



John H. Cochrane

President of the American Finance Association 2010

Presidential Address: Discount Rates

JOHN H. COCHRANE*

ABSTRACT

Discount-rate variation is the central organizing question of current asset-pricing research. I survey facts, theories, and applications. Previously, we thought returns were unpredictable, with variation in price-dividend ratios due to variation in expected cashflows. Now it seems all price-dividend variation corresponds to discount-rate variation. We also thought that the cross-section of expected returns came from the CAPM. Now we have a zoo of new factors. I categorize discount-rate theories based on central ingredients and data sources. Incorporating discount-rate variation affects finance applications, including portfolio theory, accounting, cost of capital, capital structure, compensation, and macroeconomics.

ASSET PRICES SHOULD EQUAL expected discounted cashflows. Forty years ago, Eugene Fama (1970) argued that the expected part, “testing market efficiency,” provided the framework for organizing asset-pricing research in that era. I argue that the “discounted” part better organizes our research today.

I start with facts: *how* discount rates vary over time and across assets. I turn to theory: *why* discount rates vary. I attempt a categorization based on central assumptions and links to data, analogous to Fama’s “weak,” “semi-strong,” and “strong” forms of efficiency. Finally, I point to some applications, which I think will be strongly influenced by our new understanding of discount rates. In each case, I have more questions than answers. This paper is more an agenda than a summary.

I. Time-Series Facts

A. Simple Dividend Yield Regression

Discount rates vary over *time*. (“Discount rate,” “risk premium,” and “expected return” are all the same thing here.) Start with a very simple regression of returns on dividend yields,¹ shown in Table I.

The 1-year regression forecast does not seem that important. Yes, the *t*-statistic is “significant,” but there are lots of biases and fishing. The 9% R^2 is not impressive.

*University of Chicago Booth School of Business, and NBER. I thank John Campbell, George Constantinides, Doug Diamond, Gene Fama, Zhiguo He, Bryan Kelly, Juhani Linnanmaa, Toby Moskowitz, Lubos Pastor, Monika Piazzesi, Amit Seru, Luis Viceira, Lu Zhang, and Guofu Zhou for very helpful comments. I gratefully acknowledge research support from CRSP and outstanding research assistance from Yoshio Nozawa.

¹Fama and French (1988).

Table I
Return-Forecasting Regressions

The regression equation is $R_{t \rightarrow t+k}^e = a + b \times D_t/P_t + \varepsilon_{t+k}$. The dependent variable $R_{t \rightarrow t+k}^e$ is the CRSP value-weighted return less the 3-month Treasury bill return. Data are annual, 1947–2009. The 5-year regression t -statistic uses the Hansen–Hodrick (1980) correction. $\sigma[E_t(R^e)]$ represents the standard deviation of the fitted value, $\sigma(\hat{b} \times D_t/P_t)$.

Horizon k	b	$t(b)$	R^2	$\sigma[E_t(R^e)]$	$\frac{\sigma[E_t(R^e)]}{E(R^e)}$
1 year	3.8	(2.6)	0.09	5.46	0.76
5 years	20.6	(3.4)	0.28	29.3	0.62

In fact, this regression has huge economic significance. First, the coefficient estimate is large. A one percentage point increase in dividend yield forecasts a nearly *four* percentage point higher return. Prices rise by an additional three percentage points.

Second, five and a half percentage point variation in expected returns is a lot. A 6% equity premium was already a “puzzle.”² The regression implies that expected returns *vary* by at least as much as their puzzling *level*, as shown in the last two columns of Table I.

By contrast, R^2 is a poor measure of economic significance in this context.³ The economic question is, “How much do expected returns vary over time?” There will always be lots of unforecastable return movement, so the variance of ex post returns is not a very informative comparison for this question.

Third, the slope coefficients and R^2 rise with horizon. Figure 1 plots each year’s dividend yield along with the subsequent 7 years of returns, in order to illustrate this point. Read the dividend yield as prices upside down: Prices were low in 1980 and high in 2000. The picture then captures the central fact: *High prices, relative to dividends, have reliably preceded many years of poor returns. Low prices have preceded high returns.*

B. Present Values, Volatility, Bubbles, and Long-Run Returns

Long horizons are most interesting because they tie predictability to volatility, “bubbles,” and the nature of price movements. I make this connection via the Campbell–Shiller (1988) approximate present value identity,

$$dp_t \approx \sum_{j=1}^k \rho^{j-1} r_{t+j} - \sum_{j=1}^k \rho^{j-1} \Delta d_{t+j} + \rho^k dp_{t+k}, \quad (1)$$

where $dp_t \equiv d_t - p_t = \log(D_t/P_t)$, $r_{t+1} \equiv \log R$, and $\rho \approx 0.96$ is a constant of approximation.

² Mehra and Prescott (1985).

³ Campbell (1991) makes this point, noting that a perpetuity would have very low short-run R^2 .




Figure 1. Dividend yield and following 7-year return. The dividend yield is multiplied by four. Both series use the CRSP value-weighted market index.

Now, consider regressions of weighted long-run returns and dividend growth on dividend yields:

$$\sum_{j=1}^k \rho^{j-1} r_{t+j} = a_r + b_r^{(k)} dp_t + \varepsilon_{t+k}^r, \quad (2)$$

$$\sum_{j=1}^k \rho^{j-1} \Delta d_{t+j} = a_d + b_d^{(k)} dp_t + \varepsilon_{t+k}^d, \quad (3)$$

$$dp_{t+k} = a_{dp} + b_{dp}^{(k)} dp_t + \varepsilon_{t+k}^{dp}. \quad (4)$$

The present value identity (1) implies that these long-run regression coefficients must add up to one,

$$1 \approx b_r^{(k)} - b_{\Delta d}^{(k)} + \rho^k b_{dp}^{(k)}. \quad (5)$$

To derive this relation, regress both sides of the identity (1) on dp_t .

Equations (1) and (5) have an important message. If we lived in an i.i.d. world, dividend yields would never vary in the first place. Expected future

Table II
Long-Run Regression Coefficients

Table entries are long-run regression coefficients, for example, $b_r^{(k)}$ in $\sum_{j=1}^k \rho^{j-1} r_{t+j} = a + b_r^{(k)} dp_t + \varepsilon_{t+k}^r$. See equations (2)–(4). Annual CRSP data, 1947–2009. “Direct” regression estimates are calculated using 15-year ex post returns, dividend growth, and dividend yields as left-hand variables. The “VAR” estimates infer long-run coefficients from 1-year coefficients, using estimates in the right-hand panel of Table III. See the Appendix for details.

Method and Horizon	Coefficient		
	$b_r^{(k)}$	$b_{\Delta d}^{(k)}$	$\rho^k b_{dp}^{(k)}$
Direct regression, $k = 15$	1.01	−0.11	−0.11
Implied by VAR, $k = 15$	1.05	0.27	0.22
VAR, $k = \infty$	1.35	0.35	0.00

returns and dividend growth would never change. Since dividend yields vary, they must forecast long-run returns, long-run dividend growth, or a “rational bubble” of ever-higher prices.

The regression coefficients in (5) can be read as the fractions of dividend yield variation attributed to each source. To see this interpretation more clearly, multiply both sides of (5) by $\text{var}(dp_t)$, which gives

$$\text{var}(dp_t) \approx \text{cov} \left[dp_t, \sum_{j=1}^k \rho^{j-1} r_{t+j} \right] - \text{cov} \left[dp_t, \sum_{j=1}^k \rho^{j-1} \Delta d_{t+j} \right] + \rho^k \text{cov}(dp_t, dp_{t+k}). \quad (6)$$

The empirical question is, how big is each source of variation? Table II presents long-run regression coefficients, each calculated three ways.

The long-run return coefficients, shown in the first column, are all a bit larger than 1.0. The dividend growth forecasts, in the second column, are small, statistically insignificant, and the positive point estimates go the “wrong” way—high prices relative to current dividends signal *low* future dividend growth. The 15-year dividend yield forecast coefficient is also essentially zero.

Thus, the estimates summarized in Table II say that *all price-dividend ratio volatility corresponds to variation in expected returns. None corresponds to variation in expected dividend growth, and none to “rational bubbles.”*

In the 1970s, we would have guessed exactly the opposite pattern. Based on the idea that returns are not predictable, we would have supposed that high prices relative to current dividends reflect expectations that dividends will rise in the future, and so forecast higher dividend growth. That pattern is completely absent. Instead, high prices relative to current dividends entirely forecast low returns.

*This is the true meaning of return forecastability.*⁴ This is the real measure of “how big” the point estimates are—return forecastability is “just enough”

⁴ Shiller (1981), Campbell and Shiller (1988), Campbell and Ammer (1993), Cochrane (1991a, 1992, 1994, 2005b).

to account for price volatility. This is the natural set of units with which to evaluate return forecastability. What we expected to be zero is one; what we expected to be one is zero.

Table II also reminds us that the point of the return-forecasting project is to understand *prices*, the *right-hand* variable of the regression. We put return on the left side because the forecast error is uncorrelated with the forecasting variable. This choice does not reflect “cause” and “effect,” nor does it imply that the point of the exercise is to understand ex post return variation.

How you look at things matters. The long-run and short-run regressions are equivalent, as each can be obtained from the other. Yet looking at the long-run version of the regressions shows an unexpected economic significance. We will see this lesson repeated many times.

Some quibbles: Table II does not include standard errors. Sampling variation in long-run estimates is an important topic.⁵ My point is the *economic* importance of the estimates. One might still argue that we cannot reject the alternative views. But when point estimates are one and zero, arguing we should believe zero and one because zero and one cannot be rejected is a tough sell.

The variance of dividend yields or price-dividend ratios corresponds entirely to discount-rate variation, but as much as half of the variance of *price changes* $\Delta p_{t+1} = -dp_{t+1} + dp_t + \Delta d_{t+1}$ or *returns* $r_{t+1} \approx -\rho dp_{t+1} + dp_t + \Delta d_{t+1}$ corresponds to current dividends Δd_{t+1} . This fact seems trivial but has caused a lot of confusion.

I divide by dividends for simplicity, to capture a huge literature in one example. Many other variables work about as well, including earnings and book values.

C. A Pervasive Phenomenon

This pattern of predictability is pervasive across markets. For stocks, bonds, credit spreads, foreign exchange, sovereign debt, and houses, a yield or valuation ratio translates one-for-one to expected excess returns, and does not forecast the cashflow or price change we may have expected. In each case our view of the facts has changed completely since the 1970s.

- *Stocks*. Dividend yields forecast returns, not dividend growth.⁶
- *Treasuries*. A rising yield curve signals better 1-year returns for long-term bonds, not higher future interest rates. Fed fund futures signal returns, not changes in the funds rate.⁷
- *Bonds*. Much variation in credit spreads over time and across firms or categories signals returns, not default probabilities.⁸
- *Foreign exchange*. International interest rate spreads signal returns, not exchange rate depreciation.⁹

⁵ Cochrane (2006) includes many references.

⁶ Fama and French (1988, 1989).

⁷ Fama and Bliss (1987), Campbell and Shiller (1991), Piazzesi and Swanson (2008).

⁸ Fama (1986), Duffie and Berndt (2011).

⁹ Hansen and Hodrick (1980), Fama (1984).




Figure 2. House prices and rents. OFHEO is the Office of Federal Housing Enterprise Oversight “purchase-only” price index. CSW are Case-Shiller-Weiss price data. All data are from <http://www.lincolninst.edu/subcenters/land-values/rent-price-ratio.asp>.

- *Sovereign debt.* High levels of sovereign or foreign debt signal low returns, not higher government or trade surpluses.¹⁰
- *Houses.* High price/rent ratios signal low returns, not rising rents or prices that rise forever.

Since house prices are so much in the news, Figure 2 shows house prices and rents, and Table III presents forecasting regressions. High prices relative to rents mean low returns, not higher subsequent rents, or prices that rise forever. The housing regressions are almost the same as the stock market regressions. (Not everything about house and stock data is the same of course. Measured house price data are more serially correlated.)

There is a strong common element and a strong business cycle association to all these forecasts.¹¹ Low prices and high expected returns hold in “bad times,” when consumption, output, and investment are low, unemployment is high, and businesses are failing, and vice versa.

These *facts* bring a good deal of structure to the debate over “bubbles” and “excess volatility.” High valuations correspond to low returns, and are

¹⁰ Gourinchas and Rey (2007).

¹¹ Fama and French (1989).

Table III
House Price and Stock Price Regressions

Left panel: Regressions of log annual housing returns r_{t+1} , log rent growth Δd_{t+1} , and log rent/price ratio dp_{t+1} on the rent/price ratio dp_t , $x_{t+1} = a + b \times dp_t + \varepsilon_{t+1}$ 1960–2010. Right panel: Regressions of log stock returns r_{t+1} , dividend growth Δd_{t+1} and dividend yields dp_{t+1} on dividend yields dp_t , annual CRSP value-weighted return data, 1947–2010.

	Houses			Stocks		
	b	t	R^2	b	t	R^2
r_{t+1}	0.12	(2.52)	0.15	0.13	(2.61)	0.10
Δd_{t+1}	0.03	(2.22)	0.07	0.04	(0.92)	0.02
dp_{t+1}	0.90	(16.2)	0.90	0.94	(23.8)	0.91

associated with good economic conditions. All a “price bubble” can possibly mean now is that the equivalent discount rate is “too low” relative to some theory. Though regressions do not establish causality, this equivalence guides us to a much more profitable discussion.

D. The Multivariate Challenge

This empirical project has only begun. We see that one variable at a time forecasts one return at a time. We need to understand their *multivariate* counterparts, on both the left and the right sides of the regressions.

For example, the stock and bond regressions on dividend yield and yield spread (ys) are

$$\begin{aligned} r_{t+1}^{\text{stock}} &= a_s + b_s \times dp_t + \varepsilon_{t+1}^s, \\ r_{t+1}^{\text{bond}} &= a_b + c_b \times ys_t + \varepsilon_{t+1}^b. \end{aligned}$$

We have some additional predictor variables z_t , from similar univariate or at best bivariate (i.e., including $b_s \times dp_t$) explorations:

$$r_{t+1}^{\text{stock}} = a_s + b_s \times dp_t + d_s \times z_t + \varepsilon_{t+1}^s.$$

First, which of these variables are really important in a multiple regression sense? In particular, do the variables that forecast one return forecast another? What are c_s , d_s , b_b , and d_b in regressions

$$\begin{aligned} r_{t+1}^{\text{stock}} &= a_s + b_s \times dp_t + \underline{c_s} \times \underline{ys_t} + \underline{d'_s} z_t + \varepsilon_{t+1}^s, \\ r_{t+1}^{\text{bond}} &= a_b + \underline{b_b} \times \underline{dp_t} + c_b \times ys_t + \underline{d'_b} z_t + \varepsilon_{t+1}^b? \end{aligned} \tag{7}$$

(I underline the variables we need to learn about.)

Second, how correlated are the right-hand terms of these regressions? What is the *factor structure of time-varying expected returns*? Expected returns $E_t(r_{t+1}^i)$ vary over time t . How correlated is such variation across assets and asset classes i ? How can we best express that correlation as factor structure?

As an example to clarify the question, suppose we find that the stock return coefficients are all double those of the bonds,

$$\begin{aligned} r_{t+1}^{\text{stock}} &= a_s + 2 \times dp_t + 4 \times ys_t + \varepsilon_{t+1}^s, \\ r_{t+1}^{\text{bond}} &= a_b + 1 \times dp_t + 2 \times ys_t + \varepsilon_{t+1}^b. \end{aligned}$$

We would see a one-factor model for expected returns, with stock expected returns always changing by twice bond expected returns,

$$\begin{aligned} E_t(r_{t+1}^{\text{stock}}) &= 2 \times \text{factor}_t, \\ E_t(r_{t+1}^{\text{bond}}) &= 1 \times \text{factor}_t. \end{aligned} \tag{8}$$

Third, what are the corresponding pricing factors? We relate time-varying expected returns to covariances with pricing factors or portfolio returns,

$$E_t(r_{t+1}^i) = \text{cov}_t(r_{t+1}^i \mathbf{f}_{t+1}') \boldsymbol{\lambda}_t.$$

As a small step down this road, Cochrane and Piazzesi (2005, 2008) find that forward rates of *all* maturities help to forecast bond returns of *each* maturity. Multiple regressions matter as in (7). Furthermore, the right-hand sides are almost perfectly correlated across left-hand maturities.¹² A single common factor describes 99.9% of the variance of *expected* returns as in (8). Finally, the spread in time-varying expected bond returns across maturities corresponds to a spread in covariances with a single “level” factor. The market prices of slope, curvature, and expected-return factor risks are zero.

What similar patterns hold across broad asset classes? The challenge, of course, is that there are too many right-hand variables, so we cannot simply run huge multiple regressions. But these are the vital *questions*.

E. Multivariate Prices

I advertised that much of the point of running regressions with prices on the right-hand side is to understand those prices. How will a multivariate investigation change our picture of *prices* and *long-run* returns?

Again, the Campbell–Shiller present value identity

$$dp_t \approx \sum_{j=1}^{\infty} \rho^{j-1} r_{t+j} - \sum_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j} \tag{9}$$

provides a useful way to think about these questions. Since this identity holds *ex post*, it holds for any information set. Dividend yields are a great forecasting variable because they reveal market expectations of dividend growth and returns. However, dividend yields combine the two sources of information. A variable can help the dividend yield to forecast long-run returns if it *also* forecasts long-run dividend growth. A variable can also help predict 1-year returns

¹² Hansen and Hodrick (1980) and Stambaugh (1988) find similar structures.

Table IV
Forecasting Regressions with the Consumption-Wealth Ratio

Annual data 1952–2009. Long-run coefficients in the last two rows of the table are computed using a first-order VAR with dp_t and cay_t as state variables. Each regression includes a constant. Cay is rescaled so $\sigma(cay) = 1$. For reference, $\sigma(dp) = 0.42$.

Left-Hand Variable	Coefficients		<i>t</i> -Statistics		Other Statistics		
	dp_t	cay_t	dp_t	cay_t	R^2	$\sigma[E_t(y_{t+1})]\%$	$\frac{\sigma[E_t(y_{t+1})]}{E(y_{t+1})}$
r_{t+1}	0.12	0.071	(2.14)	(3.19)	0.26	8.99	0.91
Δd_{t+1}	0.024	0.025	(0.46)	(1.69)	0.05	2.80	0.12
dp_{t+1}	0.94	−0.047	(20.4)	(−3.05)	0.91		
cay_{t+1}	0.15	0.65	(0.63)	(5.95)	0.43		
$r_t^{lr} = \sum_{j=1}^{\infty} \rho^{j-1} r_{t+j}$	1.29	0.033				0.51	
$\Delta d_t^{lr} = \sum_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j}$	0.29	0.033				0.12	

r_{t+1} without much changing long-run expected returns, if it has an offsetting effect on longer run returns $\{r_{t+j}\}$. Such a variable signals a change in the *term structure of risk premia* $\{E_t r_{t+j}\}$.

I examine Lettau and Ludvigson's (2001a, 2001b, 2005) consumption to wealth ratio cay as an example to explore these questions. Table IV presents forecasting regressions.

Cay helps to forecast one-period returns. The *t*-statistic is large, and it raises the variation of expected returns substantially. Cay only marginally helps to forecast dividend growth. (Lettau and Ludvigson report that it works better in quarterly data.)

Figure 3 graphs the 1-year return forecast using dp alone, the 1-year return forecast using dp and cay together, and the actual ex post return. Adding cay lets us forecast business-cycle frequency “wiggles” while not affecting the “trend.”

Long-run return forecasts are quite different, however. Figure 4 contrasts long-run return forecasts with and without cay . Though cay has a dramatic effect on *one-period* return r_{t+1} forecasts in Figure 3, cay has almost no effect at all on *long-run* return $\sum_{j=1}^{\infty} \rho^{j-1} r_{t+j}$ forecasts in Figure 4.

Figure 4 includes the actual dividend yield, to show (by (9)) how dividend yields break into long-run return forecasts versus long-run dividend growth forecasts. The last two rows of Table IV give the corresponding long-run regression coefficients. Essentially all price-dividend variation still corresponds to expected-return forecasts.

How can cay forecast one-year returns so strongly, but have such a small effect on the terms of the dividend yield present value identity? In the context of (9), cay alters the *term structure* of expected returns.

We can display this behavior with impulse-response functions. Figure 5 plots responses to a dividend growth shock, a dividend yield shock, and a cay shock. In each case, I include a contemporaneous return response to satisfy the return identity $r_{t+1} = \Delta d_{t+1} - \rho dp_{t+1} + dp_t$.




Figure 3. Forecast and actual 1-year returns. The forecasts are fitted values of regressions of returns on dividend yield and *cay*. Actual returns r_{t+1} are plotted on the same date as their forecast, $a + b \times dp_t$.




Figure 4. Log dividend yield dp and forecasts of long-run returns $\sum_{j=1}^{\infty} \rho^{j-1} r_{t+j}$. Return forecasts are computed from a VAR including dp , and a VAR including dp and *cay*.




Figure 5. Impulse-response functions. Response functions to dividend growth, dividend yield, and cay shocks. Calculations are based on the VAR of Table IV. Each shock changes the indicated variable without changing the others, and includes a contemporaneous return shock from the identity $r_{t+1} = \Delta d_{t+1} - \rho dp_{t+1} + dp_t$. The vertical dashed line indicates the period of the shock.

These plots answer the question: “What change in expectations corresponds to each shock?” The dividend growth shock corresponds to permanently higher expected dividends with no change in expected returns. Prices jump to their new higher value and stay there. It is thus a pure “expected cashflow” shock. The dividend yield shock is essentially a pure discount-rate shock. It shows a rise in expected returns with little change in expected dividend growth.

Though there is a completely transitory component of prices in this multivariate representation, the implied *univariate* return representation remains very close to uncorrelated. A fall in prices with no change in dividends is likely to mean-revert, but observing a fall in prices without observing dividends carries no such implication. As a result, stocks are not “safer in the long run”: Stock return variance still scales nearly linearly with horizon.

The *cay* shock in the rightmost panel of Figure 5 corresponds to a shift in expected returns from the distant future to the near future, with a small similar movement in the timing of a dividend growth forecast. It has almost no effect

on long-run returns or dividend growth. We could label it a *shock to the term structure of risk premia*.¹³

So, *cay* strongly forecasts 1-year returns, but has little effect on price-dividend ratio variance attribution. Does this pattern hold for other return forecasters? I don't know. In principle, consistent with the identity (9), other variables can help dividend yields to predict *both* long-run returns and long-run dividend growth. Consumption and dividends should be cointegrated, and since dividends are so much more volatile, the consumption-dividend ratio should forecast long-run dividend growth. Cyclical variables should work: At the bottom of a recession, both discount rates and expected growth rates are likely to be high, with offsetting effects on dividend yields. Reflecting both ideas, Lettau and Ludvigson (2005) report that "*cdy*," a cointegrating vector including dividends, forecasts long-run dividend growth in just this way. However, the lesser persistence of typical forecasters will work against their having much of an effect on price-dividend ratios. *Cay*'s coefficient of only 0.65 on its own lag, and the fact that *cay* does not forecast dividend yields in my regressions, are much of the story for *cay*'s failure to affect long-run forecasts.

Even so, additional variables can only *raise* the contribution of long-run expected returns to price-dividend variation. Additional variables do not shift variance attribution from returns to dividends. A higher long-run dividend forecast must be matched by a higher long-run return forecast if it is not to affect the dividend yield.

This is a suggestive first step, not an answer. We have a smorgasbord of return forecasters to investigate, singly and jointly, including information in additional lags of returns and dividend yields (see the Appendix). The point is this: Multivariate long-run forecasts and consequent price implications can be quite different from one-period return forecasts. As we pursue the multivariate forecasting question using the large number of additional forecasting variables, we should look at pricing implications, and not just run short-run R^2 contests.

II. The Cross-Section

In the beginning, there was chaos. Practitioners thought that one only needed to be clever to earn high returns. Then came the CAPM. Every clever strategy to deliver high average returns ended up delivering high market betas as well. Then anomalies erupted, and there was chaos again. The "value effect" was the most prominent anomaly.

Figure 6 presents the Fama–French 10 book-to-market sorted portfolios. Average excess returns rise from growth (low book-to-market, "high price") to value (high book-to-market, "low price"). This fact would not be a puzzle if the betas also rose. But the betas are about the same for all portfolios.

The fact that betas do not rise with value is really the heart of the puzzle. It is natural that stocks, which have fallen on hard times, should have higher

¹³ For impulse-responses, see Cochrane (1994). For the effect of *cay*, see Lettau and Ludvigson (2005).




Figure 6. Average returns and betas. 10 Fama–French book-to-market portfolios. Monthly data, 1963–2010.

subsequent returns. If the market declines, these stocks should be particularly hard hit. They should have higher average returns—and higher betas. All puzzles are *joint* puzzles of expected returns and betas. Beta without expected return is just as much a puzzle—and as profitable—as expected return without beta.¹⁴

Fama and French (1993, 1996) brought order once again with size and value factors. Figure 6 includes the results of multiple regressions on the market excess return and Fama and French's *hml* factor,

$$R_t^{ei} = \alpha_i + b_i \times rmrf_t + h_i \times hml_t + \varepsilon_{it}.$$

The figure shows the separate contributions of $b_i \times E(rmrf)$ and $h_i \times E(hml)$ in accounting for average returns $E(R^{ei})$. Higher average returns *do* line up well with larger values of the h_i regression coefficient.

Fama and French's factor model accomplishes a very useful data reduction. Theories now only have to explain the *hml* portfolio premium, not the expected

¹⁴ Frazzini and Pedersen (2010).

returns of individual assets.¹⁵ This lesson has yet to sink in to a lot of empirical work, which still uses the 25 Fama–French portfolios to test deeper models.

Covariance is in a sense Fama and French’s central result: If the value firms decline, they all decline together. This is a sensible result: Where there is mean, there must be comovement, so that Sharpe ratios do not rise without limit in well-diversified value portfolios.¹⁶ But theories now must also explain this *common movement* among value stocks. It is not enough to simply generate temporary price movements in individual securities, “fads” that produce high or low prices, and then fade away, rewarding contrarians. All the securities with low prices today must *rise and fall together* in the future.

Finally, Fama and French found that *other* sorting variables, such as firm sales growth, did not each require a new factor. The three-factor model took the place of the CAPM for routine risk adjustment in empirical work.

Order to chaos, yes, but once again, the world changed completely. None of the cross-section of average stock returns corresponds to market betas. All of it corresponds to *hml* and size betas.

Alas, the world is once again descending into chaos. Expected return strategies have emerged that *do not* correspond to market, value, and size betas. These include, among many others, momentum,¹⁷ accruals, equity issues and other accounting-related sorts,¹⁸ beta arbitrage, credit risk, bond and equity market-timing strategies, foreign exchange carry trade, put option writing, and various forms of “liquidity provision.”

A. The Multidimensional Challenge

We are going to have to repeat Fama and French’s anomaly digestion, but with many more dimensions. We have a lot of questions to answer:

First, which characteristics really provide *independent* information about average returns? Which are subsumed by others?

Second, does each new anomaly variable also correspond to a new factor formed on those same anomalies? Momentum returns correspond to regression coefficients on a winner–loser momentum “factor.” Carry-trade profits correspond to a carry-trade factor.¹⁹ Do accruals return strategies correspond to an accruals factor? We should routinely look.

Third, how many of these new factors are really important? Can we again account for N independent dimensions of expected returns with $K < N$ factor exposures? Can we account for accruals return strategies by betas on some other factor, as with sales growth?

¹⁵ Daniel and Titman (2006), Lewellen, Nagel, and Shanken (2010).

¹⁶ Ross (1976, 1978).

¹⁷ Jegadeesh and Titman (1993).

¹⁸ See Fama and French (2010).

¹⁹ Lustig, Roussanov, and Verdelhan (2010a, 2010b).




Figure 7. Portfolio means versus cross-sectional regressions.

Now, factor structure is neither necessary nor sufficient for factor pricing. ICAPM and consumption-CAPM models do not predict or require that pricing factors correspond to big common movements in asset returns. And big common movements, such as industry portfolios, need not correspond to any risk premium. There always is an equivalent single-factor pricing representation of any multifactor model: The mean-variance efficient portfolio return is the single factor. Still, the world would be much simpler if betas on only a few factors, important in the covariance matrix of returns, accounted for a larger number of mean characteristics.

Fourth, eventually, we have to connect all this back to the central question of finance: Why do *prices* move?

B. Asset Pricing as a Function of Characteristics

To address these questions in the zoo of new variables, I suspect we will have to use different *methods*. Following Fama and French, a standard methodology has developed: Sort assets into portfolios based on a characteristic, look at the portfolio means (especially the 1–10 portfolio alpha, information ratio, and *t*-statistic), and then see if the spread in means corresponds to a spread of portfolio betas against some factor. But we cannot do this with 27 variables.

Portfolio sorts are really the same thing as nonparametric cross-sectional regressions, using nonoverlapping histogram weights. Figure 7 illustrates the point.

For one variable, portfolio sorts and regressions both work. But we cannot chop portfolios 27 ways, so I think we will end up running multivariate regressions.²⁰ The Appendix presents a simple cross-sectional regression to illustrate the idea.

²⁰ Fama and French (2010) already run such regressions, despite reservations over functional forms.

More generally, “time-series” forecasting regressions, “cross-sectional” regressions, and portfolio mean returns are really the same thing. All we are ever really doing is understanding a big panel-data forecasting regression,

$$R_{t+1}^e = a + b'C_{it} + \varepsilon_{t+1}^i.$$

We end up describing expected returns as a function of characteristics,

$$E(R_{t+1}^e | C_t),$$

where C_t denotes some big vector of characteristics,

$$C_t = [\text{size, bm, momentum, accruals, dp, credit spread} \dots].$$

Is value a “time-series” strategy that moves in and out of a stock as that stock’s book-to-market ratio changes, or is it a “cross-sectional” strategy that moves from one stock to another following book-to-market signals? Well, both, obviously. They are the same thing. This is the managed-portfolio theorem:²¹ An instrument z_t in a time-series test $0 = E[(m_{t+1}R_{t+1}^e)z_t]$ is the same thing as a managed-portfolio return $R_{t+1}^e z_t$ in an unconditional test $0 = E[m_{t+1}(R_{t+1}^e z_t)]$.

Once we understand expected returns, we have to see if expected returns line up with covariances of returns with factors. Sorted-portfolio betas are a nonparametric estimate of this covariance function,

$$\text{cov}_t(R_{t+1}^e, f_{t+1}) = g(C_{it}).$$

Parametric approaches are natural here as well, to address a multidimensional world. For example, we can run regressions such as

$$[R_{t+1}^e - E(R_{t+1}^e | C_{it})] f_{t+1} = c + d'C_{it} + \varepsilon_{t+1}^i \Rightarrow g(C) = c + d'C.$$

(The errors may not be normal, but they are mean-zero and uncorrelated with the right-hand variable.)

We want to see if the mean return *function* lines up with the covariance *function*: Is it true that

$$E(R^e | C) = g(C) \times \lambda?$$

An implicit assumption underlies everything we do: Expected returns, variances, and covariances are stable functions of characteristics such as size and book-to-market ratio, and not security names. This assumption is why we use portfolios in the first place. Without this assumption, it is hard to tell if there is any spread in average returns at all. It means that asset pricing really is about the equality of two *functions*: The function relating means to characteristics should be proportional to the function relating covariance to characteristics.

Looking at portfolio average returns rather than forecasting regressions was really the key to understanding the *economic* importance of many effects, as was looking at long-horizon returns. For example, serial correlation with an R^2 of 0.01 does not seem that impressive. Yet it is enough to account for momentum:

²¹ Cochrane (2005b).

The last year's winners went up 100%, so an annual autocorrelation of 0.1, meaning $0.01 R^2$, generates a 10% annual portfolio mean return. (An even smaller amount of time-series cross correlation works as well.) As another classic example, Lustig, Roussanov, and Verdelhan (2010a) translate carry-trade return-forecasting regressions to means of portfolios formed on the basis of currency interest differentials. This step leads them to look for and find a factor structure of country returns that depends on interest differentials, a "high minus low" factor. This step follows Fama and French (1996) exactly, but no one thought to look for it in 30 years of running country-by-country time-series forecasting regressions.

The equivalence of portfolio sorts and regressions goes both ways. We can still calculate these measures of *economic* significance if we *estimate* panel-data regressions for means and covariances. From the spread of lagged returns and return autocorrelation, we can calculate the momentum-portfolio implications directly. The 1–10 portfolio information ratio is the same thing as the Sharpe ratio of the underlying factor, or t -statistic of the cross-sectional regression coefficient. (See the Appendix.) We can study the covariance structure of panel-data regression residuals as a function of the same characteristics rather than actually form portfolios,

$$\text{cov}_t(R_{t+1}^i, R_{t+1}^j) = h(C_{it}, C_{jt}).$$

Running multiple panel-data forecasting regressions is full of pitfalls of course. One can end up focusing on tiny firms, or outliers. One can get the functional form wrong.

However, uniting time series and cross-section will yield new insights as well. For example, variation in book-to-market over time for a given portfolio has a larger effect on returns than variation in book-to-market across the Fama–French portfolios, and a recent change in book-to-market also seems to forecast returns. (See the Appendix.)

I did not say it will be easy! But we must address the factor zoo, and I do not see how to do it by a high-dimensional portfolio sort.

C. Prices

Then, we have to answer the central question, what is the source of *price* variation?

When did our field stop being "asset pricing" and become "asset expected returning?" Why are betas exogenous?²² A lot of price variation comes from discount-factor news. What sense does it make to "explain" expected returns by the covariation of expected return shocks with market expected return shocks? Market-to-book ratios should be our *left-hand* variable, the thing we are trying to *explain*, not a sorting characteristic for expected returns.

Focusing on expected returns and betas rather than prices and discounted cashflows makes sense in a two-period or i.i.d. world, since in that case betas

²² Campbell and Mei (1993).

are all cashflow betas. It makes much less sense in a world with time-varying discount rates.

A long-run, price-and-payoff perspective may also end up being simpler. As a hint of the possibility, solve the Campbell–Shiller identity for long-run returns,

$$\sum_{j=1}^{\infty} \rho^{j-1} r_{t+j} = \sum_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j} - dp_t.$$

Long-run return uncertainty all comes from cashflow uncertainty. Long-run betas are all cashflow betas. The long run looks just like a simple one-period model with a liquidating dividend:

$$R_{t+1} = \frac{D_{t+1}}{P_t} = \left(\frac{D_{t+1}}{D_t} \right) \bigg/ \left(\frac{P_t}{D_t} \right),$$

$$r_{t+1} = \Delta d_{t+1} - dp_t.$$

A natural start is to forecast long-run returns and to form price decompositions in the cross-section, just as in the time series: to estimate forecasts such as

$$\sum_{j=1}^{\infty} \rho^{j-1} r_{t+j}^i = a + b' C_{it} + \varepsilon^i,$$

and then understand valuations with present value models as before.²³ The Appendix includes two simple examples.

In a formal sense, it does not matter whether you look at returns or prices. The expressions $1 = E_t(m_{t+1}R_{t+1})$ and $P_t = E_t \sum_{j=1}^{\infty} m_{t,t+j} D_{t+j}$ each imply the other. But, as I found with return forecasts, our economic understanding may be a lot different in a price, long-run view than if we focus on short-run returns. What constitutes a “big” or “small” error is also different if we look at prices rather than returns. At a 2% dividend yield, $D/P = (r - g)$ implies that an “insignificant” 10 bp/month expected return error is a “large” 12% price error, if it is permanent. For example, since momentum amounts to a very small time-series correlation and lasts less than a year, I suspect it has little effect on long-run expected returns and hence the level of stock prices. Long-lasting characteristics are likely to be more important. Conversely, small transient price errors can have a large impact on return measures. A tiny i.i.d. price error induces the appearance of mean reversion where there is none. Common procedures amount to taking many differences of prices, which amplify the error to signal ratio. For example, the forward spread $f_t^{(n)} - y_t^{(1)} = p_t^{(n-1)} - p_t^{(n)} + p_t^{(1)}$ is already a triple difference of price data.

III. Theories

Having reviewed a bit of *how* discount rates vary, let us think now about *why* discount rates vary so much.

²³ Vuolteenaho (2002) and Cohen, Polk, and Vuolteenaho (2003) are a start, with too few followers.

It is useful to classify theories by their main ingredient, and by which data they use to measure discount rates. My goal is to suggest for discount rates something like Fama's (1970) classification of informational possibilities. Here is an outline of the classification:

1. Theories based on fundamental investors, with few frictions.
 - (a) Macroeconomic theories. Ties to macro or microeconomic quantity data.
 - i. Consumption, aggregate risks.
 - ii. Risk sharing and background risks; hedging outside income.
 - iii. Investment and production.
 - iv. General equilibrium, including macroeconomics.
 - (b) Behavioral theories, focusing on irrational expectations. Ties to price data. Other data?
 - (c) Finance theories. Expected return-beta models, return-based factors, affine term structure models. Ties to price data, returns explained by covariances.
2. Theories based on frictions.
 - (a) Segmented markets. Different investors are active in different markets; limited risk bearing of active traders.
 - (b) Intermediated markets. Prices are set by leveraged intermediaries; funding difficulties.
 - (c) Liquidity.
 - i. Idiosyncratic liquidity: Is it easy to sell the asset?
 - ii. Systemic liquidity: How does an asset perform in times of market illiquidity?
 - iii. Trading liquidity: Is a security useful to facilitate trading?

A. Macroeconomic Theories

"Macro" theories tie discount rates to macroeconomic quantity data, such as consumption or investment, based on first-order conditions for the ultimate investors or producers.

For example, the canonical consumption-based model with power utility $E\sum_t \beta^t u(C_t)$, $u(C) = C^{1-\gamma}/(1-\gamma)$ relates discount rates to consumption growth,

$$m_{t+1} = \beta \frac{u_c(t+1)}{u_c(t)} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma},$$

$$E_t(R_{t+1}^{ei}) = R^f \text{cov}(R_{t+1}^{ei}, m_{t+1}) \approx \gamma \text{cov}(R_{t+1}^{ei} \Delta c_{t+1}),$$

where R^f is the risk-free rate, R^{ei} is an excess return, and $c = \log(C)$. High expected returns (low prices) correspond to securities that pay off poorly when consumption is low. This model combines frictionless markets, rational expectations and utility maximization, and risk sharing so that only aggregate risks matter for pricing. It evidently ties discount-rate variation to macroeconomic data.

A vast literature has generalized this framework, including (among others)²⁴ (1) nonseparability across goods, such as durable and nondurable,²⁵ or traded and nontraded goods; (2) nonseparability over time, such as habit persistence,²⁶ (3) recursive utility and long-run risks,²⁷ and (4) rare disasters, which alter measurements of means and covariances in “short” samples.²⁸

A related category of theories adds incomplete markets or frictions preventing some consumers from participating. Though they include “frictions,” I categorize such models here because asset prices are still tied to some fundamental consumer or investor’s economic outcomes. For example, if nonstockholders do not participate, we still tie asset prices to the consumption decisions of stockholders who do participate.²⁹

With incomplete markets, consumers still share risks as much as possible. The complete-market theorem that “all risks are shared,” marginal utility is equated across people i and j , $m_{t+1}^i = m_{t+1}^j$, becomes “all risks are shared as much as possible.” The projection of marginal utility on asset payoffs X is the same across people $\text{proj}(m_{t+1}^i|X) = \text{proj}(m_{t+1}^j|X) \equiv x^*$. We can still aggregate *marginal utility* *another* than aggregate *consumption* before constructing marginal utility. A discount factor $m_{t+1} = E_{t+1}(m_{t+1}^i) = \int f(i)m_{t+1}^i di$ prices assets, where E_{t+1} takes averages across people conditional on aggregates. For example, with power utility we have

$$m_{t+1} = \beta E_{t+1} \left[\left(\frac{C_{t+1}^i}{C_t^i} \right)^{-\gamma} \right].$$

The fact that we aggregate nonlinearly across people means that variation in the *distribution* of consumption matters to asset prices. Times in which there is more *cross-sectional risk* will be high discount-factor events.³⁰

Outside or nontradeable risks are a related idea. If a mass of investors has jobs or businesses that will be hurt especially hard by a recession, they avoid stocks that fall more than average in a recession.³¹ Average stock returns then reflect the tendency to fall more in a recession, in addition to market risk exposure. Though in principle, given a utility function, one could see such risks in consumption data, individual consumption data will always be so poorly measured that tying asset prices to more fundamental sources of risk may be more productive.

If we ask the “representative investor” in December 2008 why he or she is ignoring the high premiums offered by stocks and especially fixed income, the answer might well be “that’s nice, but I’m about to lose my job, and my business might go under. I can’t take any more risks right now, especially in

²⁴ See Cochrane (2007a) and Ludvigson (2011) for recent reviews.

²⁵ Eichenbaum, Hansen, and Singleton (1988); more recently, Yogo (2006).

²⁶ For example, Campbell and Cochrane (1999).

²⁷ Epstein and Zin (1989), Bansal and Yaron (2004), Hansen, Heaton, and Li (2008).

²⁸ Rietz (1988), Barro (2006).

²⁹ For example, Mankiw and Zeldes (1991), Ait-Sahalia, Parker, and Yogo (2004).

³⁰ Constantinides and Duffie (1996).

³¹ Fama and French (1996), Heaton and Lucas (2000).

securities that will lose value or become hard to sell if the recession gets worse.” These extensions of the consumption-based model all formalize this sensible intuition—as opposed to the idea that these consumers have wrong expectations, or that they would have been happy to take risks but intermediaries were making asset-pricing decisions for them.

Investment-based models link asset prices to firms’ investment decisions, and general equilibrium models include production technologies and a specification of the fundamental shocks. This is clearly the ambitious goal toward which we are all aiming. The latter tries to answer the vexing questions, where do betas come from, and what makes a company a “growth” or “value” company in the first place?³²

B. Behavioral Theories

I think “behavioral” asset pricing’s central idea is that people’s expectations are wrong.³³ It takes lessons from psychology to find systematic patterns to the “wrong” expectations. There are some frictions in many behavioral models, but these are largely secondary and defensive, to keep risk-neutral “rational arbitrageurs” from coming in and undoing the behavioral biases. Often, simple risk aversion by the rational arbitrageurs would serve as well. Behavioral models, like “rational” models, tie asset prices to the fundamental investor’s willingness, ability, or (in this case) perception of risk.

Behavioral theories are also discount-rate theories. A distorted probability with risk-free discounting is mathematically equivalent to a different discount rate:

$$p = \sum_s \pi_s m_s x_s = \frac{1}{R^f} \sum_s \pi_s^* x_s,$$

where

$$\pi_s^* \equiv \pi_s m_s R^f = \pi_s m_s / \sum_{s'} \pi_{s'} m_{s'},$$

s denote states of nature, π_s are true probabilities, m_s is a stochastic discount factor or marginal utility growth, x_s is an asset payoff in state s , and π_s^* are distorted probabilities.

It is therefore pointless to argue “rational” versus “behavioral” in the abstract. There is a discount rate and equivalent distorted probability that can rationalize any (arbitrage-free) data. “The market went up, risk aversion must have declined” is as vacuous as “the market went up, sentiment must have increased.” Any model only gets its bite by restricting discount rates or distorted expectations, ideally tying them to other data. The only thing worth arguing about is how persuasive those ties are in a given model and data set,

³² A few good examples: Gomes, Kogan, and Zhang (2003), Gala (2010), Gourio (2007).

³³ See Barberis and Thaler (2003) and Fama (1998) for reviews.

and whether it would have been easy for the theory to “predict” the opposite sign if the facts had come out that way.³⁴ And the line between recent “exotic preferences” and “behavioral finance” is so blurred³⁵ that it describes academic politics better than anything substantive.

A good question for any theory is what *data* it uses to tie down discount rates. By and large, behavioral research so far largely ties prices to other prices; it looks for price patterns that are hard to understand with other models, such as “overreaction” or “underreaction” to news. Some behavioral research uses survey evidence, and survey reports of people’s expectations are certainly unsettling. However, surveys are sensitive to language and interpretation. People report astounding discount rates in surveys and experiments, yet still own long-lived assets, houses, and durable goods. It does not take long in teaching MBAs to realize that the colloquial meanings of “expect” and “risk” are entirely different from conditional mean and variance. If people report the *risk-neutral* expectation, then many surveys make sense. An “optimistic” cashflow growth forecast is the same as a “rational” forecast, discounted at a lower rate, and leads to the correct decision, to invest more. And the risk-neutral expectation—the expectation weighted by marginal utility—is the right sufficient statistic for many decisions. Treat painful outcomes as if they were more probable than they are in fact.

Of course, “rational” theories beyond the simple consumption-based model struggle as well. Changing expectations of consumption 10 years from now (long-run risks) or changing probabilities of a big crash are hard to distinguish from changing “sentiment.” At least one can aim for more predictions than assumptions, tying together several phenomena with a parsimonious specification.

C. Finance Theories

“Finance” theories tie discount rates to broad return-based factors. That’s great for data reduction and practical applications. The more practical and “relative-pricing” the application, the more “factors” we accept on the right-hand side. For example, in evaluating a portfolio manager, hedging a portfolio, or finding the cost of capital for a given investment, we routinely include momentum as a “factor” even though we do not have a deep theory of why the momentum factor is priced.

However, we still need the deeper theories for deeper “explanation.” Even if the CAPM explained individual mean returns from their betas and the market premium, we would still have the equity premium puzzle—why is the market premium so large? (And why are betas what they are?) Conversely, even if we had the perfect utility function and a perfect consumption-based model, the fact

³⁴ Fama (1998).

³⁵ For example, which of Epstein and Zin (1989), Barberis, Santos, and Huang (2001), Hansen and Sargent (2005), Laibson (1997), Hansen, Heaton and Li (2008), and Campbell and Cochrane (1999) is really “rational” and which is really “behavioral?”

that consumption data are poorly measured means we would still use finance models for most practical applications.³⁶

A nice division of labor results. Empirical asset pricing in the Fama and French (1996) tradition boils down the alarming set of anomalies to a small set of large-scale systematic risks that generate rewards. “Macro,” “behavioral,” or other “deep” theories can then focus on why the factors are priced.

D. Theories with Frictions

Models that emphasize frictions are becoming more and more popular, especially since the financial crisis. At heart, these models basically maintain the “rational” assumption. Admittedly, there are often “irrational” agents in such models. However, these agents are usually just convenient shortcuts rather than central to the vision. A model may want some large volume of trade,³⁷ or to include some “noise traders,” while focusing clearly on the delegated management problem or the problem of leveraged intermediaries. For such a purpose, it is easy simply to assume a slightly irrational class of trader rather than spell out those trader’s motives from first principles. However, such assumptions are not motivated by deep reading of psychology or lab experiments. The focus is on the frictions and behavior of intermediaries rather than the risk-bearing ability of ultimate investors or their psychological misperceptions.

I think it is useful to distinguish three categories of frictions: (1) segmented markets, (2) intermediated markets or “institutional finance,”³⁸ and (3) liquidity. Surely, this is a broad brush categorization and more detailed divisions can usefully be made.

E. Segmented and Intermediated Markets

I distinguish “segmented markets” from “intermediated markets,” as illustrated in Figure 8. Segmentation is really about limited risk sharing among the pool of investors who are active in a particular market.³⁹ Their limited risk bearing can generate “downward-sloping demands” (in quotes, because maybe it is “supply”), and average returns that depend on a “local” factor, little and poorly linked CAPMs.⁴⁰ Given the factor zoo, which is an attractive idea.

“Intermediated markets” or “institutional finance” refers to a different, vertical rather than horizontal, separation of investors from payoffs. Investors use delegated managers. Agency problems in delegated management then spill over into asset prices. For example, suppose investors have “equity” and “debt”

³⁶ Campbell and Cochrane (2000) give a quantitative example.

³⁷ Scheinkman and Xiong (2003).

³⁸ Markus Brunnermeier coined this useful term.

³⁹ Some important examples: Burnnermeier and Pedersen (2009), Brunnermeier (2009), Gabaix, Krishnamurthy, and Vigneron (2007), Duffie and Strulovici (2011), Garleanu and Pedersen (2011), He and Krishnamurthy (2010), Krishnamurthy (2008), Froot and O’Connell (2008), Vayanos and Vila (2011).

⁴⁰ For example, Gabaix, Krishnamurthy, and Vigneron (2007).



Segmented markets**Intermediated markets**

Figure 8. Segmented markets versus intermediated markets.

claims on the managers. When losses appear, the managers stave off bankruptcy by trying to sell risky assets. But since all the managers are doing the same thing, prices fall and discount rates rise. Colorful terms such as “fire sales” and “liquidity spirals” describe this process.⁴¹

Of course, one must document and explain segmentation and intermediation. As suggested by the dashed arrows in Figure 8, there are strong incentives to undo any price anomaly induced by segmentation or intermediation. Models with these frictions often abstract from deep-pockets unintermediated investors—sovereign wealth funds, pension funds, endowments, family offices, and Warren Buffets—or institutional innovation to bridge the friction. Your “fire sale” is their “buying opportunity” and business opportunity. A little more attention to the reasons for segmentation and intermediation may help us to understand when and for how long these models apply. For example, transactions costs, attention costs, or limited expertise suggest that markets can be segmented until the “deep pockets” arrive, but that they do arrive eventually.

⁴¹ Brunnermeier (2009) and Brunnermeier and Pedersen (2009), for example.

So if this is why markets are segmented, segmentation will be more important in the short run, after unusual events, or in more obscure markets. If I try to sell a truckload of tomatoes at 2 am in front of the Booth school, I am not likely to get full price. But if I do it every night, tomato buyers will start to show up. In the flash crash, it took about 10 minutes for buyers to show up, which is either remarkably long or remarkably short, depending on your point of view.

A crucial question is, as always, what *data* will this class of theories use to measure discount rates? Arguing over puzzling patterns of prices is weak. The rational-behavioral debate has been doing that for 40 years, rather unproductively. Ideally, one should tie price or discount-rate variation to central items in the models, such as the balance sheets of leveraged intermediaries, data on who is actually active in segmented markets, and so forth. I grant that such data are hard to find.⁴²

F. Liquidity

We have long recognized that some assets have higher or lower discount rates in compensation for greater or lesser liquidity.⁴³ We have also long struggled to define and measure liquidity. There are (at least) three kinds of stories for liquidity that are worth distinguishing. Liquidity can refer to the ease of buying and selling an individual security. Illiquidity can also be systemic: Assets will face a higher discount rate if their prices fall when the market as a whole is illiquid, whether or not the asset itself becomes more or less illiquid. Finally, assets can have lower discount rates if they facilitate information trading for assets, as money facilitates physical trading of goods.

I think of “liquidity” as different from “segmentation” in that segmentation is about limited risk-bearing ability, while liquidity is about trading. Liquidity is a feature of assets, not the risks to which they are claims. Many theories of liquidity emphasize asymmetric information, not limited risk-bearing ability—assets become illiquid when traders suspect that anyone buying or selling knows something, not because traders are holding too much of a well-understood risk. Some kinds of liquidity, such as the off-the-run Treasury spread, refer to different prices of economically identical claims. Understanding liquidity requires us to unravel the puzzle of why real people and institutions trade so much more than they do in our models.

G. Efficiency and Discount Rates

All of these theories and related facts are really about discount rates, expected returns, risk bearing, risk sharing, and risk premiums. None are

⁴² Mitchell, Pedersen, and Pulvino (2007) is a good example. They document who was active in convertible arbitrage markets through two episodes in which specialized hedge funds left the market and it took months for multi-strategy funds to move in.

⁴³ Acharya and Pedersen (2005), Amihud, Mendelson, and Pedersen (2005), Cochrane (2005a), Pastor and Stambaugh (2003), Vayanos and Wang (2011).

fundamentally about slow or imperfect diffusion of cashflow information, informational “inefficiency” as Fama (1970) defined it. Informational efficiency is not wrong or disproved. Efficiency basically won, and we moved on. When we see information, it is quickly incorporated into asset prices. There is a lot of asset-price movement not related to visible information, but Hayek (1945) told us that would happen, and we learned that a lot of such price variation corresponds to expected returns. Little of the (large) gulf between the above models is really about information. Seeing the facts and the models as categories of discount-rate variation seems much more descriptive of most (not all) theory and empirical work.

Informational efficiency is much easier for markets and models to obtain than wide risk sharing or desegmentation, which is perhaps why it was easier to verify. A market can become informationally efficient with only one informed trader, who does not need to actually buy anything or take any risk. He should run into a wall of indexers, and just bid up the asset he knows is underpriced.⁴⁴ Though in reality price discovery seems to come with a lot of trading, it does not have to do so. Risk sharing needs *everyone* to change their portfolios and bear a risk in order to eliminate segmentation. For example, if the small-firm effect came from segmentation, the passively managed small stock fund should have ended it—but it also took the invention and marketing of such funds to end it. The actions of small numbers of arbitrageurs could not do so.

IV. Recent Performance

This is not the place for a deep review of theory and empirical work supporting or confronting theories. Instead, I think it will be more productive to think informally about how these classes of models might be able to handle big recent events.

A. Consumption

I still think the macro-finance approach is promising. Figure 9 presents the market price-dividend ratio, and aggregate consumption relative to a slow-moving “habit.” The habit is basically just a long moving average of lagged consumption, so the surplus-consumption ratio line is basically detrended consumption.⁴⁵

Consumption and stock market prices *did* collapse together in 2008. Many high average-return securities and strategies (stocks, mortgage-backed securities, low-grade bonds, momentum, currency carry) collapsed more than low average-return counterparts. The basic consumption-model logic—securities must pay higher returns, or fetch lower prices, if their values fall more when consumption falls—is not drastically wrong.

Furthermore, the habit model captures the idea that people become more risk averse as consumption falls in recessions. As consumption nears habit,

⁴⁴ Milgrom and Stokey (1982).

⁴⁵ Campbell and Cochrane (1999).




Figure 9. Surplus-consumption ratio and price-dividend ratio. The price-dividend ratio is that of the CRSP NYSE value-weighted portfolio. The surplus consumption is formed from monthly real nondurable consumption using the Campbell and Cochrane (1999) specification and parameters, multiplied by three to fit on the same scale.

people are less willing to take risks that involve the same proportionate risk to consumption. Discount rates rise, and prices fall. Lots of models have similar mechanisms, especially models with leverage.⁴⁶ In the habit model, the price-dividend ratio is a nearly log-linear function of the surplus-consumption ratio. The fit is not perfect, but the general pattern is remarkably good, given the hue and cry about how the crisis invalidates all traditional finance.

B. Investment

The Q theory of investment is the off-the-shelf analogue to the simple power-utility model from the producer point of view. It predicts that investment should be low when valuations (market to book) are low, and vice versa,

$$1 + \alpha \frac{i_t}{k_t} = \frac{\text{market value}_t}{\text{book value}_t} = Q_t, \quad (10)$$

where i_t = investment and k_t = capital.

⁴⁶ For example, Longstaff (2008).




Figure 10. Investment-capital ratio, price-dividend ratio, and market-to-book ratio. Investment is real private nonresidential fixed investment. Capital is cumulated from investment with an assumed 10% annual depreciation rate. The price-dividend ratio is that of the CRSP S&P500 portfolio. The market-to-book ratio comes from Ken French's website.

Figure 10 contrasts the investment-capital ratio, market-to-book ratio, and price-dividend ratio. The simple Q theory also links asset prices and investment better than you probably supposed, both in the tech boom and in the financial crisis.

Many finance puzzles are stated in terms of returns. To make that connection, one can transform (10) into a relation linking asset returns to investment growth. Many return puzzles are mirrored in investment growth as the Q theory suggests.⁴⁷

Q theory also reminds us that supply as well as demand matters in setting asset prices. If capital could adjust freely, stock values would never change, no matter how irrational investors are. Quantities would change instead.

I do not argue that consumption or investment *caused* the boom or the crash. Endowment-economy causal intuition does not hold in a production economy. These first-order conditions are happily consistent with a view, for example, that losses on subprime mortgages were greatly amplified by a run on the shadow banking system and flight to quality,⁴⁸ which certainly qualifies as a

⁴⁷ Cochrane (1991b, 1996, 2007a), Lamont (2000), Li, Livdan, and Zhang (2008), Liu, Whited, and Zhang (2009), Belo (2010), Jermann (2010), Liu and Zhang (2011).

⁴⁸ Cochrane (2011).

“friction.” The first-order conditions are consistent with many other views of the fundamental determinants of both prices and quantities. But the graphs do argue that asset prices and discount rates are much better *linked* to big macroeconomic events than most people think (and vice versa). They suggest an important amplification mechanism: If people did not become more risk averse in recessions, and if firms could quickly transform empty houses into hamburgers, asset prices would not have declined as much.

I do not pretend to have perfect versions of either of these first-order conditions, let alone a full macro model that captures value or the rest of the factor zoo. These are very simple and rejectable models. Each makes a 100% R^2 prediction that is easy, though a bit silly, to formally reject: The habit model predicts that the price-dividend ratio is a function of the surplus-consumption ratio, with no error. The Q theory predicts that investment is a function of Q, with no error, as in (10). The point is only that research and further elaboration of these kinds of models, as well as using their basic intuition as an important guide to events, is not a hopeless endeavor.

C. Comparisons

Conversely, I think the other kinds of models, though good for describing particular anomalies, will have greater difficulty accounting for big-picture asset-pricing events, even the huge movements of the financial crisis.

We see a *pervasive, coordinated* rise in the premium for systematic risk,⁴⁹ common across all asset classes, and present in completely unintermediated and unsegmented assets. For example, Figure 11 plots government and corporate rates, and Figure 12 plots the BAA-AAA spread with stock prices. You can see a huge credit spread open up and fade away along with the dip in stock prices.

Behavioral ideas—narrow framing, salience of recent experience, and so forth—are good at generating anomalous prices and mean returns in individual assets or small groups. They do not easily generate this kind of coordinated movement across all assets that looks just like a rise in risk premium. Nor do they naturally generate covariance. For example, “extrapolation” generates the slight autocorrelation in returns that lies behind momentum. But why should all the momentum stocks then rise and fall together the next month, just as if they are exposed to a pervasive, systematic risk?

Finance models do not help, of course, because we are looking at variation of the factors, which those models take as given.

Segmented or institutional models are not obvious candidates to understand broad market movements. Each of us can easily access stocks and bonds through low-cost indices.

And none of these models naturally describe the strong correlation of discount rates with macroeconomic events. Is it a coincidence that people become

⁴⁹ The “systematic” adjective is important. People don’t seem to drive more carefully in recessions.




Figure 11. BAA, AAA, and Treasury yields. Source: Board of Governors of the Federal Reserve via Fred website.




Figure 12. Common Risk Premiums. P/D is the S&P500 price-dividend ratio from CRSP. S&P500 is the level of the S&P500 index from CRSP. BAA-AAA is that bond spread, from the Board of Governors. P/D is divided by 15 and the S&P500 is divided by 500 to fit on the same scale.

irrationally pessimistic when the economy is in a tailspin, and they could lose their jobs, houses, or businesses if systematic events get worse?

Again, macro is not everything—understanding the smaller puzzles is important. The point is that *looking* for macro underpinnings discount-rate variation, even through fairly simple models, is not as hopelessly anachronistic as many seem to think.

D. Arbitrages?

One of the nicest pieces of evidence for segmented or institutional views is that arbitrage relationships were violated in the financial crisis.⁵⁰ Unwinding the arbitrage opportunities required one to borrow dollars, which intermediary arbitrageurs could not easily do.

Figure 13 gives one example. CDS plus Treasury should equal a corporate bond, and usually does. Not in the crisis.

Figure 14 gives another example: covered interest parity. Investing in the United States versus investing in Europe and returning the money with forward rates should give the same return. Not in the crisis.

Similar patterns happened in many other markets, including even U. S. Treasuries.⁵¹ Now, any arbitrage opportunity is a dramatic event. But in each case here the *difference* between the two ways of getting the same cashflow is dwarfed by the *overall* change in prices. And, though an “arbitrage,” the price differences are not large enough to attract long-only deep-pocket money. If your precious cash is in a U.S. money market fund, 20 basis points in the depth of a financial crisis is not enough to get you to listen to the salesman offering offshore investing with an exchange rate hedging program.

So maybe it is possible that the “macro” view still builds the benchmark story of *overall* price change, with very interesting *spreads* opening up due to frictions. At least we have a theory for that. Constructing a theory in which the arbitrage spreads drive the coordinated rise in risk premium seems much harder.

The price of coffee displays arbitrage opportunities across locations at the ASSA meetings. (The AFA gave it away for free downstairs while Starbucks was selling it upstairs.) The arbitrage reflects an interesting combination of transactions costs, short-sale constraints, consumer biases, funding limits, and other frictions. Yet we do not dream that this fact matters for big-picture variation in worldwide commodity prices.

E. Liquidity Premia and Trading Value

Trading-related liquidity does strike me as potentially important for the big picture, and a potentially important source of the low discount rates in “bubble” events.⁵²

⁵⁰ See also Fleckenstein, Longstaff, and Lustig (2010),

⁵¹ Hu, Pan, and Wang (2011).

⁵² Cochrane (2001, 2003, 2005a), Garber (2000), Krishnamurthy (2002), O’Hara (2008), Scheinkman and Xiong (2003).



Figure 13. Citigroup CDS and Bond Spreads. Source: Fontana (2010).




Figure 14. Three-month Libor and FX Swap Rates. Source: Baba and Packer (2009).

I am inspired by one of the most obvious “liquidity” premiums: Money is overpriced—it has a lower discount rate—relative to government debt, though they are claims to the same payoff in a frictionless market. And this liquidity spread can be huge, hundreds of percent in hyperinflations.

Now, money is “special” for its use in transactions. But many securities are “special” in trading. Trading needs a certain supply of their physical shares. We cannot support large trading volumes by recycling one outstanding share at arbitrarily high speed. Even short sellers must hold a share for a short period of time.

When share supply is small, and trading demand is large, markets can drive down the discount rate or drive up the price of highly traded securities, as they do for money. These effects have long been seen in government bonds, for example, in the Japanese “benchmark” effect, the spreads between on-the-run and off-the-run Treasuries, or the spreads between Treasury and agency bonds.⁵³ Could these effects extend to other assets?

Figures 15 and 16 are suggestive. The stock *price* rise and fall of the late 1990s was concentrated in NASDAQ and NASDAQ Tech. The stock *volume* rise and fall was concentrated in the same place. *Every* asset price “bubble”—defined here by people’s use of the label—has coincided with a similar trading frenzy, from Dutch tulips in 1620 to Miami condos in 2006.

Is this a coincidence? Do prices rise and fall for other reasons, and large trading volume follows, with no effect on price? Or is the high price—equivalently a low discount rate—explained at least in part by the huge volume, that is, by the value of shares in facilitating a frenzy of information trading?

To make this a deep theory, we must answer why people trade so much. At a superficial level, we know the answer: The markets we study *exist* to support information-based trading. Yet, we really do not have good models of information-based trading.⁵⁴ Perhaps the question of how information is incorporated in asset markets will come back to the center of inquiry.

V. Applications

Finance is about practical application, not just deep explanation. Discount-rate variation will change applications a lot.

A. Portfolio Theory

A huge literature explores how investors should exploit the market-timing and intertemporal-hedging opportunities implicit in time-varying expected returns.⁵⁵

⁵³ Boudoukh and Whitelaw (1991), Longstaff (2004), Krishnamurthy and Vissing-Jorgensen (2010).

⁵⁴ Milgrom and Stokey (1982).

⁵⁵ Merton (1971), Barberis (2000), Brennan, Schwartz, and Lagnado (1997), Campbell and Viceira (1999, 2002), Pastor (2000); see a review in Cochrane (2007b).




Figure 15. NASDAQ Tech, NASDAQ, and NYSE indices. Source: Cochrane (2003).




Figure 16. Dollar volume in NASDAQ Tech, NASDAQ, and NYSE. Source: Cochrane (2003).




Figure 17. Multifactor efficient frontiers. Investors minimize variance given mean and covariance with the extra factor. A three-fund theorem emerges (left). The market portfolio is multifactor efficient, but not mean-variance efficient (right).

But *the average investor must hold the market portfolio*. We cannot all time the market, we cannot all buy value, and we cannot all be smarter than average. We cannot even all rebalance. No portfolio advice other than “hold the market” can apply to everyone. A useful and durable portfolio theory must be consistent with this theorem. Our discount-rate facts and theories suggest one, built on *differences* between people.

Consider Fama and French’s (1996) story for value. The average investor is worried that value stocks will fall at the same time his or her human capital falls. But then some investors (“steelworkers”) will be more worried than average, and should *short* value despite the premium; others (“tech nerds”) will have human capital correlated with growth stocks and buy lots of value, effectively selling insurance. A two-factor model implies a three-fund theorem, and a three-dimensional multifactor efficient frontier as shown in Figure 17.⁵⁶ It is not easy for an investor to figure out how much of three funds to hold.

And now we have *dozens* of such systematic risks for each investor to consider. Time-varying opportunities create more factors, as habits or leverage shift some investors’ risk aversion over time more or less than others. *Unpriced* factors are even more important. Our steelworker should start by shorting a steel industry portfolio, even if it has zero alpha. Zero-alpha portfolios are attractive, as they provide actuarially fair insurance. We academics should understand the variation across people in risks that are hedgeable by systematic factors, and find low-cost portfolios that span that variation.⁵⁷ Yet we have spent all our time looking for priced factors that are only interesting for the measure-zero mean-variance investor!

⁵⁶ See Fama (1996), and Cochrane (2007b).

⁵⁷ Heaton and Lucas (2000).

All of this sounds hard. That's good! We finally have a reason for a fee-based "tailored portfolio" industry to exist, rather than just to deplore it as folly. We finally have a reason for us to charge large tuitions to our MBA students. We finally have an interesting portfolio theory that is not based on chasing zero-sum alpha.

A.1. State Variables

Discount-rate variation means that state-variable hedging should matter. It is almost completely ignored in practice. Almost all hedge funds, active managers, and institutions still use mean-variance optimizers. This is particularly striking given that they follow active strategies, predicated on the idea that expected returns and variances vary a lot over time!

Perhaps state-variable hedging seems nebulous, and therefore maybe small and easy to ignore. Here is a story to convince you otherwise. Suppose you are a highly risk-averse investor, with a 10-year horizon. You are investing to cover a defined payment, say your 8-year-old child's future tuition at the University of Chicago. The optimal investment is obviously a 10-year zero-coupon indexed Treasury (TIP).⁵⁸ Figure 18 tracks your investment over time.

Suppose now that bond prices plunge, and volatility surges, highlighted in the graph. Should you sell in a panic to avoid the risk of further losses? No. You should tear up the statement. "Short-term volatility" is irrelevant. Every decline in price comes with a corresponding rise in expected return. Evaluating bonds with a one-period mean-variance, alpha-beta framework is silly—though a surprising amount of the bond investing world does it!

That is pretty obvious, but now imagine that you are a stock investor in December 2008—say, you are running your university's endowment. As shown in Figure 19, stocks plummeted, and stock volatility, rose dramatically, from 16% to 70%.

Should you sell, to avoid the risks of further losses? The standard formula says so. Picking a mean return and risk aversion to justify a 60% equity share in normal times, you should reduce the portfolio's equity share to 4%:

$$\text{Equity Share} = \frac{1}{\gamma} \frac{E(R^e)}{\sigma^2(R^e)} \cdot 0.6 = \frac{1}{2} \frac{0.04}{0.18^2} \Rightarrow \frac{1}{2} \frac{0.04}{\mathbf{0.70^2}} = \mathbf{0.04}.$$

(You might object that mean returns rose too. But they would had to have risen to $4 \times 0.70^2/0.18^2 = 60\%$ for this formula to tell you not to change allocation. You also may object that many investors had leverage, tenured professor salaries to pay, or other habit-like considerations for becoming more risk averse. Fair enough, but then one-period mean-variance theory is particularly inappropriate in the first place.)

But not everyone can do this. The market did not fall $1 - 4/60 = 93\%$. If you are selling, who is buying? Is everyone else being stupid? Does it make any sense to think that the market was irrationally *over*-valued in the midst of the financial crisis?

⁵⁸ Campbell and Viceira (2001), Wachter (2003).




Figure 18. Bond Price Example. The solid line tracks the price of a zero-coupon bond that matures in year 10. The dashed lines illustrate the yield, and the fact that the bond price will recover by year 10. The shaded area represents a “crisis” in which interest rates rise and interest-rate volatility rises.

The answer, of course, is that one-period mean-variance analysis is completely inappropriate. If the world were i.i.d., volatility could not change in the first place. Stocks are a bit like bonds: Price-dividend drops increase expected returns.⁵⁹ To some extent, “short-run volatility” does not matter to a long-run investor. *State-variable hedging matters a lot, even for simple real-world applications.* And, by ICAPM logic, we should therefore expect to observe multiple priced factors. Time-series predictability should be a strong source of additional pricing factors in the “cross section.”

A.2. Prices and Payoffs

Or maybe not. Telling our bond investor to hold 10-year zeros because their “price happens to covary properly with state variables for their investment opportunities” just confuses the obvious. It is much clearer to look at the final payoff and tell him to *ignore* price fluctuations. Maybe dynamic portfolio theory overall might get a lot simpler if we look at payoff streams rather than look at dynamic trading strategies that achieve those streams.

⁵⁹ Campbell and Vuolteenaho (2004).



Figure 19. S&P500 Price Index and index volatility. Realized volatility is a 20-day average.




Figure 20. Volatility. A 20-day realized daily volatility, and VIX index.

If you look at payoff streams, it is obvious that an indexed perpetuity (or annuity) is the risk-free asset for long-term investors, despite arbitrary time-varying return moments, just as the 10-year zero was obviously the risk-free asset for my bond investor. It is interesting that coupon-only TIPS are perceived to be an exotic product, or a way to speculate on inflation, not the benchmark risk-free asset for every portfolio in place of a money-market investment.

How about risky investments? Here is a simple and suggestive step.⁶⁰ If utility is quadratic,

$$\max_{\{c_t\}} E \sum_{t=0}^{\infty} \beta^t \left(-\frac{1}{2} \right) (c_t - c^*)^2,$$

it turns out that we can still use two-period mean-variance theory to think about *streams of payoffs* (loosely, streams of dividends), no matter how much expected returns vary over time. To be specific,

1. *Every optimal payoff stream combines an indexed perpetuity and a claim to the aggregate dividend stream.* Less risk-averse investors hold more of the claim to aggregate dividends, and vice versa.
2. *Optimal payoffs lie on a long-run mean/long-run variance frontier.* The “long-run mean” $\tilde{E}(x)$ in this statement sums over time as well as across states,

$$\tilde{E}(x) = \frac{1}{1-\beta} \sum_{j=0}^{\infty} \beta^j E(x_{t+j}).$$

3. *State variables disappear from portfolio theory*, just as they did for our 10-year TIP investor, once he looked at the 10-year problem.

If our stock market investor thought this way when facing the crisis, he would answer: “I invested in the aggregate dividend stream. Why should I buy or sell? I don’t look at the statements.” This is a lot simpler to explain and implement than state-variable identification, deep time-series modeling of investment opportunities, value function calculations, and optimal hedge portfolios!

If investors have outside income in this framework, they first short a payoff stream most correlated with their outside income stream, and then hold the mean-variance efficient payoffs. Calculating long-run correlations of income streams this way may be easier than trying to impute discount-rate induced changes in the present value of outside income streams, which are needed to calculate return-based hedge portfolios.

If investors have no outside income, long-run expected returns (payoffs divided by initial prices) line up with long-run market betas. A CAPM emerges, despite arbitrary time variation in expected returns and variances. ICAPM pricing factors fade away as we look at longer horizons. If investors do have outside income, the payoff corresponding to average outside income payoff

⁶⁰ Results from Cochrane (2008).




Figure 21. Hedge Fund Returns. One-year excess returns of the “equity market neutral” hedge fund index and the CRSP value-weighted portfolio. Data source: hedgeindex.com and CRSP.

emerges as a second priced factor, in the style of Fama and French’s (1996) human capital story for the value effect.

Of course, quadratic utility is a troublesome approximation, especially for long-term problems. Still, this simple example captures the possibility that a price and payoff approach can give a much simpler view of pricing and portfolio theory than we get by focusing on the high-frequency dynamic trading strategy that achieves those payoffs in a given market structure.

B. Alphas, Betas, and Performance Evaluation

In the 1970 view, there is one source of systematic risk, the market index. Active management chases “alpha,” which means uncovering assets whose prices do not reflect available information.

Now we have dozens of dimensions of systematic risks. Many hedge fund strategies include an element of option writing. For example, Figure 21 shows the annual returns of “equity market neutral” hedge funds together with the market return. “Providing liquidity” looks a lot like writing out-of-the-money!⁶¹

⁶¹ Mitchell and Pulvino (2001), Asness, Krail, and Liew (2001), Agarwal, and Naik (2004).

I tried telling a hedge fund manager, “You don’t have alpha. Your returns can be replicated with a value-growth, momentum, currency and term carry, and short-vol strategy.” He said, “‘Exotic beta’ is my alpha. I understand those systematic factors and know how to trade them. My clients don’t.” He has a point. How many investors have even thought through their exposures to carry-trade or short-volatility “systematic risks,” let alone have the ability to program computers to execute such strategies as “passive,” mechanical investments? To an investor who has not heard of it and holds the market index, a new factor is alpha. And that alpha has nothing to do with informational inefficiency.

Most active management and performance evaluation today just is not well described by the alpha–beta, information-systematic, selection-style split anymore. There is no “alpha.” There is just beta you understand and beta you do not understand, and beta you are positioned to buy versus beta you are already exposed to and should sell.

C. Procedures, Corporate Finance, Accounting, and Regulation

Time-varying discount rates and multiple factors deeply change many applications.

The first slide in a capital budgeting lecture looks something like this

$$\text{Value of investment} = \frac{\text{Expected payout}}{R^f + \beta [E(R^m) - R^f]},$$

with a 6% market premium. All of which, we now know, is completely wrong. The market premium is not always 6%, but varies over time by as much as its mean. (And uncertainty about the market premium is also several percentage points.) Expected returns do not line up with CAPM betas, but rather with multifactor betas to the extent we understand them at all. And since expected returns change over time, the discount rate is different for cashflows at different horizons.

It is interesting that investment decisions get so close to right anyway, with high investment when stock prices are high. (Remember Figure 10.) Evidently, a generation of our MBAs figured out how to jigger the numbers and get the right answer despite a wrong model. Perhaps what we often call “irrational” cashflow forecasts, optimistic in good times and pessimistic in bad times, are just a good way to offset artificially constant discount rates. Or perhaps they understood the Q theory lecture and just follow its advice.

I do not think the answer lies in multifactor betas⁶² or discounting with dynamic present value models and time-varying risk premia, at least not yet. Capital budgeting is a “relative pricing” exercise—we want to use available information in asset markets to help us decide what the discount rate for a given project should be. For this purpose, simply looking at average returns of “similar” securities is enough. Understanding discount rates as a function of characteristics—or, better, understanding valuations directly as a function

⁶² Fama and French (1997) try.

of characteristics (the use of “comparables”)—may end up being more fruitful. We do not have to *explain* discount rates—relate expected returns to betas and understand their deep economics—in order to *use* them. We do not need an all-purpose model of everything to extend prices from known assets to a new one. Even when discount rates are explained, the characterization (characteristic models) may be a better measure for practical relative pricing than the explanation (beta models). Conversely, capital budgeting gives the same answer if discount rates are “wrong.” When you shop for a salad, all you care about is the price of tomatoes. Whether tomatoes are expensive because the trucks got stuck in bad weather or because of an irrational bubble in the tomato futures market makes no difference to your decision.

Many procedures in accounting, regulation, and capital structure implicitly assume that returns are independent over time, and hence that prices only reflect cashflow information.

Suppose that a firm has a single cashflow in 10 years, and is funded by a zero-coupon bond and equity. In most accounting, capital structure, and regulation procedures we would use the stock and bond prices to calculate the probability of and distance to default. But if prices decline because discount rates rise, that fact has no implication for the probability of or distance to default.

Perhaps banks' complaint that low asset prices represent “illiquidity” or “temporarily depressed valuations” rather than insolvency—a lesser chance of making future interest and principal repayments—make some sense. Perhaps capital requirements do not have to respond immediately to such events. Perhaps “hold to maturity” accounting is not as silly as it sounds. Perhaps the fact that firms change capital structures very slowly in response to changes in equity valuations makes some sense.⁶³

Of course, in such an event the *risk-neutral* probability of default has risen. Maybe regulators, bondholders, and capital structure *should* respond to a rise in the state price of the default event exactly as they respond to a rise in the real probability of that event. Maybe, but at least it is a very different issue and worth asking the question.

I am not arguing that mark-to-market accounting is bad, or that fudging the numbers is a good idea. The point is only that what you *do* with a mark-to-market number might be quite different in a world driven by discount-rate variation than in a world driven by cashflow variation.⁶⁴ The mark-to-market value is no longer a sufficient statistic. A loss of value coincident with a rise in expected return has different implications than a loss of value with a decline in expected return. Decisions need to incorporate more information, not less.

The view that the stock price is driven by expected earnings lies behind stock-based executive compensation as well. It is already a bit of a puzzle that executives holding company stock or options should therefore hold systematic risks due to recessions, market betas, or commodity price exposures, about

⁶³ Welch (2004).

⁶⁴ Heaton, Lucas, and McDonald (2009).

which they can do nothing. Understanding that a large fraction of stock returns reflect changes in discount rates or new-factor beta exposures makes the logic of such incentives even more curious. Perhaps stock-based compensation has less to do with effort and operating performance, and more with tax treatment or incentives for risk management.

D. Macroeconomics

Large variation in risk premia implies exciting changes for macroeconomics. Most of macroeconomics focuses on variation in a single intertemporal price, “the” interest rate, which intermediates saving and investment. Yet in the recent recession, as shown in Figure 11, interest rates paid by borrowers (and received by any investors willing to lend) spiked up, while short-term government rates went down. Recessions are all about changes in credit spreads, about the willingness to bear risk and the amount of risk to be borne, far more than they are about changes in the desire for current versus future certain consumption. Most of the Federal Reserve’s response consisted of targeting risk premiums, not changing the level of interest rates or addressing a transactions demand for money.

Macroeconomics and finance have thought very differently about consumer (we call them investors) and firm behavior. For example, the consumers in the Campbell and Cochrane (1999) habit model balance very strong precautionary saving motives with very strong intertemporal substitution motives, and have large and time-varying risk aversion. Their behavior is very far from the permanent-income intuition, its borrowing-constrained alternative, or central role of simple intertemporal substitution (the modern “IS” curve) in macroeconomic thinking.

As one simple example, macroeconomists often think about how consumers will respond to a change in “wealth” coming from a change in stock prices or house prices. Financial economists might suspect that consumers will respond quite differently to a decline in value coming from a discount-rate rise—a temporary change in price with no change in capital stock or cashflow—than one that comes from a change in expected cashflows or destruction of physical capital stock.

Financial models also emphasize adjustment costs or irreversibilities: If firms can freely transform consumption goods to capital, then stock prices (Q) are constant. Yet, most “real business cycle” literature following King, Plosser, and Rebelo (1988) left out adjustment costs because they did not need them to match basic quantity correlations. The first round of “new-Keynesian” literature abstracted from capital altogether, and much work in that tradition continues to do so. Figure 10 pretty strongly suggests that Q is not constant. This figure, the regression evidence of Table I, and interest rates in Figure 11, together suggest that variations in the *risk premium* drive investment, not variation in the level of risk-free interest rates. Adjustment costs lead to basic differences in analysis. For example, without adjustment costs, the marginal product of capital can be less than one, clashing with the zero bound on nominal rates.

With adjustment costs, the price of capital can fall, giving a positive real rate of interest.

Formal macroeconomics has started to introduce some of the same ingredients that macro-finance researchers are using to understand discount-rate variation, including “new” preferences, adjustment cost or other frictions in capital formation.⁶⁵ Since the 2008 financial crisis, there has been an explosion of macroeconomic models with financial frictions, especially credit-market frictions. Still, we are a long way from a single general-equilibrium model that matches basic quantity and price facts.

Everyone is aware of the question. The job is just hard. Macroeconomic models are technically complicated. Macroeconomic models with time-varying risk premia are even harder. Adding financial frictions while maintaining the models’ dynamic intertemporal character is harder still. At a deeper level, successful “grand synthesis” models do not consist of just mixing all the popular ingredients together and stirring the pot; they must maintain the clear quantitative-parable feature of good economic analysis.

An asset-pricing perspective also informs monetary economics, and the interaction between monetary and fiscal policy. From a finance perspective, nominal government debt is “equity” in the government: it is the residual claim to primary fiscal surpluses. Hence, the price level must satisfy the standard asset-pricing equation:

$$\frac{\text{Debt}_t}{\text{Price level}_t} = E_t \sum_{j=0}^{\infty} m_{t,t+j} (\text{real primary surplus}_{t+j}). \quad (11)$$

Inflation can absorb shocks to surpluses, just as equity absorbs shocks to profit streams. This fact is at least an important constraint on monetary policy, especially in a time of looming deficits.⁶⁶ It suggests there is a large component of inflation that the Fed is powerless to avoid. The analogy to stocks also suggests that variation in the *discount rate* $m_{t,t+j}$ for government debt is important. A “flight to quality” lowers the discount rate for government debt, raising the right-hand side of (11). People want to hold more government debt, which means getting rid of goods and services. This story links the “rising risk premium” which finance people see as the core of a recession with the “decline in aggregate demand” which macroeconomists see. The standard corporate finance perspective also illuminates the choice of government debt maturity structure and denomination. Indexed or foreign-currency debt is debt, which must be repaid or defaulted. Domestic-currency debt is equity, which can be

⁶⁵ For example Christiano, Eichenbaum, and Evans (2005). However, this is also a good example of remaining differences between macroeconomic and finance models. They use a one-period habit, which does not generate slow-moving expected excess returns. They also use an adjustment cost tied to investment growth rather than to the investment-capital ratio, which does not generate the Q theory predictions of Figure 10 and related finance literature. Of course, their choices produce better quantity dynamics, so we need to meet in the middle somewhere.

⁶⁶ Sargent and Wallace (1981), Cochrane (2011).

inflated. Long-term debt allows bond prices to temporarily absorb shocks to future surplus, rather than result in inflation immediately when short-term debt cannot be rolled over.

VI. Conclusion

Discount rates vary a lot more than we thought. Most of the puzzles and anomalies that we face amount to discount-rate variation we do not understand. Our theoretical controversies are about how discount rates are formed. We need to recognize and incorporate discount-rate variation in applied procedures.

We are really only beginning these tasks. The facts about discount-rate variation need at least a dramatic consolidation. Theories are in their infancy. And most applications still implicitly assume i.i.d. returns and the CAPM, and therefore that price changes only reveal cashflow news. Throughout, I see hints that discount-rate variation may lead us to refocus analysis on prices and long-run payoff streams rather than one-period returns.

Appendix

A. Return Forecasts and VARs

A.1. Identities

Campbell and Shiller (1988) Taylor-expand the definition of log return to obtain the approximation

$$r_{t+1} = \kappa - \rho dp_{t+1} + dp_t + \Delta d_{t+1}, \quad (\text{A1})$$

where $dp_t \equiv \log(D_t/P_t)$, $r_t = \log(R_t)$, $d_t = \log(D_t)$, $\rho \equiv \frac{PD}{1+PD}$, PD denotes the point about which one takes the approximation, and $\kappa = -(1 - \rho)\log(1 - \rho) - \rho\log(\rho)$. The point of approximation PD need not be the mean, so one can use this and following identities to examine cross-sectional variation in dividend yields without security-specific approximation points. When we are only interested in second moments, we interpret symbols as deviations from means, leaving

$$r_{t+1} = -\rho dp_{t+1} + dp_t + \Delta d_{t+1}. \quad (\text{A2})$$

Iterating (A2) forward, we obtain

$$dp_t = \sum_{j=1}^k \rho^{j-1} r_{t+j} - \sum_{j=1}^k \rho^{j-1} \Delta d_{t+j} + \rho^k dp_{t+k}.$$

Assuming the latter term goes to zero—the absence of “rational bubbles”—we obtain the Campbell–Shiller present value identity,

$$dp_t = \sum_{j=1}^{\infty} \rho^{j-1} r_{t+j} - \sum_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j}. \quad (\text{A3})$$

This is an ex post identity, so it also holds in conditional expectations using any information set.

A.2. Dividend Construction

I create dividends from the CRSP annual return series with and without dividends. These series are defined as

$$R_{t+1} \equiv \frac{P_{t+1} + D_{t+1}}{P_t}; \quad RX_{t+1} \equiv \frac{P_{t+1}}{P_t}.$$

Then, I construct the dividend yield as

$$\frac{D_{t+1}}{P_{t+1}} = \frac{R_{t+1}}{RX_{t+1}} - 1.$$

By using an annual horizon, I avoid the strong seasonal in dividend payments even when using monthly observations.

This definition reinvests dividends to the end of the year.⁶⁷ Therefore, dividend growth also includes some return information. Annual sums of dividends are a good deal less volatile. However, the identity $R_{t+1} = (P_{t+1} + D_{t+1})/P_t$ and therefore (A2), (A3), and following calculations do not hold using sums of dividends, or dividends reinvested at the risk-free rate.

This definition of dividends has a small practical advantage as well. The resulting dividend-price ratio is a better univariate return forecaster because it removes a good deal of 1-year dividend growth forecastability. For example, consider the sharp stock market decline in Fall 2008. Using a simple sum of past dividends, we would see a large decline in the price-dividend ratio. But much of the stock price decline surely reflected news that dividends in 2009 were going to fall. Reinvested dividends were lower than the sum, and the resulting price-dividend ratio thus included the information that dividends would decline in 2009.

I construct dividend growth by

$$\frac{D_{t+1}}{D_t} = \frac{(D_{t+1}/P_{t+1})}{(D_t/P_t)}.$$

For the VAR in Tables II–IV, I use instead dividend growth implied by the identity (A1),

$$\Delta d_{t+1} = \kappa + r_{t+1} + \rho dp_{t+1} - dp_t.$$

Actual dividend growth gives very similar results. However, this construction means that Campbell–Shiller approximate identities hold exactly, so it is easier to see the identities in the results. To make identities hold, it is better to use “pure” returns rather than infer returns from dividend growth, otherwise approximation errors can show up as magic investment opportunities.

⁶⁷ The algebra is in Cochrane (1992).

A.3. Shock Definition

I identify the dividend growth, dividend yield, and *cay* shocks in Figure 5 by setting changes to the other variables to zero in each case.

By (A2), this dividend growth shock must come with a contemporaneous return shock,

$$\varepsilon_{t+1}^d = 1, \quad \varepsilon_{t+1}^{dp} = 0, \quad \varepsilon_{t+1}^{cay} = 0, \quad \varepsilon_{t+1}^r = 1,$$

and this dividend yield shock must come with a contemporaneous negative return shock,

$$\varepsilon_{t+1}^d = 0, \quad \varepsilon_{t+1}^{dp} = 1, \quad \varepsilon_{t+1}^{cay} = 0, \quad \varepsilon_{t+1}^r = -\rho.$$

Since this *cay* shock affects neither dividend growth nor dividend yield, it also comes with no change in return r_t ,

$$\varepsilon_{t+1}^d = 0, \quad \varepsilon_{t+1}^{dp} = 0, \quad \varepsilon_{t+1}^{cay} = 1, \quad \varepsilon_{t+1}^r = 0.$$

I choose this definition of shocks because it leads to nicely interpretable responses as “cashflow” and “discount rate” news. Because dividends remain roughly unpredictable, this “short-run” identification gives almost the same result as a “long-run” identification. If, rather than define the first shock as $\varepsilon_{t+1}^{dp} = 0$ and $\varepsilon_{t+1}^{cay} = 0$, we had identified the first shock by

$$(E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^{j-1} r_{t+j} = 0,$$

we would have gotten nearly the same result. The resulting shocks are nearly uncorrelated, which is also convenient. However, there is no guarantee that these elegant properties will survive as we add more variables.

This VAR is very simple, since I left dividend growth and returns out of the right-hand side. My purpose is to distill the essential message of more complex VARs, not to deny that there may be some information in additional lags of these variables.

A.4. Identities in the *cay* VAR

The present-value identity (A3), conditioned down and reproduced here,

$$dp_t = E \left[\sum_{j=1}^{\infty} \rho^{j-1} r_{t+j} \mid I_t \right] - E \left[\sum_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j} \mid I_t \right],$$

implies that an extra variable can only help dp to forecast long-horizon returns if it also helps dp to forecast long-horizon dividend growth. An extra variable can help to forecast 1-year returns by changing the term structure of return forecasts as well. Here I show how this intuition applies algebraically to multiple regression coefficients and impulse-response functions.

Regressing both sides of (A3) on dp_t and z_t ,

$$r_t^{lr} = \sum_{j=1}^{\infty} \rho^{j-1} r_{t+j} = a_r + b_r^{lr} dp_t + c_r^{lr} z_t + \varepsilon^r,$$

$$\Delta d_t^{lr} = \sum_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j} = a_d + b_d^{lr} dp_t + c_d^{lr} z_t + \varepsilon^d,$$

we obtain the generalized restriction on long-run multiple regression coefficients,

$$1 = b_r^{lr} - b_d^{lr}, \quad (\text{A4})$$

$$0 = c_r^{lr} - c_d^{lr}. \quad (\text{A5})$$

Equation (A4) is the same as before, now applied to the multiple regression coefficient. Equation (A5) expresses the idea that a new variable can only help to forecast long-run returns if it also helps to forecast long-run dividend growth. The extra dividend growth and return forecasts will be perfectly negatively correlated. In this way, extra long-run dividend growth forecastability means *more* long-run return forecastability, not less.

In terms of individual long-horizon regressions

$$r_{t+j} = b_r^{(j)} dp_t + c_r^{(j)} z_t + \varepsilon_{t+j}^r,$$

etc., (A3) similarly implies

$$1 = \sum_{j=1}^{\infty} \rho^{j-1} b_r^{(j)} - \sum_{j=1}^{\infty} \rho^{j-1} b_d^{(j)},$$

$$0 = \sum_{j=1}^{\infty} \rho^{j-1} c_r^{(j)} - \sum_{j=1}^{\infty} \rho^{j-1} c_d^{(j)}.$$

A variable can help to forecast *1-year* returns, $c_r^{(1)} \neq 0$, only if it correspondingly changes the forecast of longer horizon returns or dividend growth.

Since impulse-response functions are the same as regression coefficients of future variables such as r_{t+j} on shocks at time t , the impulse-response functions must obey the same relation,

$$1 = \sum_{j=1}^{\infty} \rho^{j-1} e_{dp \rightarrow r}^{(j)} - \sum_{j=1}^{\infty} \rho^{j-1} e_{dp \rightarrow \Delta d}^{(j)},$$

$$0 = \sum_{j=1}^{\infty} \rho^{j-1} e_{z \rightarrow r}^{(j)} - \sum_{j=1}^{\infty} \rho^{j-1} e_{z \rightarrow \Delta d}^{(j)},$$

where $e_{dp \rightarrow r}^{(j)}$ denotes the response of r_{t+j} to a dp_t shock. This fact lets me easily interpret the change in forecastability by adding *cay*, in the context of

Table AI
Return Forecasts Using Additional Predictors

The return forecasts are of the form

$$r_{t+1} = a + b \times dp_t + c \times z_t + \varepsilon_{t+1}.$$

Data are from Welch and Goyal (2008), 1947–2009. I calculate the variance of long-horizon expected returns and dividend growth from a bivariate VAR, and using actual (not identity) dividend growth forecasts. EQUIS, percentage equity issuance, is the ratio of equity issuing activity as a fraction of total issuing activity. SVAR is stock variance, computed as the sum of squared daily returns on the S&P500. IK is the investment to capital ratio, the ratio of aggregate (private nonresidential fixed) investment to aggregate capital for the whole economy. DFY, the default yield spread, is the difference between BAA- and AAA-rated corporate bond yields.

	<i>dp</i>	<i>cay</i>	<i>EQUIS</i>	<i>SVAR</i>	<i>IK</i>	<i>DFY</i>
<i>c</i>		2.21	−0.71	1.48	−5.30	5.25
<i>t(c)</i>		(1.73)	(−2.53)	(3.40)	(−0.85)	(1.86)
<i>b</i>	0.13	0.10	0.19	0.15	0.11	0.13
<i>t(b)</i>	(2.61)	(1.82)	(3.75)	(3.05)	(2.16)	(2.53)
<i>R</i> ²	0.10	0.16	0.19	0.15	0.11	0.13
$\sigma(E_t \Sigma_{j=1}^{\infty} \rho^{j-1} r_{t+j})$	0.52	0.46	0.49	0.42	0.53	0.49
$\sigma(E_t \Sigma_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j})$	0.17	0.13	0.16	0.11	0.17	0.14

the present value identity, by plotting the impulse responses. The numbers in Figure 5 are terms of this decomposition.

A.5. Results Using Welch–Goyal Predictors

To see if the pattern of the *cay* VAR holds more generally, I tried a number of the forecasting variables in Welch and Goyal (2008). The results are in Table AI. Each of these variables helps substantially to forecast one-period returns. Yet the variables mean-revert quickly and do not forecast dividends much, so the contribution to the variance of dividend yields still comes mostly from the variance of long-run expected returns.

A.6. More Lags of Dividend Growth and Returns

An obvious first source of additional variables is less restrictive VARs than the simple first-order VAR that I presented in the text. As usual in VARs, individual coefficients largely enter insignificantly, so it takes some art or prior information to see robust patterns. And by (A2), the same forecasts can be interpreted as regressions using additional lags of any two of returns, dividend yield, and dividend growth.

The second lag of dividend yields is at least economically important. Table AII presents the regressions. The change in dividend yield helps the return forecast, increasing *R*² from 0.09 to 0.15, and correspondingly increasing the more interesting measures of expected return variation. The change in dividend yield really helps to forecast dividend growth, with a 3.27 *t*-statistic,

5% standard deviation of forecast, and a forecast that varies by 90% of the mean. However, the 0.10 autocorrelation in Δdp_t suggests that this will be a very short-lived signal, one with little impact on forecasts of long-run dividend growth or returns, and thus to our view of the sources of price-dividend ratio volatility.

Similarly, while *individual* r_{t-j} and Δd_{t-j} coefficients do not look big and do not have much pattern, they can nonetheless help as a *group*, especially if one sensibly restricts the pattern of lagged coefficients. In this vein, Lacerda and Santa-Clara (2010) and Koijen and van Binsbergen (2010) find that moving averages of past dividend growth help to forecast both returns and dividend growth (as they must, given the present value identity), almost doubling the return-forecast R^2 .

A.7. VAR Calculations

To find long-run regression coefficients implied by a first-order VAR as in Table II, I run

$$r_{t+1} = b_r dp_t + \varepsilon_{t+1}^r, \quad (\text{A6})$$

$$\begin{aligned} \Delta d_{t+1} &= b_d dp_t + \varepsilon_{t+1}^d, \\ dp_{t+1} &= \phi dp_t + \varepsilon_{t+1}^{dp}. \end{aligned} \quad (\text{A7})$$

I then report

$$b_r^{(k)} = b_r \frac{1 - (\rho\phi)^k}{1 - \rho\phi}.$$

To calculate long-run regression coefficients as in Table IV, with $z = cay$, I write the VAR as

$$\begin{aligned} \begin{bmatrix} dp_{t+1} \\ z_{t+1} \end{bmatrix} &= \Phi \begin{bmatrix} dp_t \\ z_t \end{bmatrix} + \begin{bmatrix} \varepsilon_{t+1}^{dp} \\ \varepsilon_{t+1}^z \end{bmatrix}, \\ \begin{bmatrix} r_{t+1} \\ \Delta d_{t+1} \end{bmatrix} &= B \begin{bmatrix} dp_t \\ z_t \end{bmatrix} + \begin{bmatrix} \varepsilon_{t+1}^r \\ \varepsilon_{t+1}^d \end{bmatrix}. \end{aligned}$$

Table AII
Return Forecasts with Additional Lags

Forecasts using dividend yield and change in dividend yield. CRSP value-weighted return, 1947–2009. $\Delta dp_t = dp_t - dp_{t-1}$.

Left-Hand Variable	dp_t	Δdp_t	$t(dp_t)$	$t(\Delta dp_t)$	R^2	$\sigma[E_t(y)]\%$	$\frac{\sigma[E_t(y)]}{E(y)}$
r_{t+1}	0.13	0.26	(2.45)	(1.83)	0.15	6.76	0.65
Δd_{t+1}	0.03	0.35	(0.62)	(3.27)	0.14	4.98	0.90
dp_{t+1}	0.93	0.10	(24.7)	(0.85)	0.91		
Δdp_{t+1}	−0.07	0.10	(−1.85)	(0.85)	0.06		

I then report

$$E_t \begin{bmatrix} r_t^{lr} \\ \Delta d_t^{lr} \end{bmatrix} = B \sum_{j=1}^{\infty} \rho^{j-1} \Phi^{j-1} \begin{bmatrix} dp_t \\ z_t \end{bmatrix} = B(I - \rho\Phi)^{-1} \begin{bmatrix} dp_t \\ z_t \end{bmatrix}.$$

B. Asset Pricing as a Function of Characteristics

B.1. Portfolio Spreads

In the text, I related 1–10 portfolio means to Sharpe ratios of underlying factors. Here is the result. Suppose that expected returns rise with a characteristic C_i (for example, log book-to-market ratio)

$$E(R^{ei}) = a + b \times C_i,$$

and this variation corresponds exactly to a factor (for example, hml),

$$R_t^{ei} = \beta^i \times f_t + \varepsilon_t,$$

with betas that also rise with the characteristic

$$\beta^i = \frac{a}{E(f)} + \frac{b}{E(f)} \times C_i,$$

and with residuals that are uncorrelated

$$\text{cov}(\varepsilon^i, \varepsilon^j) = 0.$$

Now, consider the usual 1–10 or 1–20 spread portfolio. Its mean and variance are

$$E(R^{ei} - R^{ej}) = b(C_i - C_j),$$

$$\sigma^2(R^{ei} - R^{ej}) = (\beta^i - \beta^j)^2 \sigma(f)^2 + 2 \frac{\sigma_\varepsilon^2}{N} = \frac{b^2}{E(f)^2} (C_i - C_j)^2 \sigma(f)^2 + 2 \frac{\sigma_\varepsilon^2}{N},$$

where N is the number of securities in each portfolio. The Sharpe ratio, which is proportional to the t -statistic for the mean spread-portfolio return, is

$$\frac{E(R^{ei} - R^{ej})}{\sigma(R^{ei} - R^{ej})} = \frac{E(f)}{\sigma(f)} \frac{b(C_i - C_j)}{\sqrt{b^2 (C_i - C_j)^2 + 2 \frac{\sigma_\varepsilon^2}{N} \frac{E(f)^2}{\sigma(f)^2}}}.$$

This Sharpe ratio rises as we look at further-separated portfolios. As $C_i - C_j$ increases, it approaches the pure Sharpe ratio of the factor $E(f)/\sigma(f)$. The Sharpe ratio does not increase forever. Splitting into finer portfolios can get the magic 1% per month portfolio mean or alpha, but cannot arbitrarily raise Sharpe ratios or t -statistics. Splitting the portfolio more finely reduces N , so




Figure A1. Value effect before and after 1963. Average returns on Fama–French 10 portfolios sorted by book-to-market equity versus CAPM betas. Monthly data. Source: Ken French’s website.

splits that are too fine end up reducing the Sharpe ratio and t -statistic by including too much idiosyncratic risk.

Having seen this analysis, however, it would seem more efficient to use the information in all assets, not just the tail portfolios, by examining the cross-sectional regression coefficient $\hat{b}$. Since $\hat{b} = \text{cov}(E(R^i), C^i) / \text{var}(C^i) = E(R^i \times [C^i - E(C^i)]) / \text{var}(C^i)$, this regression coefficient is the mean of a factor that is formed as a linear function of the characteristic.

B.2. Value, Betas, and Samples

In the text, I emphasized that all puzzles are joint puzzles of expected returns and betas, and cautioned that the value puzzle does not hold in pre-1963 U.S. data. Figure A1 presents the CAPM in the Fama–French 10 book-to-market portfolios before and after 1963. In the left-hand panel, you see the familiar failure of the CAPM—average returns are higher in the value portfolio, but there is no association between the wide spread in average returns and market betas. The right-hand panel shows average returns and betas before 1963. Here the CAPM works remarkably well. The big change is not in the pattern of average returns. Value still earns more than growth. The big change is betas: Value firms have higher betas than growth firms in the pre-1963 period.⁶⁸

B.3. Time Series and Cross-Section

As a first step toward understanding mean returns as a function of characteristics, and to help make the ideas concrete, Table AIII presents regressions

⁶⁸ Davis, Fama, and French (2000), Campbell and Vuolteenaho (2004), Ang and Chen (2007), Fama and French (2006).

using the Fama–French 25 size and book-to-market portfolios. I use log book-to-market and log size relative to the market portfolio.

The first row of Table AIII gives a pure cross-sectional regression. The regression fits the portfolio average returns quite well, with a 77% R^2 . (One does better still with a $size \times bm$ cross-term, allowing the growth portfolios to have a different slope on size than the value portfolios.)

The second row of Table AIII gives a pooled forecasting regression, which is the most natural way to integrate time-series and cross-sectional approaches. The *size* coefficient is a little smaller, and the *bm* coefficient is much larger.

To diagnose the difference between the cross-sectional and pooled regressions, rows 3 and 4 present a regression with time dummies and a regression with portfolio dummies respectively. Variation *over time* in a *given* portfolio's book-to-market ratio is a much stronger signal of return variation than the same variation *across portfolios* in average book-to-market ratio.

When we run such regressions for individual firms, we cannot use dummies, since the average return of a specific company over the whole sample is meaningless. The goal of this regression is to mirror portfolio formation and remove firm name completely from the list of characteristics. The last line of Table AIII gives a way to capture the difference between time-series and cross-sectional approaches without dummies: It allows an independent effect of recent changes in the characteristics. This specification accounts quite well for the otherwise unpalatable time and portfolio dummies. The portfolio dummy regression coefficient that captures time-series variation is quite similar to the sum of the level and recent-change coefficients. It is also gratifyingly similar to the “recent-change” effect in aggregate dividend yield regressions of Table AII. One could of course capture the same phenomenon with portfolios, by sorting based on level and recent change of characteristics. But my goal is to explore the other direction of this equivalence.

Next, we want to run regressions like this on individual data, and find a similar characterization of the covariance matrix as a function of characteristics. Then, we can expand the analysis to multiple right-hand variables.

B.4. Prices in the Cross-Section

Section II.C of the text suggested that we try to understand the variation in prices (price-dividend ratios) across time and portfolios by exploring long-run return predictability in the cross-section. How much of the difference between one asset's price-dividend ratio (or price earnings, book to market, etc.) and another's is due to variation in expected returns, and how much is due to expected dividend growth or other cashflow expectations?⁶⁹

To explore this question and clarify the idea, I examine the 10 Fama–French book-to-market portfolios. Eventually, we want to do this analysis in individual security data and avoid the use of portfolios altogether, but the portfolios are a

⁶⁹ Vuolteenaho (2002), Cohen, Polk, and Vuolteenaho (2003).

Table AIII
Characteristic Regressions

The regression specification in the first row is

$$E(R_{t+1}^{ei}) = a + b \times E(\text{size}_{it}) + c \times E(\text{bm}_{it}) + \varepsilon_i; \quad i = 1, 2, \dots, 25.$$

The remaining rows are

$$R_{t+1}^{ei} = a + (a_t) + (a_i) + b \times \text{size}_{i,t} + c \times \text{bm}_{i,t} + d \times (\text{size}_{i,t} - \text{size}_{i,t-12}) + e \times (\text{bm}_{i,t} - \text{bm}_{i,t-12}).$$

Terms in parentheses only appear in some regressions. *size* is log(market equity/total market equity). *bm* is log(book equity/market equity). Monthly data, 1947–2009. Data from Ken French’s website.

Method	Regression Coefficients			
	<i>size_t</i>	<i>bm_t</i>	Δsize_t	Δbm_t
1. Cross-section	−0.030	0.27		
2. Pooled	−0.022	0.55		
3. Time dummies	−0.031	0.29		
4. Portfolio dummies	−0.087	1.48		
5. Pooled	−0.030	0.46	−0.38	1.11

simple place to start. Figure A2 presents the average return, dividend growth, and dividend yield of the portfolios.

Over long horizons, dividend yields are stationary so long-term average returns come from dividend yields and dividend growth. Taking unconditional means of the return identity (A2), and imposing stationarity so $E(dp_t) = E(dp_{t+1})$,

$$E(r^i) = (1 - \rho) E(dp^i) + E(\Delta d^i). \quad (\text{A8})$$

Figure A2 shows that value portfolio returns come roughly half from greater dividend growth and half from a larger average dividend yield.

Our objective is to produce variance decompositions over time and across securities as I did with the market return. Flipping (A8) around, we have

$$E(dp^i) = \frac{1}{1 - \rho} [E(r^i) - E(\Delta d^i)]. \quad (\text{A9})$$

The same observation about the sources of return gives a fairly extreme version of the usual surprising result about prices. Low prices—high dividend yields—correspond to *high* dividend growth and thus to *even higher* returns.⁷⁰

The first column of Table AIV, Panel A expresses the same idea in a purely cross-sectional regression. From (A9), the coefficients in such a regression obey

$$1 = \frac{b_r^{cs}}{1 - \rho} - \frac{b_d^{cs}}{1 - \rho}, \quad (\text{A10})$$

⁷⁰ Chen, Petkova, and Zhang (2008, section 2.2.2) discuss this puzzle.




Figure A2. Components of Average Returns. Average return r_{t+1} , dividend growth Δd_{t+1} , and dividend yield dp_t for the Fama–French 10 book-to-market portfolios, 1947–2009. The dashed Δd line gives mean dividend growth implied by the approximate identity $\Delta d_{t+1} = r_{t+1} - \kappa + \rho dp_{t+1} - dp_t$.

where the b are the cross-sectional regression coefficients of the terms in (A9). We can interpret these coefficients as the fraction of *cross-sectional* dividend yield variation driven by discount rates and the fraction driven by dividend growth. (Vuolteenaho (2002) uses a different present value identity to understand variation in the book-to-market ratio directly, rather than use dividend yields as I have. This is a better procedure for individual stocks, which often do not pay dividends. I use dividend yields here for simplicity.) The results are quite similar to the time-series regressions for the market portfolio from Tables II to IV: More than 100% of the cross-sectional variation in average dividend yields of these portfolios comes from cross-sectional variation in expected returns (1.33). Expected dividend growth goes “the wrong way”—low prices correspond to high dividend growth, as seen in Figure A2. (Sample means obey the identity

$$E(dp_t^i) = \frac{1}{1-\rho} \left[E(r_{t+1}^i) - E(\Delta d_{t+1}^i) + \rho \frac{1}{T} (dp_T^i - dp_1^i) \right].$$

The last term is not zero, which is why the coefficients in the $b/(1-\rho)$ column do not add up following (A10).)

Table AIV
Long-Run Panel-Data Regressions

The regression specification in the first column is

$$E(R_{t+1}^{ei}) = a + b \times E(dp_{it}) + \varepsilon_i; \quad i = 1, 2, \dots, 10,$$

and similarly for Δd . The specification in the remaining columns is

$$R_{t+1}^{ei} = a + (a_i) + (a_i) + b \times dp_{i,t} + (c \times \Delta dp_{i,t}) + \varepsilon_{i,t+1},$$

and similarly for Δd and dp . ϕ represents the coefficient of dp_{t+1}^i on dp_t^i in the same regression. Terms in parentheses only appear in some regressions. Annual data on 10 Fama–French size and book-to-market sorted portfolios, 1947–2009. Data from Ken French’s website. $\Delta d_{t+1} = r_{t+1} - \kappa + \rho dp_{t+1} - dp_t$. I use $\rho = 0.96$.

Left- Hand Variable	Cross- Section		Portfolio Dummies		Time Dummies		Pooled		Pooled	
	b	$\frac{b}{1-\rho}$	b	$\frac{b}{1-\rho\phi}$	b	$\frac{b}{1-\rho\phi}$	b	$\frac{b}{1-\rho\phi}$	dp	Δdp
Panel A: Book-to-Market Portfolios										
r	0.053	1.33	0.107	0.90	0.044	0.33	0.095	0.97	0.090	0.074
Δd	0.026	0.64	−0.011	−0.10	−0.092	−0.68	−0.003	−0.03	−0.012	0.076
dp			0.92		0.90		0.94		0.94	0.002
Panel B: Size Portfolios										
r	−0.014	−0.36	0.077	1.02	0.023	0.27	0.067	0.95		
Δd	−0.048	−1.20	0.002	0.02	−0.063	−0.73	−0.004	−0.05		
dp			0.963		0.952		0.968			

We can, of course, ask how much of the *time variation* in these dividend yields around their *portfolio average* corresponds to return versus dividend growth forecasts. A regression that includes portfolio dummies, shown next in Table AIV, Panel A, addresses this question. The 0.11 return-forecasting coefficient for portfolios is almost the same as the return-forecasting coefficient for the market as a whole seen in Tables II–IV. The dividend growth forecast is also nearly zero. So *all variation in book-to-market sorted portfolio dividend yields over time, about portfolio means, corresponds to variation in expected returns*, much like that of market returns.

The regression with time dummies, next in Table AIV, Panel A, paints a different picture. The return coefficient is smaller at 0.044, and ϕ is smaller as well, so expected returns only account for 33% of the variation in dividend yields. We finally see an important dividend growth forecast, with the right sign, −0.09, accounting for 68% of dividend yield volatility. The strong contrast of this result with the pure cross-sectional regression means that *a time of unusually large cross-sectional dispersion in dividend yields corresponds to unusually high dispersion in dividend growth forecasts*.

This is an important regression, in that it shows a component of variation in valuations that does correspond to dividend growth forecasts. The unusual

dispersion in dividend growth forecasts adds up to zero, so this kind of variation cannot be seen in the aggregate dividend yield and its forecasting relations. There is variation in forecastable dividend growth, which drives some individual variation in dividend-price ratios. But it averages out across all securities, so that the aggregate dividend yield is driven primarily by expected returns.

A pooled regression with no dummies looks much like the time-series regression with portfolio dummies. There is more time variation in dividend yields than cross-sectional variation, so, adding them up evenly, the time variation dominates the pooled regression.

The last column of Table AIV, Panel A follows Table AIII, to try to unite time-series and cross-sectional variation without using dummies. It shows a very similar result, with the Δdp variable accounting for much of the dividend growth forecastability. The next step is to calculate the price implications of this multivariate regression, as I did with *cay*, but that takes us too far afield of this simple example.

The Fama–French size portfolios, shown in Table AIV, Panel B, paint a quite different picture. The pure cross-sectional regression (first column) shows cash-flow effects: Higher prices (low dividend yields) are associated with higher subsequent dividend growth, which by one measure fully accounts for the dividend yield variation! However, with portfolio dummies we again see that practically all dividend yield variation *over time for a given portfolio* comes from expected return variation, just as for the market as a whole. With time dummies, variation *across portfolios in a given time period* is split between return and dividend growth forecasts.

REFERENCES

- Ait-Sahalia, Yacine, Jonathan Parker, and Motohiro Yogo, 2004, Luxury goods and the equity premium, *Journal of Finance* 59, 2959–3004.
- Acharya, Viral, and Lasse Heje Pedersen, 2005, Asset pricing with liquidity risk, *Journal of Financial Economics* 77, 375–410.
- Agarwal, Vikas, and Narayan Naik, 2004, Risk and portfolio decisions involving hedge funds, *Review of Financial Studies* 17, 63–98.
- Amihud, Yakov, Haim Mendelson, and Lasse Heje Pedersen, 2005, Liquidity and asset prices, *Foundations and Trends in Finance* 1, 269–364.
- Ang, Andrew, and Joe Chen, 2007, CAPM over the long run: 1926–2001, *Journal of Empirical Finance* 14, 1–40.
- Asness, Clifford, Robert Krail, and John Liew, 2001, Do hedge funds hedge? *Journal of Portfolio Management* 28, 6–19.
- Baba, Naohiko, and Frank Packer, 2009, Interpreting deviations from covered interest parity during the financial market turmoil of 2007–08, *Journal of Banking and Finance* 33, 1953–1962.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.
- Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225–264.
- Barberis, Nicholas, Tano Santos, and Ming Huang, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1–53.

- Barberis, Nicholas, and Richard Thaler, 2003, A survey of behavioral finance, in George Constantinides, René Stulz, and Milton Harris, eds. *Handbook of the Economics of Finance* (North Holland, Amsterdam).
- Belo, Frederico, 2010, Production-based measures of risk for asset pricing, *Journal of Monetary Economics* 57, 146–163.
- Boudoukh, Jacob, and Robert Whitelaw, 1991, The benchmark effect in the Japanese government bond market, *Journal of Fixed Income* 1, 52–59.
- Brennan, Michael J., Eduardo S. Schwartz, and Roland Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Brunnermeier, Markus, 2009, Deciphering the liquidity and credit crunch 2007–2009, *Journal of Economic Perspectives* 23, 77–100.
- Brunnermeier, Markus, and Lasse H. Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.
- Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.
- Campbell, John Y., and John Ammer, 1993, What moves the stock and bond markets? A variance decomposition for long-term asset returns, *Journal of Finance* 48, 3–37.
- Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and John H. Cochrane, 2000, Explaining the poor performance of consumption-based asset pricing models, *Journal of Finance* 55, 2863–2878.
- Campbell, John Y., and Jianping Mei, 1993, Where do betas come from? Asset pricing dynamics and the sources of systematic risk, *Review of Financial Studies* 6, 567–592.
- Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Campbell, John Y., and Robert J. Shiller, 1991, Yield spreads and interest rate movements: A bird's eye view, *Review of Economic Studies* 58, 495–514.
- Campbell, John Y., and Luis M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics* 114, 433–495.
- Campbell, John Y., and Luis M. Viceira, 2001, Who should buy long-term bonds? *American Economic Review* 91, 99–127.
- Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, *Clarendon Lectures in Economics* (Oxford University Press, Oxford).
- Campbell, John Y., and Tuomo Vuolteenaho, 2004, Bad beta, good beta, *American Economic Review* 94, 1249–1275.
- Chen, Long, Ralitsa Petkova, and Lu Zhang, 2008, The expected value premium, *Journal of Financial Economics* 87, 269–280.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, 2005, Nominal rigidities and the dynamic effects of a shock to monetary policy, *Journal of Political Economy* 113, 1–45.
- Cochrane, John H., 1991a, Volatility tests and efficient markets: A review essay, *Journal of Monetary Economics* 27, 463–485.
- Cochrane, John H., 1991b, Production-based asset pricing and the link between stock returns and economic fluctuations, *Journal of Finance* 46, 207–234.
- Cochrane, John H., 1992, Explaining the variance of price-dividend ratios, *Review of Financial Studies* 5, 243–280.
- Cochrane, John H., 1994, Permanent and transitory components of GNP and stock prices, *Quarterly Journal of Economics* 109, 241–266.
- Cochrane, John H., 1996, A cross-sectional test of an investment-based asset pricing model, *Journal of Political Economy* 104, 572–621.
- Cochrane, John H., 2001, Review of 'Famous First Bubbles: The Fundamentals of Early Manias', *Journal of Political Economy* 109, 1150–1154.

- Cochrane, John H., 2003, Stock as money: Convenience yield and the tech-stock bubble, in William C. Hunter, George G. Kaufman, and Michael Pomerleano, eds. *Asset Price Bubbles* (MIT Press, Cambridge).
- Cochrane, John H., 2005a, Liquidity, trading, and asset prices. NBER asset pricing program report, *NBER Reporter* (National Bureau of Economic Research, Cambridge).
- Cochrane, John H., 2005b, *Asset Pricing* (Princeton University Press, Princeton).
- Cochrane, John H., 2006, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533–1575.
- Cochrane, John H., 2007a, Financial markets and the real economy, in Rajnish Mehra, ed. *Handbook of the Equity Risk Premium* (Elsevier, Amsterdam).
- Cochrane, John H., 2007b, Portfolio theory, Working paper, University of Chicago.
- Cochrane, John H., 2008, A mean-variance benchmark for intertemporal portfolio theory, Working paper, University of Chicago.
- Cochrane, John H., 2011, Understanding fiscal and monetary policy in the great recession: Some unpleasant fiscal arithmetic, *European Economic Review* 55, 2–30.
- Cochrane, John H., and Monika Piazzesi, 2005, Bond risk premia, *American Economic Review* 95, 138–160.
- Cochrane, John H., and Monika Piazzesi, 2008, Decomposing the yield curve, Working paper, University of Chicago.
- Cohen, Randolph B., Christopher K. Polk, and Tuomo Vuolteenaho, 2003, The value spread, *Journal of Finance* 58, 609–642.
- Constantinides, George M., and Darrell Duffie, 1996, Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219–240.
- Davis, James L., Eugene F. Fama, and Kenneth R. French, 2000, Characteristics, covariances and average returns: 1929–1997, *Journal of Finance* 55, 389–406.
- Daniel, Kent, and Sheridan Titman, 2006, Testing factor-model explanations of market anomalies, Working paper, Columbia University and University of Texas.
- Duffie, Darrell, and Antje Berndt, 2011, Explaining corporate debt returns, Working paper, Stanford University.
- Duffie, Darrell, and Bruno Strulovici, 2011, Capital mobility and asset pricing, Working paper, Stanford University and Northwestern University.
- Eichenbaum, Martin S., Lars Peter Hansen, and Kenneth J. Singleton, 1988, A time series analysis of representative agent models of consumption and leisure choice under uncertainty, *Quarterly Journal of Economics* 103, 51–78.
- Epstein, Larry G., and Stanley Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.
- Fama, Eugene F., 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383–417.
- Fama, Eugene F., 1984, Forward and spot exchange rates, *Journal of Monetary Economics* 14, 319–338.
- Fama, Eugene F., 1986, Term premiums and default premiums in money markets, *Journal of Financial Economics* 17, 175–196.
- Fama, Eugene F., 1996, Multifactor portfolio efficiency and multifactor asset pricing, *Journal of Financial and Quantitative Analysis* 31, 441–465.
- Fama, Eugene F., 1998, Market efficiency, long-term returns, and behavioral finance, *Journal of Financial Economics* 49, 283–306.
- Fama, Eugene F., and Robert R. Bliss, 1987, The information in long-maturity forward rates, *American Economic Review* 77, 680–692.
- Fama, Eugene F., and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.
- Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stock and bonds, *Journal of Financial Economics* 33, 3–56.

- Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55–84.
- Fama, Eugene F., and Kenneth R. French, 1997, Industry costs of equity, *Journal of Financial Economics* 43, 153–193.
- Fama, Eugene F., and Kenneth R. French, 2006, The value premium and the CAPM, *Journal of Finance* 61, 2163–2185.
- Fama, Eugene F., and Kenneth R. French, 2010, Dissecting anomalies, *Journal of Finance* 63, 1653–1678.
- Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig, 2010, Why does the Treasury issue TIPS? The TIPS-Treasury bond puzzle, Working paper, UCLA Anderson School of Management.
- Fontana, Alessandro, 2010, The persistent negative CDS-bond basis during the 2007/08 financial crisis, Working paper, Università Ca Foscari, Venice.
- Frazzini, Andrea, and Lasse Hege Pedersen, 2010, Betting against beta, Working paper, New York University.
- Froot, Kenneth A., and P. O'Connell, 2008, On the pricing of intermediated risks: Theory and application to catastrophe reinsurance, *Journal of Banking and Finance* 32, 69–85.
- Gabaix, Xavier, Arvind Krishnamurthy, and Olivier Vigneron, 2007, Limits of arbitrage: Theory and evidence from the mortgage backed securities market, *Journal of Finance*, 62, 557–596.
- Gala, Vito, 2010, Irreversible investment and the cross section of stock returns in general equilibrium, Working paper, London Business School.
- Garber, Peter M., 2000, *Famous First Bubbles: The Fundamentals of Early Manias* (MIT Press, Cambridge).
- Garleanu, Nicolae, and Lasse Heje Pedersen, 2011, Margin-based asset pricing and deviations from the law of one price, *Review of Financial Studies* 24, 1980–2022.
- Gomes, Joao, Leonid Kogan, and Lu Zhang, 2003, Equilibrium cross-section of returns, *Journal of Political Economy* 111, 693–732.
- Gourinchas, Pierre-Olivier, and Hélène Rey, 2007, International financial adjustment, *Journal of Political Economy* 115, 665–703.
- Gourio, Francois, 2007, Labor leverage, firms' heterogeneous sensitivities to the business cycle, and the cross-section of returns, Working paper, Boston University.
- Gromb, Denis, and Dimitri Vayanos, 2010, Limits of arbitrage: The state of the theory, *Annual Review of Financial Economics* 2, 251–275.
- Hansen, Lars Peter, John C. Heaton, and Nan Li, 2008, Consumption strikes back? Measuring long-run risk, *Journal of Political Economy* 116, 260–302.
- Hansen, Lars Peter, and Robert J. Hodrick, 1980, Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy* 88, 829–853.
- Hansen, Lars Peter, and Thomas J. Sargent, 2005, Recursive robust estimation and control under commitment, *Journal of Economic Theory* 124, 258–301.
- Hayek, Friedrich A. Von, 1945, The use of knowledge in society, *American Economic Review* 35, 519–530.
- Heaton, John C., and Deborah J. Lucas, 2000, Portfolio choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance* 55, 1163–1198.
- Heaton, John C., Deborah J. Lucas, and Robert L. McDonald, 2009, Is mark-to-market accounting destabilizing? Analysis and implications for policy, *Journal of Monetary Economics* 57, 64–75.
- He, Zhiguo, and Arvind Krishnamurthy, 2010, Intermediary asset pricing, Working paper, University of Chicago and Northwestern University.
- Hu, Xing, Jun Pan, and Jiang Wang, 2011, Noise as information for illiquidity, NBER Working Paper 16468.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.
- Jermann, Urban, 2010, The equity premium implied by production, *Journal of Financial Economics* 98, 279–296.

- King, Robert G., Charles I. Plosser, and Sergio T. Rebelo, 1988, Production, growth and business cycles: I. The basic neoclassical model, *Journal of Monetary Economics* 21, 195–232.
- Koijen, Ralph, and Jules van Binsbergen, 2009, Predictive regressions: A present-value approach, *Journal of Finance* 65, 1439–1471.
- Krishnamurthy, Arvind, 2002, The bond/old-bond spread, *Journal of Financial Economics* 66, 463–506.
- Krishnamurthy, Arvind, 2008, Fundamental value and limits to arbitrage, Working paper, Kellogg Graduate School of Management.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2010, The aggregate demand for Treasury debt, Working paper, Kellogg Graduate School of Management.
- Lacerda, Filipe, and Pedro Santa-Clara, 2010, Forecasting dividend growth to better predict returns, Working paper, Universidade Nova de Lisboa.
- Laibson, David, 1997, Golden eggs and hyperbolic discounting, *Quarterly Journal of Economics* 62, 443–477.
- Lamont, Owen, 2000, Investment plans and stock returns, *Journal of Finance* 55, 2719–2745.
- Lettau, Martin, and Sydney Ludvigson, 2001a, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815–849.
- Lettau, Martin, and Sydney Ludvigson, 2001b, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238–1287.
- Lettau, Martin, and Sydney Ludvigson, 2005, Expected returns and expected dividend growth, *Journal of Financial Economics* 76, 583–626.
- Lettau, Martin, and Jessica Wachter, 2011, The term structures of equity and interest rates, *Journal of Financial Economics* 101, 90–113.
- Lewellen, Jonathan, Stefan Nagel, and Jay Shanken, 2010, A skeptical appraisal of asset pricing tests, *Journal of Financial Economics* 96, 175–194.
- Li, Erica X., Dmitry Livdan, and Lu Zhang, 2008, Anomalies, *Review of Financial Studies* 22, 4301–4334.
- Liu, Laura Xiaolei, Toni M. Whited, and Lu Zhang, 2009, Investment-based expected stock returns, *Journal of Political Economy* 117, 1105–1139.
- Liu, Laura Xiaolei, and Lu Zhang, 2011, Investment-based momentum profits, Working paper, Hong Kong University of Science and Technology and Ohio State University.
- Longstaff, Francis A., 2004, The flight to liquidity premium in U. S. Treasury bond prices, *Journal of Business* 77, 511–526.
- Longstaff, Francis A., 2008, Flight-from-leverage in distressed financial markets, Working paper, UCLA Anderson School of Management.
- Ludvigson, Sydney, 2011, Advances in consumption-based asset pricing: Empirical tests, in George Constantinides, Milton Harris, and Rene Stulz, eds. *Handbook of the Economics of Finance, Volume 2* (Elsevier, Amsterdam).
- Lustig, Hanno, Nikolai L. Roussanov, and Adrien Verdelhan, 2010a, Common risk factors in currency markets, Working paper, UCLA, MIT, and Wharton School.
- Lustig, Hanno, Nikolai L. Roussanov, and Adrien Verdelhan, 2010b, Countercyclical currency risk premia, Working paper, UCLA, MIT, and Wharton School.
- Mankiw, N. Gregory, and Stephen Zeldes, 1991, The consumption of stockholders and non-stockholders, *Journal of Financial Economics* 29, 97–112.
- Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.
- Merton, Robert C., 1971, Optimum consumption and portfolio rules in a continuous time model, *Journal of Economic Theory* 3, 373–413.
- Milgrom, Paul, and Nancy Stokey, 1982, Information, trade and common knowledge, *Journal of Economic Theory* 26, 17–27.
- Mitchell, Mark, Lasse Heje Pedersen, and Todd Pulvino, 2007, Slow-moving capital, *American Economic Review* 97, 215–220.
- Mitchell, Mark, and Todd Pulvino, 2001, The characteristics of risk and return in risk arbitrage, *Journal of Finance* 56, 2135–2176.

- O'Hara, Maureen, 2008, Bubbles: Some perspectives (and loose talk) from history, *Review of Financial Studies* 21, 11–17.
- Pastor, Lubos, 2000, Portfolio selection and asset pricing models, *Journal of Finance* 55, 179–223.
- Pastor, Lubos, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.
- Piazzesi, Monika, and Eric Swanson, 2008, Futures prices as risk-adjusted forecasts of monetary policy, *Journal of Monetary Economics* 55, 677–691.
- Rietz, Thomas A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.
- Ross, Stephen A., 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–360.
- Ross, Stephen A., 1978, A simple approach to the valuation of risky streams, *Journal of Business* 51, 453–475.
- Sargent, Thomas J., and Neil Wallace, 1981, Some unpleasant monetarist arithmetic, *Federal Reserve Bank of Minneapolis Quarterly Review* 5, 1–17.
- Scheinkman, José A., and Wei Xiong, 2003, Overconfidence and speculative bubbles, *Journal of Political Economy* 111, 1183–1219.
- Shiller, Robert, 1981, Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421–436.
- Stambaugh, Robert, 1988, The information in forward rates: Implications for models of the term structure, *Journal of Financial Economics* 21, 41–70.
- Vayanos, Dimitri, and Jean-Luc Vila, 2011, A preferred-habitat model of the term structure of interest rates, Working paper, London School of Economics.
- Vayanos, Dimitri, and Jiang Wang, 2011, Liquidity and asset prices: A unified framework—Part II, Working paper, London School of Economics and MIT.
- Vuolteenaho, Tuomo, 2002, What drives firm-level stock returns? *Journal of Finance* 57, 233–264.
- Wachter, Jessica, 2003, Risk aversion and allocation to long-term bonds, *Journal of Economic Theory* 112, 325–333.
- Welch, Ivo, 2004, Capital structure and stock returns, *Journal of Political Economy* 112, 106–131.
- Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.
- Yogo, Motohiro, 2006, A consumption-based explanation of expected stock returns, *Journal of Finance* 61, 539–580.