a "cash-flow beta," corresponding to the covariance of dividend growth with consumption growth, and a "valuation beta," corresponding to the covariance of valuation shocks with consumption growth. Substituting Equation (25) into Equation (34), we have

$$E_t[R] - r = \text{Cov}_t\left[\frac{dD}{D}, \frac{dC}{C}\right] + \text{Cov}_t\left[\frac{dV}{V}, \frac{dC}{C}\right]. \tag{35}$$

We can evaluate the terms on the right-hand side using the dividend-growth and consumption processes in Equations (2) and (5),

$$\text{Cov}_t\left[\frac{dD}{D}, \frac{dC}{C}\right] = \sigma^2 s = \delta s, \tag{36}$$

$$\text{Cov}_t\left[\frac{dV}{V}, \frac{dC}{C}\right] = \text{Cov}_t\left[\frac{dD}{D}, \frac{dV}{V}\right](2s - 1)$$

$$= (1 - 2s)\left[\delta(1 + (1 - s)) + \frac{1}{1-s}\frac{\ln(s)}{V}\right]. \tag{37}$$

In the last equality, we have substituted from Equation (30).

The "cash-flow beta" expresses what would happen if price-dividend ratios were constant. Then, the covariance of returns with consumption growth would be exactly proportional to the covariance of dividend growth with consumption growth. This covariance would of course be larger precisely as the first tree's dividend provides a larger share of consumption. Thus, the "cash-flow beta" is linear in the share.

The "valuation beta" is more interesting, as it captures the fact that price-dividend ratios change as well, and changes in valuation that covary with the consumption growth generate a risk premium. Valuation betas capture the fact that return dynamics—changes in expected return, which change valuations—spill over to the level of the expected return, as in Merton (1973).

Figure 3 plots expected excess returns from Equation (35), the risk-free rate from Equation (10), and the "cash-flow" and "valuation betas" from Equations (36) and (37), respectively. As shown, the expected excess return starts at zero, but then increases rapidly as the dividend share increases. Expected excess returns rise uniformly following a positive dividend shock, so we expect to see the positive autocorrelation dynamics throughout the share range using this measure (of course, expected excess returns remain positive as $s \to 0$ if one allows correlated cash flows).
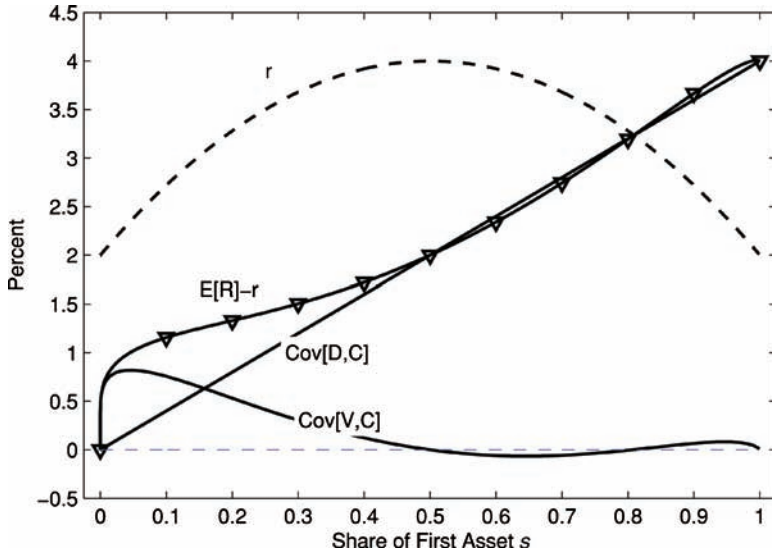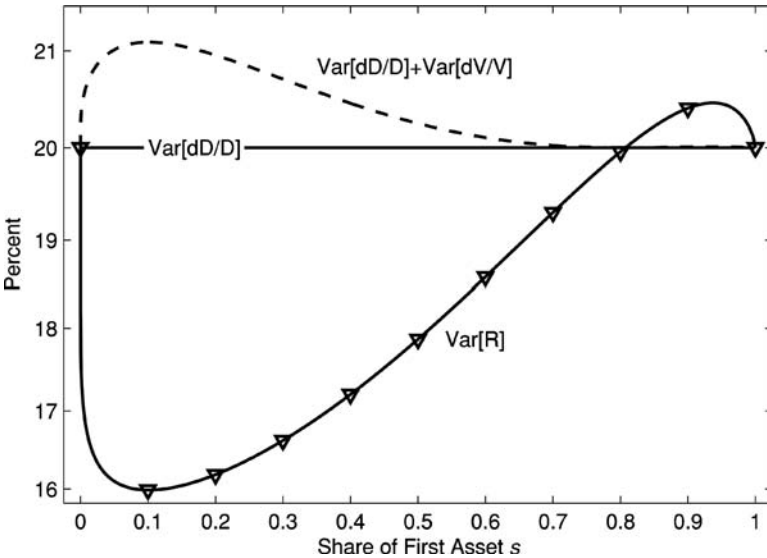
Figure 3
Expected excess return and components
$E[R] - r$ gives the expected excess return of the first asset as a function of its share. Cov[D,C] gives the covariance of dividend and consumption-growth shocks, Cov[dD/D,dC/C]. Cov[V,C] gives the covariance of price-dividend ratio and consumption-growth shocks, Cov[dV/V,dC/C]. The latter two components add up to the expected excess return. The riskless rate is given by r.

The "valuation beta" can have either sign. It is slightly negative for dividend share values between about 0.50 and 0.80. For small shares, the "valuation beta" is dominant. For small shares, the risk premium is due primarily to changes in valuations correlated with the market "discount rate" effect—rather than changes in dividends or cash flows correlated with the market. An observer might be puzzled why there is so much return correlation, beta, and expected return in the face of so little correlation of cash flows.

In the limit as $s \to 0$, the expected excess return collapses to zero. (One can show this fact analytically by taking limits of the above expressions.) In that limit, the first tree's dividends are completely uncorrelated with consumption. However, as Figure 3 demonstrates, the expected excess return rises very quickly from zero and is substantial even for very small shares. For example, the expected excess return is already 1/2% at $s = 0.1\%$. Formally, the derivative of expected excess return with respect to share rises to infinity as $s \to 0$. Therefore, the market-clearing-induced risk premium and return dynamics remain important for "small" assets.

Figure 4 plots the return volatility given in Equation (33) along with the components of that volatility from Equation (27). Dividend-growth Var[dD/D] gives a constant contribution of 20% volatility. Changes in valuation Var[dV/V] add a small amount of volatility at small share values, where the valuation in Figure 1 is a strong function of the share. The negative
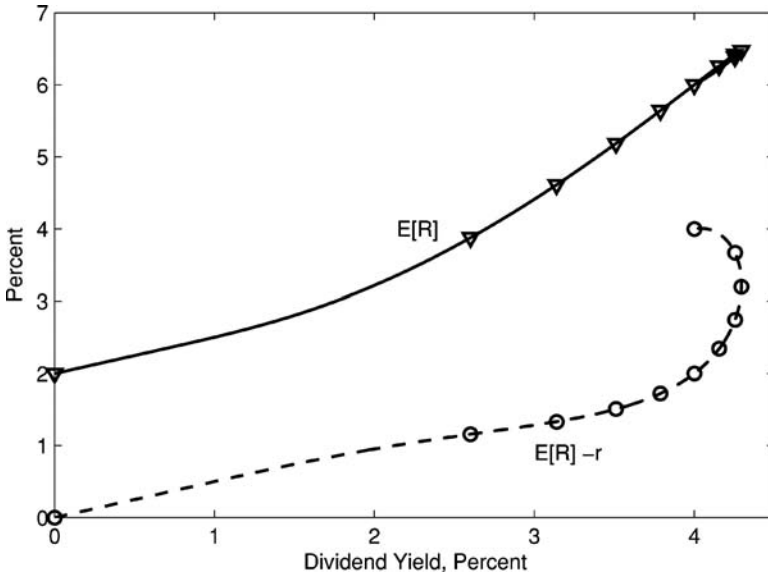
Figure 4
Return volatility and components
The solid line with symbols labeled Var[*R*] gives the variance of the first asset's return as a function of the first asset's share. (The vertical axis labels give percent standard deviation. However, the vertical axis plots variance, so that the variance components add up.) The horizontal solid line labeled Var[d*D/D*] gives the dividend-growth component of return variance. The dashed line labeled Var[d*D/D*]+Var[d*V/V*] adds the dividend-growth variance and the variance of the price-dividend ratio. Adding the third, usually negative, component Cov[d*D/D*,d*V/V*] gives the overall return variance.

value of the covariance term in Equation (33) is a larger effect and pushes over-all variance below dividend-growth variance for $s < 0.80$. The price-dividend ratio is a declining function of share here, so a positive dividend shock lowers the price-dividend ratio (Figure 1). As a result, returns (price plus dividend) move less than dividends themselves. As the share increases, however, the covariance term eventually becomes positive, where the price-dividend ratio rises with the share, and adds to the total return variance. Thus, there is a small region of "excess volatility" in which the volatility of the asset's return exceeds the volatility of the underlying cash flows or dividends. Again, the derivative of volatility with respect to share becomes infinite as $s \to 0$, so market-clearing effects apply to very small assets.

Because the price-dividend ratio, expected return, and expected excess return are all functions of the share, we can substitute out the share and plot expected returns and excess returns as functions of the dividend yield. Figure 5 presents the results.

Both expected returns and expected excess returns are increasing functions of the dividend yield. Therefore, dividend yields (or price-dividend ratios) forecast returns in the time series and in the cross section. Expected returns show a nicely linear relation to dividend yields through most of the relevant range. Expected *excess* returns show an intriguing nonlinear relation. The slope

Figure 5
Dividend yields and expected returns
The solid line plots the first asset's expected return versus its dividend yield. The dashed line plots the first asset's expected excess return versus its dividend yield. Symbols mark the points $s = 0$, $s = 0.1$, $s = 0.2$, etc., starting from the left.

of the return line is about 1.6 through the linear portion. A slope of one means that higher dividend yields translate to higher expected returns one-for-one. Higher slopes mean that a high dividend yield forecasts valuation increases as well.
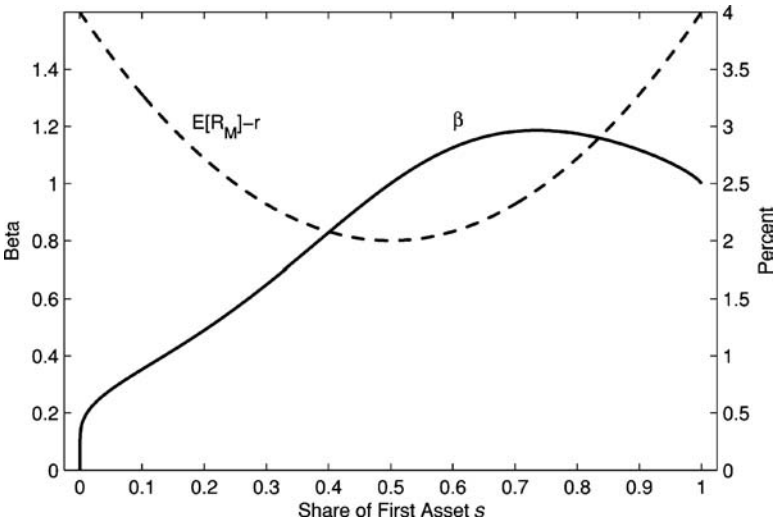
1.8 Market betas

With log utility, expected returns follow a conditional CAPM and consumption CAPM. Thus, we can also understand expected excess returns by reference to the asset's beta and the market expected excess return.

In the single-tree model, the asset is the market, and its beta equals one. In the two-tree model, the beta of each asset varies over time with the dividend share. Using the fact from Equation (12) that the market return equals consumption growth plus the discount rate, we have

$$\beta = \frac{\text{Cov}_t [R, R_M]}{\text{Var}_t [R_M]} = \frac{\text{Cov}_t [R, dC/C]}{\text{Var}_t [R_M]}. \tag{38}$$

Substituting the covariances from Equations (36) and (37), and using the market return variance in Equation (18), we obtain

$$\beta = \frac{2\sigma^2(1 - s)^2 + 1 - \frac{s}{1-s}\frac{\ln(s)}{V}}{\sigma^2 \left[s^2 + (1 - s)^2 \right]}. \tag{39}$$

Figure 6
Market betas and market risk premium
The solid line labeled β gives the beta on the market portfolio of the first asset return. Its values are plotted on the left vertical axis. The dashed line labeled $E[R_M] - r$ gives the expected excess return of the market portfolio. Its values are plotted on the right vertical axis.

(We present Equation (39) in terms of σ^2, which is more intuitive for a second moment, but this formula is only valid under the restriction $\sigma^2 = \delta$ of this section.)

Figure 6 plots beta as a function of the dividend share, revealing interesting dynamics. As shown, the beta is zero when the share is zero. As the share increases, the beta rises quickly, in fact infinitely quickly as $s \to 0$. As we can see in Figure 3, this rise is due to the large "valuation beta" for small assets. As the share rises, the beta continues to rise almost linearly. Here, the nearly linear "cash-flow beta" of Figure 3 is at work: the first asset contributes more to the total market return and its beta begins to increase correspondingly.

At a share $s = 0.5$, the beta becomes greater than one, and then declines until it becomes one again when the first asset is the entire market. As before, we can start to understand this nonmonotonic behavior by aggregation: the share-weighted average beta must be one, so if the small asset has a beta less than one, the larger asset must have a beta larger than one. As the share approaches one, however, the beta begins to decrease and converges to one because the first asset becomes the market as $s \to 1$.

The expected excess return of Figure 3 is equal to the beta of Figure 6 times the market expected excess return; $E_t[R_M] - r = \sigma^2 \left[s^2 + (1-s)^2 \right]$ from Equations (18) and (19) and is also included in Figure 6. The decline in the market expected excess return from $s = 0$ to $s = 1/2$ accounts for the

slightly lower rise in expected excess return in Figure 3 compared to the rise in market beta in Figure 6 in this region. The rise in market expected excess return from $s = 1/2$ to $s = 1$ offsets the decline in beta shown in Figure 6, allowing the nearly linear rise in expected excess return shown in Figure 3.

In sum, although a conditional CAPM and consumption CAPM hold in this model, one must make reference to time-varying expected excess returns, expected excess market returns, and market betas in order to see the relations predicted by the CAPM or consumption CAPM.

1.9 Return correlations

The *returns* of the two assets can be correlated even though their *dividends* are independent. To see this fact, we can write from Equation (25),

$$
\text{Cov}_t\left[R_1, R_2\right] = \text{Cov}_t\left[\frac{dD_1}{D_1}, \frac{dD_2}{D_2}\right] + \text{Cov}_t\left[\frac{dD_1}{D_1}, \frac{dV_2}{V_2}\right]
$$
$$
+ \text{Cov}_t\left[\frac{dD_2}{D_2}, \frac{dV_1}{V_1}\right] + \text{Cov}_t\left[\frac{dV_1}{V_1}, \frac{dV_2}{V_2}\right]. \quad (40)
$$

Because the dividends are independent, the first term on the right-hand side is zero. If the price-dividend ratios V_i for the two assets were constants, the remaining three terms on the right-hand side would also be zero and the returns of the two assets would be uncorrelated. However, the price-dividend ratios vary over time and are correlated with each other and with the dividends. Thus, the correlation of the assets' returns is generally not equal to zero.

Figure 7 plots the correlation between the assets' returns as a function of the dividend share. As shown, the returns have a correlation above 25% for most of the range of dividend shares, even though the underlying cash flows are not correlated.

The mechanisms are straightforward. If tree two enjoys a positive dividend shock, that event raises asset two's return. However, it also lowers asset one's share. Lowering the share typically raises the price-dividend ratio—i.e., gives rise to a positive return for asset one. This story underlies the $\text{Cov}_t\left[dD_2/D_2, dV_1/V_1\right]$ component of Equation (40), graphed as the "$[V_1, D_2]$" line of Figure 7. As expected it gives a positive contribution to correlation for shares below $s = 0.8$, in which the price-dividend ratio in Figure 1 is rising in the share. Adding the symmetrical contribution to correlation from the effect of asset one's dividends on asset two's valuations gives the dashed line marked $[V_1, D_2] + [V_2, D_1]$ in Figure 7, and we see that these two effects are most of the story for the overall correlation, marked $[R_1, R_2]$ in Figure 7.

Again, the correlations are zero in the limit as $s \to 0$ or $s \to 1$, but the derivative becomes infinite at these limits so market-clearing-induced correlations are large for vanishingly small assets.
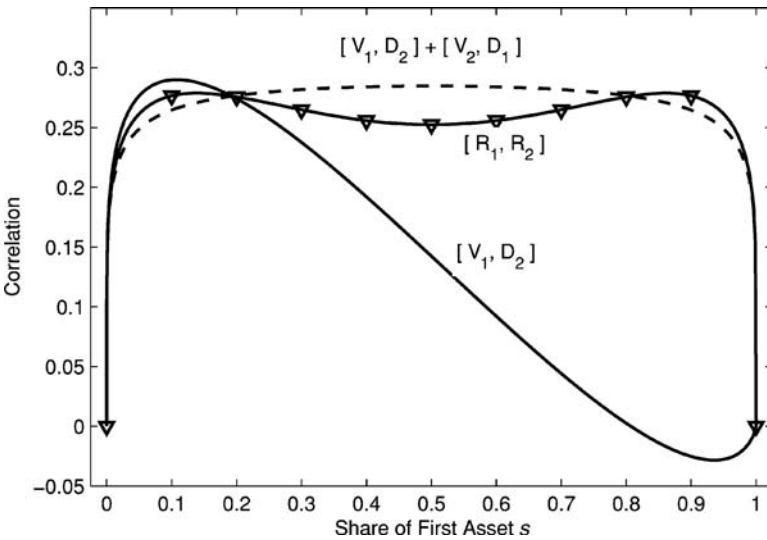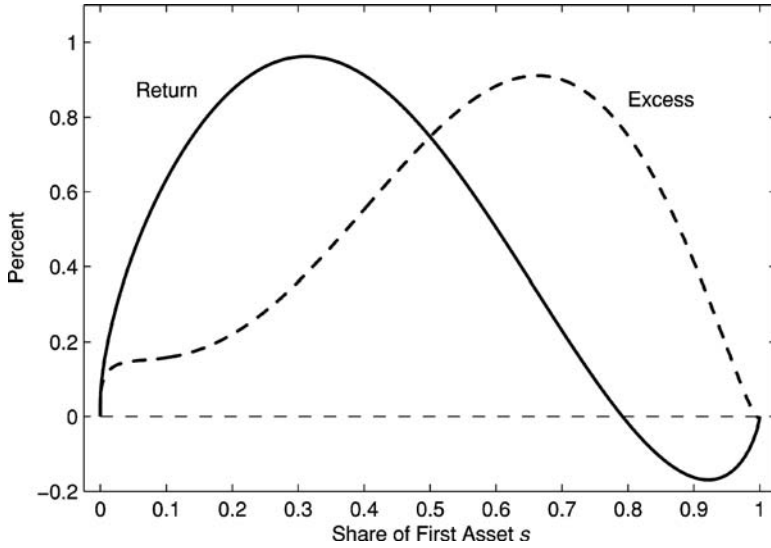
Figure 7
Return correlation
The solid line with triangles labeled $[R_1, R_2]$ gives the conditional correlation between the two assets' returns, given the dividend share s of the first asset. The remaining lines give components of that correlation, using the decomposition of Equation (40). For example, $[V_1, D_2]$ gives the component of correlation corresponding to $\mathrm{Cov}\left[\frac{dV_1}{V_1}, \frac{dD_2}{D_2}\right]$.

1.10 Autocorrelation

How important are the return dynamics in the two-tree model? How much do expected returns vary over time? Do market-clearing return dynamics vanish for small assets?

As one way to answer these questions, Figure 8 presents the instantaneous autocorrelation of returns. In discrete time, the autocorrelation is the regression coefficient of future returns on current returns, $\mathrm{Cov}[R_{t+1}, R_t]/\mathrm{Var}[R_t] = \mathrm{Cov}[E_t(R_{t+1}), R_t]/\mathrm{Var}[R_t]$. We compute a continuous-time conditional counterpart to the second expression, $\mathrm{Cov}_t\ [dER_t, R_t]/\mathrm{Var}_t[R_t]$. R_t denotes the instantaneous expected return. ER_t denotes the expected return that is a function of the state as plotted in Figure 2. Thus, we can apply Itô's Lemma to find dER_t and the covariance follows. The result expresses how much a return shock raises subsequent *expected* returns.

The pattern of autocorrelation shown in Figure 8 is consistent with the patterns of expected returns and excess returns illustrated in Figures 2 and 3. The autocorrelation of returns is positive where the expected return in Figure 2 rises with the share, and the autocorrelation is negative where expected returns in Figure 2 decline with share. Expected *excess* returns in Figure 3 rise with share throughout, and we see a positive autocorrelation throughout. This result is not automatic, as a positive return can also be caused by an increase in the other asset's dividend, which typically lowers the expected return. Figure 8

Figure 8
Return autocorrelation
The "return" line is calculated as $\text{Cov}[dER_t, R_t]/\text{Var}[R_t]$ where R_t is the instantaneous return, and $ER_t = E_t[R_t]$ is the expected value of the instantaneous return. It measures how much a unit instantaneous return raises the subsequent expected return. The "excess" line does the same calculation for expected excess returns.

shows that this effect is dominated by the own-dividend shock, allowing a positive autocorrelation to emerge. The magnitude of the autocorrelation is small, reaching only one percentage point.[2]

In the limit $s \to 0$, both autocorrelations vanish. However, the slope of the autocorrelation tends to infinity as $s \to 0$, so again market-clearing effects remain important for vanishingly small assets.

2. The General Model

This section presents the general case of the two-asset model. Dividends still follow geometric Brownian motions, but we allow different parameters,

$$\frac{dD_i}{D_i} = \mu_i \, dt + \sigma_i \, dZ_i. \tag{41}$$

The correlation between dZ_1 and dZ_2 is ρ, not necessarily zero. We also allow the discount rate and the volatility of dividend growth to differ, $\sigma^2 \neq \delta$.

---

[2] Autocorrelation is a common and intuitive measure of return predictability. However, lagged returns do not contain all information that is available at time t. It is possible for returns to be predictable—for example, by dividend yields—while not autocorrelated. The variance of expected returns, which is the variance of the numerator of R^2 in a multivariate return-forecasting regression, is a more comprehensive measure. However, our computation of the continuous-time counterpart to this quantity $\text{Var}[dER_t]$ does not differ enough from Figure 8 to warrant presentation.

2.1 Consumption dynamics

Applying Itô's Lemma to $C = D_1 + D_2$ implies

$$\frac{dC}{C} = \left[\mu_1 s + \mu_2(1 - s)\right] dt + \sigma_1 s\, dZ_1 + \sigma_2(1 - s)\, dZ_2. \qquad (42)$$

As before, consumption growth is no longer i.i.d. through time. The mean consumption growth,

$$E_t \left[\frac{dC}{C}\right] = \mu_1 s + \mu_2(1 - s), \qquad (43)$$

is the share-weighted mean of the dividend-growth rates and so is no longer constant. Consumption volatility,

$$\text{Var}_t \left[\frac{dC}{C}\right] = \sigma_1^2 s^2 + \sigma_2^2(1 - s)^2 + 2\rho\sigma_1\sigma_2 s(1 - s), \qquad (44)$$

is again a convex quadratic function of the share, lower where consumption is diversified across the two trees.

2.2 The riskless rate

Substituting the consumption moments in Equations (43) and (44) into the expression for the riskless rate in Equation (8) gives the riskless rate in the general two-asset model,

$$r = \delta + \mu_1 s + \mu_2(1 - s) - \sigma_1^2 s^2 - \sigma_2^2(1 - s)^2 - 2\rho\sigma_1\sigma_2 s(1 - s). \qquad (45)$$

The riskless rate is again a quadratic function of the dividend share. If the means or volatilities of the dividend streams differ, it is no longer symmetric, however.

2.3 Market price and dynamics

The market price and its dynamics are virtually the same as in the simple case. The price-dividend ratio V_M for the market is still $1/\delta$, and the instantaneous return R_M on the market equals the percent change in aggregate consumption plus the dividend yield as in Equation (12). Thus, the expected return and variance of the market differ only in that the moments of the consumption process differ in the general model,

$$E_t[R_M] = \delta + \mu_1 s + \mu_2 (1 - s), \qquad (46)$$
$$\text{Var}_t[R_M] = \sigma_1^2 s^2 + \sigma_2^2 (1 - s)^2 + 2\rho\sigma_1\sigma_2 s (1 - s). \qquad (47)$$

The equity premium again equals the variance of the market as in Equation (19).

2.4 Dividend share dynamics

An application of Itô's Lemma gives the dynamics of the dividend share,[3]

$$ds = s(1-s)\left[\mu_1 - \mu_2 - s\sigma_1^2 + (1-s)\sigma_2^2 + 2(s-1/2)\rho\sigma_1\sigma_2\right]dt$$
$$+ s(1-s)(\sigma_1\,dZ_1 - \sigma_2\,dZ_2). \tag{48}$$

The drift of this dividend share process is zero when $s = 0$, κ, or 1, where

$$\kappa = \frac{\mu_1 - \mu_2 + \sigma_2^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}. \tag{49}$$

When κ lies between zero and one, the drift is positive from zero to κ, bringing the share up toward κ, and negative from κ to one, bringing the share down toward κ. We see a more general version of the same S-shaped mean reversion that characterizes our simple case. The diffusion coefficient in Equation (48) is again quadratic, implying that changes in the dividend share are most volatile when $s = 1/2$. As in the simple case, one tree eventually will dominate the other so this model does not possess a stationary share distribution.[4]

2.5 Asset prices

The first asset's price-dividend ratio is still derived as a discounted "present value" of future shares as in Equation (21). The share process has changed to Equation (48), however, so the form of the solution changes. The Appendix shows that in this general case, the price-dividend ratio of the first asset V can be expressed as

$$V = \frac{1}{\psi(1-\gamma)(1-s)}\,F\left(1,\ 1-\gamma;\ 2-\gamma;\ \frac{s}{s-1}\right)$$
$$+ \frac{1}{\psi\theta s}\,F\left(1,\ \theta;\ 1+\theta;\ \frac{s-1}{s}\right), \tag{50}$$

where

$$\psi = \sqrt{\nu^2 + 2\delta\eta^2}, \qquad \gamma = \frac{\nu - \psi}{\eta^2}, \qquad \theta = \frac{\nu + \psi}{\eta^2},$$

and

$$\nu = \mu_2 - \mu_1 - \sigma_2^2/2 + \sigma_1^2/2, \qquad \eta^2 = \sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2.$$

---

[3] The share process is a member of the Wright-Fisher class of diffusions. These types of diffusions are often applied in genetic theory to characterize the evolution of genes in a population of two genetic types. For example, Karlin and Taylor (1981, ch. 15, pp. 184–88) present an example in which population shares follow a diffusion similar to Equation (48). Also see Crow and Kimura (1970) for examples and a discussion of the asymptotic properties of these models. The cubic drift of our share process is also closely related to that of the stochastic Ginzburg-Landau diffusion used in superconductivity physics to model phase transitions. See Kloeden and Platen (1992) and Katsoulakis and Kho (2001).

[4] This feature parallels the asymptotic properties of Wright-Fisher gene frequency models in which one of the two gene types ultimately becomes fixed in the population.

$F(a, b; c; z)$ is the standard hypergeometric function (see Abramowitz and Stegun, (1970), Chapter 15). The hypergeometric function is defined by the power series

$$F(a, b; c; z) = 1 + \frac{a \cdot b}{c \cdot 1} z + \frac{a(a+1) \cdot b(b+1)}{c(c+1) \cdot 1 \cdot 2} z^2$$
$$+ \frac{a(a+1)(a+2) \cdot b(b+1)(b+2)}{c(c+1)(c+2) \cdot 1 \cdot 2 \cdot 3} z^3 + \cdots \quad (51)$$

The hypergeometric function has an integral representation, which can be used for numerical evaluation and as an analytic continuation beyond $\|z\| < 1$,

$$F(a, b; c; z) = \frac{\Gamma(c)}{\Gamma(b)\Gamma(c-b)} \int_0^1 w^{b-1}(1-w)^{c-b-1}(1-wz)^{-a} \, dw, \quad (52)$$

where $Re(c) > Re(b) > 0$. The derivative of the hypergeometric function, needed for Itô's Lemma calculations, has the simple form

$$\frac{d}{dz} F(a, b; c; z) = \frac{ab}{c} F(a+1, b+1; c+1; z). \quad (53)$$

This formula can be derived by differentiating the terms of the power series in Equation (51) (see also Gradshteyn and Ryzhik, 2000, 9.100, 9.111). Though the hypergeometric function may be unfamiliar to many readers, it has appeared in a number of important asset-pricing contexts, including Merton (1973), Ingersoll (1977), Ingersoll and Ross (1992), Albanese et al. (2001), Longstaff (2005), and many others.

2.6 Asset returns

Given the explicit price function in Equation (50) and the functional form of its derivatives from Equation (53), the Appendix shows that the instantaneous return on the first asset R can be given by a direct application of Itô's Lemma:

$$R = \left[\delta + \mu_1 s + \mu_2(1 - s) + (\rho\sigma_1\sigma_2 - \sigma_2^2 + \eta^2 s) \, \Phi(s) \right] dt$$
$$+ \sigma_1[s + \Phi(s)] \, dZ_1 - \sigma_2[s - 1 + \Phi(s)] \, dZ_2, \quad (54)$$

where

$$\Phi(s) = \frac{A(s)}{B(s)},$$

$$A(s) = \frac{1}{1-\gamma} \left(\frac{s}{1-s} \right) F \left(1, \ 1-\gamma; \ 2-\gamma; \ \frac{s}{s-1} \right)$$
$$- \frac{1}{2-\gamma} \left(\frac{s}{1-s} \right)^2 F \left(2, \ 2-\gamma; 3-\gamma; \ \frac{s}{s-1} \right)$$
$$+ \frac{1}{1+\theta} \left(\frac{1-s}{s} \right) F \left(2, \ 1+\theta; \ 2+\theta; \ \frac{s-1}{s} \right),$$

$$B(s) = \frac{1}{1-\gamma}\left(\frac{s}{1-s}\right) F\left(1,\ 1-\gamma;\ 2-\gamma,\ \frac{s}{s-1}\right)$$
$$+ \frac{1}{\theta}\ F\left(1,\ \theta;\ 1+\theta;\ \frac{s-1}{s}\right).$$

Taking moments, we can now express how the expected return, expected excess return, and return volatility vary with the state variable s:

$$E_t[R] = \delta + \mu_1 s + \mu_2(1-s) + \left(\rho\sigma_1\sigma_2 - \sigma_2^2 + \eta^2 s\right)\Phi(s), \quad (55)$$

$$E_t[R] - r = \sigma_1^2 s^2 + \sigma_2^2(1-s)^2 + 2\rho\sigma_1\sigma_2 s(1-s)$$
$$+ \left(\rho\sigma_1\sigma_2 - \sigma_2^2 + \eta^2 s\right)\Phi(s), \quad (56)$$

$$\text{Var}_t[R] = \sigma_1^2[s + \Phi(s)]^2 + \sigma_2^2[s - 1 + \Phi(s)]^2$$
$$- 2\rho\sigma_1\sigma_2[s + \Phi(s)][s - 1 + \Phi(s)]. \quad (57)$$

2.7 Limits

In evaluating these formulas, it is useful to have exact expressions for limits as $s \to 0$. These limits also allow us to make some general statements about the behavior of "small" assets in this model. The Appendix shows that if $\theta \le 1$, we have

$$\lim_{s\to 0} V = \infty, \quad (58)$$

$$\lim_{s\to 0} \Phi(s) = 0, \quad (59)$$

and hence,

$$\lim_{s\to 0} E_t[R] = \delta + \mu_2 + \left(\rho\sigma_1\sigma_2 - \sigma_2^2\right)\theta, \quad (60)$$

$$\lim_{s\to 0} E_t[R] - r = \left(\rho\sigma_1\sigma_2 - \sigma_2^2\right)\theta + \sigma_2^2. \quad (61)$$

If $\theta > 1$, we instead have

$$\lim_{s\to 0} V = (\delta + \nu - \eta^2/2)^{-1}, \quad (62)$$

$$\lim_{s\to 0} \Phi(s) = 1, \quad (63)$$

and hence,

$$\lim_{s\to 0} E_t[R] = \delta + \mu_2 + \rho\sigma_1\sigma_2 - \sigma_2^2, \quad (64)$$

$$\lim_{s\to 0} E_t[R] - r = \rho\sigma_1\sigma_2. \quad (65)$$

The limit as $s \to 1$ is the one-tree model, of course.

The risk premium (expected excess return) of the first asset can be greater than zero even in the limit $s \to 0$, in two ways. First of all, it can obviously be greater than zero if the cash flows are correlated. If $\rho > 0$ in Equation (65),

there still is a risk premium, but one that comes entirely from the covariance of the first asset's dividend growth with the dividend growth of the second asset, which is now the market. Second, though, the limiting risk premium can be positive even with $\rho = 0$ per Equation (61). In this case, "valuation betas" are positive and lead to a positive risk premium even in the limit. However, the price-dividend limits show us that this result holds exactly when the price-dividend ratio, though finite for all nonzero s, tends to an infinite limit as $s \to 0$. The special case studied in the previous section has $\theta = 1$. In that case, the price-dividend ratio (just) tends to infinity, the risk premium of the first asset approaches zero at $s \to 0$, but its derivative with respect to s is infinite.

2.8 An example

Both as an example of the general case and for its own interest, we model one asset as a real perpetuity with $\mu_1 = 0$, $\sigma_1 = 0$. In an economy in which the second asset has $\mu_2 = 0.04$, $\sigma_2 = 0.20$, and $\delta = 0.04$, the top panel of Figure 9 presents the risk premium (expected excess return) of this perpetuity. In this case, all risk premium comes from "valuation betas" or discount-rate changes, because the dividend is constant. Nonetheless, there is interesting variation in the risk premium as a function of share, and the risk premium takes on both signs, as do bond risk premia in the data.

We can trace much of the behavior of the risk premium back to the valuation ratio, as usual. The bottom panel of Figure 9 presents the dividend yield— the inverse of the valuation ratio—along with the risk-free rate. Befitting a perpetuity, its dividend (coupon) yield moves with the risk-free rate. However, it does not move one-for-one, as the yield curve is not flat in this model. The region of positive risk premium corresponds to the region of rising dividend yield or declining price-dividend ratio. All returns here are due to shocks to the second asset's dividends. If that dividend increases, the share of the first asset decreases, raising the price and hence return of the first asset. Hence, the first asset return is positively correlated with consumption growth, and generates a positive risk premium. The converse logic holds in the region $s > 1/2$ with a rising price-dividend ratio, declining dividend-price ratio, and negative risk premium.

3. Concluding Remarks

We extend the classic single-asset Lucas-tree pure-exchange framework to the case of two assets and solve the model in closed form. Our two-tree model has the simplest ingredients, log utility, and i.i.d. normal dividend growth. Nonetheless, market-clearing logic and a fixed-share supply generate interesting and complex patterns of time-varying asset prices, expected returns, risk premia, variances, covariances, and correlations.

Expected Excess Return

Dividend Yield and Risk-Free Rate
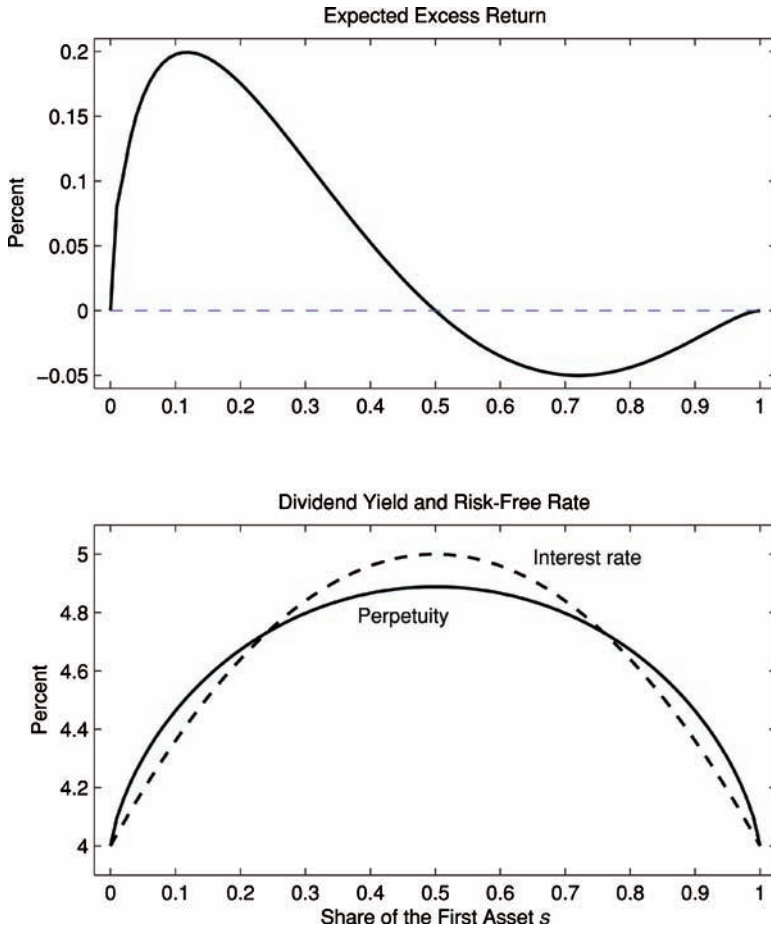
Figure 9
Expected excess returns and dividend yield for a perpetuity
The first asset is a perpetuity that pays a constant dividend stream. Its parameters are $\mu_1 = 0$, $\sigma_1 = 0$. The second, risky, asset has parameters $\mu_2 = 0.04$, $\sigma_2 = 0.20$. The discount factor is $\delta = 0.04$. The top panel displays the expected excess return of the perpetuity as a function of its dividend share. The bottom panel plots the dividend-price ratio of the perpetuity, along with the instantaneous interest rate.

3.1 Summary and intuition

With the results in hand, we can restate the intuition more clearly, emphasizing the quantitatively important channels. Start at the left-hand side of the plots, for small dividend shares. As the dividend share of the first tree increases from zero, the first tree becomes a larger part of the total, so its beta and risk premium naturally increase. Its expected excess return therefore rises from zero as shown in Figure 3. Also, as the first tree becomes a larger part of the total, that total becomes less risky by diversifying across the two trees. This change raises

interest rates by lowering the precautionary savings motive, again shown in Figure 3. The expected return is the sum of the expected excess return and the interest rate, and therefore rises even more steeply with share as shown in Figure 2. With a constant dividend-growth rate, higher expected returns mean lower price-dividend ratios, which is why the price-dividend ratio falls with the share in Figure 1.

The rise of the expected return with share and the decline of the price-dividend ratio with share underlie many of the dynamics we find. They also express the underlying market-clearing intuition. If there is a shock to the first tree's dividend, investors want to rebalance. Equivalently, they want to spread some of their larger wealth across both trees. They cannot collectively rebalance, so prices and expected returns must adjust. The expected return of the first tree must rise, making it more attractive to hold the larger share, and thus its valuation must fall. The expected return of the second tree falls (or, if you wish, the expected return of the first tree when there is a shock to the second tree) and its price rises. Investors want to buy more of the second tree but cannot, forcing its price to rise. Equivalently, the expected return of the second tree must fall so investors will hold it in its now smaller proportions.

However, overall betas that drive risk premia, as shown in Figure 6, depend not only on this "cash-flow beta" intuition, shown as $Cov[D, C]$ in Figure 3, but also on valuation betas, the tendency of the price-dividend ratio to rise or fall when the market and total consumption rise or fall. When the first tree's share is small, most increases in aggregate consumption come from increases in the second tree's dividend. Such an increase decreases the first tree's share, which increases the first tree's price-dividend ratio (Figure 1). Therefore, valuation betas are positive and large in the region of small shares, and they are a large component of risk premia. We see this in the large increase of expected excess returns in the left-hand side of Figure 3, and its decomposition into compensation for cash-flow risk labeled $Cov[D, C]$ and valuation risk labeled $Cov[V, C]$. The same nonlinearity is apparent in the beta of Figure 6. Figure 9 treats the case that the first tree is a perpetuity with a constant dividend stream. Here, the entire risk premium shown in the top panel derives from this valuation-beta mechanism.

Now let us move to the middle and right-hand parts of the graphs, where the first tree becomes a larger and larger share of the total. When the share reaches one, the first tree is the market, so it must have the market price-dividend ratio, as shown in Figure 1. However, there must be a region as shown in the right half of Figure 1 in which the first tree's price-dividend ratio is *less* than that of the market, because here the second tree's price-dividend ratio, with a small share, is *greater* than that of the market. Thus, the price-dividend ratio is not monotonic with the share, as shown in Figure 1. In terms of risk premia, this behavior is driven by interest rates (Figure 3) for our canonical example. The expected excess return is driven almost entirely by cash-flow betas at this

point, but interest rates fall with the share because the market becomes less diversified and more volatile as the share rises past 1/2. The lower interest rates eventually overcome the higher expected excess returns and cause a slight rise in the price-dividend ratio for shares above 0.9 (Figure 1) and a slight decline in the expected return (Figure 2).

The "valuation beta" is not always positive—it is possible that a rise in the market dividend lowers the price-dividend ratio of the first tree. Though quantitatively small, this effect can be seen in Figure 3 for shares between about 0.5 and 0.8. In this region, the first tree is more than half the market, so a rise in the market typically means a rise in the first tree's dividend and its share. The price-dividend ratio is still downward sloping in this region (Figure 1), so a rise in the share means a decline in the price-dividend ratio.

The possibility of a negative "valuation beta" is quantitatively more interesting in the bond-stock case of Figure 9. Here, all movements in total consumption come from movements in the second dividend. When the share is above one half, the dividend yield shown in the bottom of Figure 9 declines, so the U-shaped price-dividend ratio of the first tree rises. Thus, an increase in the market, which lowers the first tree's share, will also lower the first tree's price— a negative valuation beta. This negative beta generates the negative expected excess return shown in the right half of the top panel of Figure 9. Intuitively, when the second (small) tree's share rises under this circumstance it is still true that agents want to spread their increase in wealth across both trees, which should raise the price of the first tree. However, the interest rate also changes, because both mean and variance of consumption growth have changed, and this change more than offsets the rebalancing desire.

The remaining graphs draw out the dynamic implications of these effects. Because expected returns (Figure 2) and excess returns (Figure 3) vary with the share, and because the dividend yield (Figure 1) varies with the share, dividend yields forecast stock returns and excess returns as shown in Figure 5. Returns are cross-correlated despite no correlation in dividend growth, as shown in Figure 7. When the first tree's share is small, an increase in the second tree's dividend lowers the first tree's share, which (Figure 1) substantially raises its price-dividend ratio and thus gives a shock to the first tree's return, shown in the $[V_1, D_2]$ line of Figure 7. The symmetric effect generates a positive correlation when the first asset has a large share. Returns are correlated over time as shown in Figure 8. An increase in dividends increases today's return. This also increases the share, which increases expected returns (Figure 2) and excess returns (Figure 3), leading to positive autocorrelation. The small region in which expected returns (Figure 2) decline with share, driven by the decline in interest rates with share (Figure 3), generates a small, but theoretically interesting, region of negative return autocorrelation in Figure 8.

3.2 Discussion

A natural question is: Are the effects generated by the model quantitatively important and empirically relevant, and if not, what conclusions should one come to?

Our first answer is that this is the wrong question. Our aim is theory, to understand and characterize the dynamics induced by the market-clearing mechanism with the simplest possible preference and technology structure. Our aim is not to provide a calibrated model that replicates the full range of asset-pricing facts and puzzles. If the predictions of this model are counter to fact, one cannot conclude that we should ignore market-clearing dynamics. As a matter of logic, market-clearing dynamics will be present in any model that does not allow instantaneous aggregate rebalancing. If the model's predictions are false, one can only conclude that other ingredients are present in the real-world economy, which is obviously true.

That said, however, the realism or unrealism of the model's predictions should be addressed to some extent. If the model gives unrealistic predictions, it is worth speculating whether more general versions of market-clearing dynamics might address them, including more trees and higher risk aversion, or at what point preferences or technologies with dynamic elements will have to be part of the story.

The magnitudes are small. The relation between price-dividend ratios and subsequent returns in our model is about the same as that found in empirical time-series or cross-sectional (value-growth) return forecasts, but other effects such as the autocorrelation of returns are smaller in our model than many claims in the empirical literature.

However, we only use log utility, because we are not able to solve the model for higher risk aversion. It is natural to speculate that risk premia will be larger with higher risk aversion. Similarly, with log utility, our model is obviously inconsistent with the observed equity premium and low consumption volatility, but all the familiar equity-premium-boosting ingredients are likely to change that with two trees, as they do with one tree.

More seriously than magnitudes, some of the patterns in our model seem at odds with the data. For example, "small" firms are also "growth" firms in the simple parameterization of our model, with high valuations and low expected returns. Clearly, a "small firm effect" requires some other mechanism, such as a larger covariance of the underlying cash flows with the aggregate.

More generally, our model has only one state variable, the dividend share, which is both the only aggregate state variable and the only variable describing cross-sectional variation. Our model also has only two shocks. Taken literally to the data, it is as easy to reject these predictions as it is to reject the "prediction" that there are only two assets.

One deep question is whether market-clearing effects apply to individual assets, i.e., for very small share values, or whether they only apply to large

aggregates that are substantial shares of aggregate wealth. If the latter conclusion holds, then market-clearing dynamics may not be an important part of the story for many empirical findings that are concentrated in the very smallest of firms. One must recognize, however, that the question is not well posed. How can aggregates show effects that are not present in the individual constituents of those aggregates? If they matter for aggregates, ipso facto, in some sense they must matter for individual firms that compose those aggregates. Also, much empirical work on "individual firm" behavior in fact studies the behavior of portfolios of such firms, which constitute an aggregate with non-negligible share. Thus, the question is not easy to answer in the abstract. A full answer must await the analysis of an N-tree model, and must be specific about which fact one wants to address.

We address some of this issue in our study of the limits as $s \to 0$. In our simple model, though expected returns, risk premia, and autocorrelations do go to zero as $s \to 0$, the derivatives go to infinity at that point for many sets of parameters, so vanishingly small assets can have substantial risk premia and return dynamics induced by market clearing.

The dividend process and consequent share process in our model are obviously chosen for transparency, not realism. First of all, as in all endowment economies, it is unrealistic that there is no investment or share issue and repurchase at all. Our model is clearly best applied to thinking about dynamics that occur in the short run before share issuance/repurchase or investment and disinvestment can take place. For this reason, we do not think it too troubling that plots of price-dividend ratios or average returns versus shares in the data do not look like those of our model.[5] It seems likely that a model with, say, adjustment costs that slow down investment and disinvestment will generate market-clearing dynamics in the short run, meaning that *changes* in shares are associated with *changes* in valuations and average returns, while reverting in the long run to an i.i.d. economy in which there is no association between the *level* of the share and the *level* of valuations and average returns.

Finally, it may seem a little unsettling that the share in our model does not have a stationary distribution. This result is an inescapable implication of the geometric Brownian motion for dividends: one of two random walks will eventually end up dominating the other one. It is not obvious what the "right" assumption is here. In the end the car industry did dominate the horse buggy industry, so perhaps birth of new industries rather than mean reversion of dividends is the right way to generate long-run nondegenerate shares.

The question for us is: To what extent does the fact that the distribution of shares tends to two points (zero and one) have on the short-run asset pricing dynamics we have characterized, such as autocorrelation, predictability,

---

[5] We thank Ravi Bansal for pointing this out.

price-underreaction, and so forth? Is this model a good parable for what would happen with, say, a very long-run mean reversion in the shares? The basic asset pricing, Equation (21), reproduced here,

$$\frac{P}{C} = E_t \left[\int_0^\infty e^{-\delta\tau} s_{t+\tau} \, d\tau \right], \tag{66}$$

gives some comfort on this point. What matters for asset pricing is the conditional *mean* of the share, which always tends to 1/2. The fact that the *distribution* underlying that mean tends to two points, 0 and 1, does not have any effect on asset prices. We can also document that the share tends to its endpoints very slowly. For example, when the initial dividend share is 0.50, there is only about a 15% chance of the dividend share being below 0.05 or being above 0.95 after 100 years. Even when the initial share is 0.05, there is only about a 52% chance of the share being below 0.05 after 100 years, and less than a 1% chance of being above 0.95 after 100 years. Because the present value of cash flows beyond 100 years has only a negligible effect on asset values, these simulations indicate clearly that the asymptotic nonstationarity property of the dividend share is not driving the results. We have also repeated the analysis by simulation with cash flows truncated at 100 years, and find no differences worth reporting. As a result, we suspect that our core analysis would be very little affected if one changed to a dividend process with long-run mean reversion, though of course until the alternative model is solved one cannot say for sure.

Appendix A

A1 Derivation of asset prices, simplified case
Here we prove Equation (22) for the price-dividend ratio of the first asset in the simplified ($\delta = \sigma^2$) two-tree economy. From Equation (21), the price-consumption ratio Y for the first asset is

$$Y_t \equiv \frac{P_t}{C_t} = E_t \left[\int_0^\infty e^{-\delta\tau} s_{t+\tau} \, d\tau \right]. \tag{A1}$$

This ratio solves

$$E\left[dY\right] = (\delta Y - s) \, dt. \tag{A2}$$

From Equation (A1), Y depends only on the current value and future distribution of s. Thus, because s follows a Markov process (from Equation (4)), Y must be a function $Y(s)$ of the state variable. Using Itô's Lemma,

$$E\left[dY(s)\right] = Y'(s) \, E\left[ds\right] + \frac{1}{2} Y''(s) \, E\left[ds^2\right]. \tag{A3}$$

Putting together the last two equations, Y solves the differential equation

$$Y'(s) \, E\left[ds\right] + \frac{1}{2} Y''(s) \, E[ds^2] = (\delta Y - s) \, dt. \tag{A4}$$

Now we add the share process,

$$ds = s(1 - s)(1 - 2s)\sigma^2 dt + \sigma s(1 - s)(dZ_1 - dZ_2). \tag{A5}$$

Specializing to $\delta = \sigma^2$, the differential equation becomes

$$s(1 - s)(1 - 2s)Y'(s) + s^2(1 - s)^2 Y''(s) = Y(s) - \frac{s}{\delta}. \tag{A6}$$

We conjecture a solution and take derivatives:

$$Y(s) = \frac{1}{2\delta} \left(1 + \frac{1 - s}{s} \ln(1 - s) - \frac{s}{1 - s} \ln(s) \right), \tag{A7}$$

$$Y'(s) = -\frac{1}{2\delta} \frac{1}{s(1 - s)} \left(1 + \frac{1 - s}{s} \ln(1 - s) + \frac{s}{1 - s} \ln(s) \right), \tag{A8}$$

$$Y''(s) = -\frac{1}{\delta} \frac{1}{s^2(1 - s)^2} \left((2s - 1) - \frac{(1 - s)^2}{s} \ln(1 - s) + \frac{s^2}{1 - s} \ln(s) \right). \tag{A9}$$

Substituting into the differential equation and multiplying by 2δ,

$$0 = -(1 - 2s) \left(1 + \frac{1 - s}{s} \ln(1 - s) + \frac{s}{1 - s} \ln(s) \right)$$

$$- 2 \left((2s - 1) - \frac{(1 - s)^2}{s} \ln(1 - s) + \frac{s^2}{1 - s} \ln(s) \right)$$

$$- \left(1 + \frac{1 - s}{s} \ln(1 - s) - \frac{s}{1 - s} \ln(s) \right) + 2s. \tag{A10}$$

Grouping terms,

$$0 = -(1 - 2s) - 2(2s - 1) - 1 + 2s$$

$$- \left((1 - 2s)\frac{1 - s}{s} - 2\frac{(1 - s)^2}{s} + \frac{1 - s}{s} \right) \ln(1 - s)$$

$$- \left((1 - 2s)\frac{s}{1 - s} + 2\frac{s^2}{1 - s} - \frac{s}{1 - s} \right) \ln(s). \tag{A11}$$

Each term is zero, verifying the conjectured solution. Dividing the price-consumption ratio by s gives the price-dividend ratio in Equation (22).

A2 Derivation of asset prices, general case

The price-consumption ratio of the first asset is given by

$$\frac{P}{C} = E_t \left[\int_0^\infty e^{-\delta\tau} \frac{D_{t+\tau}}{C_{t+\tau}} d\tau \right] = E_t \left[\int_0^\infty e^{-\delta\tau} \frac{1}{1 + \frac{D_{2,t+\tau}}{D_{1,t+\tau}}} d\tau \right]$$

$$= E_t \left[\int_0^\infty e^{-\delta\tau} \frac{1}{1 + qe^u} d\tau \right], \tag{A12}$$

where q is the initial dividend ratio $D_{2,t}/D_{1,t}$ and u is a normally distributed random variate with mean $\nu\tau$ and variance $\eta^2\tau$, and where

$$\nu = \mu_2 - \mu_1 - \sigma_2^2/2 + \sigma_1^2/2,$$

$$\eta^2 = \sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2.$$

Note that $\nu\, dt = E[\ln(D_2/D_1)]$ and $\eta^2\, dt = \text{Var}[\ln(D_2/D_1)]$. Introducing the density for u into the last integral gives

$$\frac{P}{C} = \int_0^\infty \int_{-\infty}^\infty e^{-\delta\tau} \frac{1}{\sqrt{2\pi\eta^2\tau}} \frac{1}{1+qe^u} \exp\left(\frac{-(u-\nu\tau)^2}{2\eta^2\tau}\right)\, du\, d\tau. \tag{A13}$$

Interchanging the order of integration and collecting terms in τ gives

$$\frac{P}{C} = \int_{-\infty}^\infty \frac{1}{\sqrt{2\pi\eta^2}} \frac{1}{1+qe^u} \exp\left(\frac{\nu u}{\eta^2}\right)$$

$$\times \int_0^\infty \tau^{-1/2} \exp\left(-\frac{u^2}{2\eta^2}\frac{1}{\tau} - \frac{\nu^2+2\delta\eta^2}{2\eta^2}\tau\right)\, d\tau\, du. \tag{A14}$$

From Equation (3.471.9) of Gradshteyn and Ryzhik (2000), this expression becomes

$$\frac{P}{C} = \int_{-\infty}^\infty \frac{2}{\sqrt{2\pi\eta^2}} \frac{1}{1+qe^u} \exp\left(\frac{\nu u}{\eta^2}\right) \left(\frac{u^2}{\nu^2+2\delta\eta^2}\right)^{1/4}$$

$$\times K_{1/2}\left(2\sqrt{\frac{u^2(\nu^2+2\delta\eta^2)}{4\eta^4}}\right)\, du, \tag{A15}$$

where $K_{1/2}(\cdot)$ is the modified Bessel function of order $1/2$ (see Abramowitz and Stegun, 1970, ch. 9). From the identity relations for Bessel functions of order equal to an integer plus one half given in Gradshteyn and Ryzhik Equation (8.469.3), however, the above expression can be expressed as

$$\frac{P}{C} = \frac{1}{\psi} \int_{-\infty}^\infty \frac{1}{1+qe^u} \exp\left(\frac{\nu u}{\eta^2}\right) \exp\left(-\frac{\psi}{\eta^2}|u|\right)\, du, \tag{A16}$$

where

$$\psi = \sqrt{\nu^2+2\delta\eta^2}.$$

In turn, Equation (A16) can be written as

$$\frac{P}{C} = \frac{1}{\psi} \int_0^\infty \frac{1}{1+qe^u} \exp(\gamma u)\, du + \frac{1}{\psi} \int_{-\infty}^0 \frac{1}{1+qe^u} \exp(\theta u)\, du, \tag{A17}$$

where

$$\gamma = \frac{\nu - \psi}{\eta^2},$$

$$\theta = \frac{\nu + \psi}{\eta^2}.$$

Define $w = e^{-u}$. By a change of variables, Equation (A17) can be written as

$$\frac{P}{C} = \frac{1}{q\psi} \int_0^1 \frac{1}{1 + w/q} \, w^{-\gamma} \, dw + \frac{1}{\psi} \int_0^1 \frac{1}{1 + qw} \, w^{\theta - 1} \, dw. \tag{A18}$$

From Abramowitz and Stegun Equation (15.3.1), this expression becomes

$$\frac{P}{C} = \frac{1}{q\psi(1 - \gamma)} \, F(1, \; 1 - \gamma; \; 2 - \gamma; \; -1/q) + \frac{1}{\psi\theta} \, F(1, \; \theta; \; 1 + \theta; \; -q). \tag{A19}$$

Finally, substituting $q = (1 - s)/s$ into Equation (A19) and dividing by s gives the price-dividend ratio of the first asset,

$$V = \frac{1}{\psi(1 - \gamma)(1 - s)} \, F\left(1, \; 1 - \gamma, \; 2 - \gamma; \; \frac{s}{s - 1}\right)$$

$$+ \frac{1}{\psi\theta s} \, F\left(1, \; \theta; \; 1 + \theta; \; \frac{s - 1}{s}\right), \tag{A20}$$

which is Equation (50).

The special case solution for V given in Equation (22) can also be obtained directly from the general solution above. To see this, note that the parameter restrictions in the special case imply that $\theta = 1$ and $\gamma = -1$. Substituting these values into Equation (A20) results in hypergeometric functions of the form $F(1, 2; 3; \; \cdot \;)$ and $F(1, 1; 2; \; \cdot \;)$. A repeated application of the relations for contiguous hypergeometric functions described in Abramowitz and Stegun (1970) to $F(1, 2; 3; \; \cdot \;)$ allows V to be expressed entirely in terms of $F(1, 1; 2; \; \cdot \;)$. From Abramowitz and Stegun Equation (15.1.3), however, $F(1, 1; 2; z) = -\ln(1 - z)/z$. Substituting this into the expression for V leads immediately to Equation (22).

By symmetry, the price-dividend ratio of the second asset is given by

$$V_2 = \frac{1}{\psi(1 + \theta)s} \, F\left(1, \; 1 + \theta; \; 2 + \theta; \; \frac{s - 1}{s}\right)$$

$$- \frac{1}{\psi\gamma(1 - s)} \, F\left(1, \; -\gamma; \; 1 - \gamma; \; \frac{s}{s - 1}\right). \tag{A21}$$

Applying the recurrence relations for contiguous hypergeometric functions presented in Abramowitz and Stegun (1970), Equations (15.2.18) and (15.2.20) give the result

$$P_1 + P_2 = \frac{C}{\delta} = \frac{D_1 + D_2}{\delta} = P_M. \tag{A22}$$

To solve for the returns on the first asset, it is convenient to define $Y = P/C$ and to define an alternative state variable $x = D_1/D_2$. Note that $s = (1 + x)/x$. An application of Itô's Lemma gives

$$dx = \left(\mu_1 - \mu_2 + \sigma_2^2 - \rho\sigma_1\sigma_2\right) x \, dt + \sigma_1 \, x \, dZ_1 - \sigma_2 \, x \, dZ_2. \tag{A23}$$

Now applying Itô's Lemma to Y gives

$$dY = \left(\left(\mu_1 - \mu_2 + \sigma_2^2 - \rho\sigma_1\sigma_2\right) x \, Y_x + \left(\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2\right)x^2 Y_{xx}/2\right) dt$$

$$+ \sigma_1 \, x \, Y_x \, dZ_1 - \sigma_2 \, x \, Y_x \, dZ_2. \tag{A24}$$

Because $P = CY$, the above dynamics imply

$$\frac{dP}{P} = \left(\mu_1 s + \mu_2(1-s) + \left(\mu_1 - \mu_2 + \left(\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2 \right)s \right) xY_x/Y \right.$$

$$\left. + \left(\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2 \right) x^2 Y_{xx}/(2Y) \right) dt + \sigma_1 \left(xY_x/Y + s \right) dZ_1$$

$$- \sigma_2 \left(xY_x/Y + s - 1 \right) dZ_2. \tag{A25}$$

An argument similar to that in Section A1 of the Appendix can be used to show that Y satisfies the differential equation

$$\left(\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2 \right) x^2 Y_{xx}/2 = -\left(\mu_1 - \mu_2 - \rho\sigma_1\sigma_2 + \sigma_2^2 \right) xY_x$$

$$+ \delta Y - x/(1+x). \tag{A26}$$

Substituting out for Y_{xx} using the above expression allows us to rewrite Equation (A25) as

$$\frac{dP}{P} = \left(\delta + \mu_1 s + \mu_2(1-s) + \left(\rho\sigma_1\sigma_2 - \sigma_2^2 \right. \right.$$

$$\left. + \left(\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2 \right)s \right) xY_x/Y \right) dt + \sigma_1 \left(xY_x/Y + s \right) dZ_1$$

$$- \sigma_2 \left(xY_x/Y + s - 1 \right) dZ_2. \tag{A27}$$

Applying the expression for the derivative of the hypergeometric function repeatedly to the equation for Y given in Equation (A19), and then substituting out for x using $x = s/(1-s)$, shows that xY_X/Y is simply $\Phi(s)$ as given in Equation (54). Substituting $\Phi(s)$ into Equation (A27) and using the definition for η^2 gives the result.

Finally, while many numerical software programs calculate the hypergeometric function, it is very simple to evaluate the function by numerically integrating the integral representation in Equation (52). In doing so, however, it is important to use a numerical integration algorithm that provides robust results for functions with endpoint singularities. As one example of such an algorithm, see Piessens et al. (1983).

A3 Limits

In this section, we derive limits for price-dividend ratios as $s \to 0$ and $s \to 1$. Also, we derive limits for the function $\Phi(s)$ that figures prominently in the asset-price dynamics in Equation (54). We focus on the first asset, as the second is symmetric.

From the power series expression for the hypergeometric function, $F(a, b; c; 0) = 1$. Because of this result, it is useful to apply the linear transformation formula given in Abramowitz and Stegun Equation (15.3.7) so that the argument of the hypergeometric function goes to zero at the limit being evaluated:

$$F(a, b; c; z) = \frac{\Gamma(c)\Gamma(b-a)}{\Gamma(b)\Gamma(c-a)} (-z)^{-a} F(a, 1-c+a; 1-b+a; 1/z)$$

$$+ \frac{\Gamma(c)\Gamma(a-b)}{\Gamma(a)\Gamma(c-b)} (-z)^{-b} F(b, 1-c+b; 1-a+b; 1/z). \tag{A28}$$

To obtain the limit of the price-dividend ratio V as $s \to 0$, we use the linear transformation formula to rewrite Equation (A20) as

$$
V = \frac{1}{\psi(1-\gamma)} \left(\frac{1}{1-s} \right) F\left(1,\ 1-\gamma,\ 2-\gamma;\ \frac{s}{s-1} \right)
$$

$$
+ \frac{1}{\psi\theta} \left(\frac{\theta}{\theta-1} \right) \left(\frac{1}{1-s} \right) F\left(1,\ 1-\theta;\ 2-\theta;\ \frac{s}{s-1} \right)
$$

$$
+ \frac{1}{\psi\theta} \left(\frac{1}{s} \right) \Gamma(\theta+1)\, \Gamma(1-\theta) \left(\frac{s}{1-s} \right)^{\theta} F\left(\theta,\ 0;\ \theta;\ \frac{s}{s-1} \right).
$$

$$(A29)$$

From this expression, it is readily seen that

$$
\lim_{s \to 0} V =
\begin{cases}
\infty, & \text{if } \theta \le 1; \\[2mm]
\frac{1}{\delta + \nu - \eta^2/2}, & \text{if } \theta > 1.
\end{cases}
$$

$$(A30)$$

To obtain the limit of the price-dividend ratio as $s \to 1$, we again use the linear transformation formula and rewrite Equation (A20) as

$$
V = -\frac{1}{\psi(1-\gamma)} \left(\frac{1-\gamma}{\gamma} \right) \left(\frac{1}{s} \right) F\left(1,\ \gamma;\ 1-\gamma;\ \frac{s-1}{s} \right)
$$

$$
+ \frac{1}{\psi(1-\gamma)} \left(\frac{1}{1-s} \right) \Gamma(2-\gamma)\, \Gamma(\gamma) \left(\frac{1-s}{s} \right)^{1-\gamma}
$$

$$
\times F\left(1-\gamma,\ 0;\ 1-\gamma;\ \frac{s-1}{s} \right)
$$

$$
+ \frac{1}{\psi\theta} \left(\frac{1}{s} \right) F\left(1,\ \theta;\ 1+\theta;\ \frac{s-1}{s} \right).
$$

$$(A31)$$

From this, it immediately follows that

$$
\lim_{s \to 1} V = \frac{1}{\delta}.
$$

$$(A32)$$

A similar approach can be used to show that

$$
\lim_{s \to 0} V_2 = \frac{1}{\delta},
$$

$$(A33)$$

and that

$$
\lim_{s \to 1} V_2 =
\begin{cases}
\frac{1}{\delta - \nu - \eta^2/2}, & \text{if } \gamma < -1; \\[2mm]
\infty, & \text{if } \gamma \ge -1.
\end{cases}
$$

$$(A34)$$

Finally, the use of l'Hopital's rule and the repeated application of the linear transformation formula gives

$$
\lim_{s \to 0} \Phi(s) =
\begin{cases}
\theta, & \text{if } \theta \le 1; \\[2mm]
1, & \text{if } \theta > 1,
\end{cases}
$$

$$(A35)$$

and

$$\lim_{s \to 1} \Phi(s) = 0. \tag{A36}$$

Substituting the limiting values of $\Phi(s)$ into the asset-price dynamics in Equation (54) allows us to fully characterize the properties of these price dynamics as $s \to 0$ and $s \to 1$.

References

Abramowitz, M., and I. A. Stegun. 1970. *Handbook of Mathematical Functions*. New York: Dover Publications.

Albanese, C., G. Campolieti, P. Carr, and A. Lipton. 2001. Black Scholes Goes Hypergeometric. *Risk* 14:99–103.

Bansal, R., R. F. Dittman, and C. T. Lundblad. 2002. Consumption, Dividends, and the Cross Section of Stock Returns. Working Paper, Duke University.

Cox, J., J. E. Ingersoll, and S. A. Ross. 1985. An Intertemporal General Equilibrium Model of Asset Prices. *Econometrica* 53:363–84.

Crow, J. F., and M. Kimura. 1970. *An Introduction to Population Genetics Theory*. Minneapolis: Burgess.

Gala, V. D. 2006. Investment and Returns. Working Paper, Graduate School of Business, University of Chicago.

Gomes, J. F., L. Kogan, and L. Zhang. 2003. Equilibrium Cross-Section of Returns. *Journal of Political Economy* 111:693–732.

Gradshteyn, I. S., and I. M. Ryzhik. 2000. *Tables of Integrals, Series, and Products*, 6th ed. New York: Academic Press.

Guibaud, S., and N. Coeurdacier. 2006. A Dynamic Equilibrium Model of Imperfectly Integrated Financial Markets. Working Paper, Princeton University.

Hau, H., and H. Rey. 2004. Can Portfolio Rebalancing Explain the Dynamics of Equity Returns, Equity Flows, and Exchange Rates? *American Economic Review* 94:126–33.

Ingersoll, J. E. 1977. An Examination of Corporate Call Policies on Convertible Securities. *Journal of Finance* 32:463–78.

Ingersoll, J. E., and S. A. Ross. 1992. Waiting to Invest: Investment and Uncertainty. *Journal of Business* 65:1–29.

Karlin, S., and H. M. Taylor. 1981. *A Second Course in Stochastic Processes*. London: Academic Press.

Katsoulakis, M. A., and A. T. Kho. 2001. Stochastic Curvature Flows: Asymptotic Derivation, Level Set, and Formulation and Numerical Experiments. *Interfaces and Free Boundaries* 3:265–90.

Kloeden, P. E., and E. Platen. 1992. *Numerical Solution of Stochastic Differential Equations*. Berlin: Springer-Verlag.

Longstaff, F. A. 2005. Asset Pricing in Markets with Illiquid Assets. Working Paper, University of California, Los Angeles.

Longstaff, F. A., and M. Piazzesi. 2004. Corporate Earnings and the Equity Premium. *Journal of Financial Economics* 74:401–21.

Lucas, R. E., Jr. 1978. Asset Prices in an Exchange Economy. *Econometrica* 46:1429–45.

Menzly, L., T. Santos, and P. Veronesi. 2004. Understanding Predictability. *Journal of Political Economy* 112:1–47.

Merton, R. C. 1973. An Intertemporal Capital Asset Pricing Model. *Econometrica* 41:867–88.

Pavlova, A., and R. Rigobon. 2007. Asset Prices and Exchange Rates. *Review of Financial Studies* 20:1139–81.

Piessens, R., E. deDoncker-Kapenga, C. W. Uberhuber, and D. K. Kahaner. 1983. *QUADPACK*. New York: Springer-Verlag.

Santos, T., and P. Veronesi. 2006. Labor Income and Predictable Stock Returns. *Review of Financial Studies* 19:1–44.