

Expected Return, Realized Return, and Asset Pricing Tests

EDWIN J. ELTON*

ONE OF THE FUNDAMENTAL ISSUES in finance is what the factors are that affect expected return on assets, the sensitivity of expected return to those factors, and the reward for bearing this sensitivity. There is a long history of testing in this area, and it is clearly one of the most investigated areas in finance.

Almost all of the testing I am aware of involves using realized returns as a proxy for expected returns. The use of average realized returns as a proxy for expected returns relies on a belief that information surprises tend to cancel out over the period of a study and realized returns are therefore an unbiased estimate of expected returns. However, I believe that there is ample evidence that this belief is misplaced. There are periods longer than 10 years during which stock market realized returns are on average less than the risk-free rate (1973 to 1984). There are periods longer than 50 years in which risky long-term bonds on average underperform the risk free rate (1927 to 1981).¹ Having a risky asset with an expected return above the riskless rate is an extremely weak condition for realized returns to be an appropriate proxy for expected returns, and 11 and 50 years is an awfully long time for such a weak condition not to be satisfied. In the recent past, the United States has had stock market returns of higher than 30 percent per year while Asian markets have had negative returns. Does anyone honestly believe that this is because this was the riskiest period in history for the United States and the safest for Asia? Furthermore, there is a large body of evidence we find anomalous. This includes the effect of inflation on asset pricing and the failure of the generalized expectation theory to explain term premiums. Changing risk premiums and conditional asset pricing theories may be a way of "explaining" some of the anomalous results; however, this does not explain returns on risky assets that are less than the riskless rate for the long periods when it has occurred.

It seems to me that the more logical explanation for these anomalous results is that realized returns are a very poor measure of expected returns and that information surprises highly influence a number of factors in our

*Nomura Professor of Finance, New York University. A special thanks to Martin Gruber for a lifelong friendship and productive collaboration. Pierluigi Balduzzi, Steve Brown, and Matt Richardson were especially helpful on this manuscript. Thanks to Deepak Agrawal for computational assistance and thoughtful comments. I also thank Yakov Amihud, Anthony Lynch, Jennifer Carpenter, Paul Wachtel, and Cliff Green for their comments and help. As always, none of the aforementioned are responsible for any opinions expressed or any errors.

¹See Ibbotson (1995) and compare long-term government bond returns to T-bills returns.



Edwin J. Elton
President of the American Finance Association
1998

asset pricing model. I believe that developing better measures of expected return and alternative ways of testing asset pricing theories that do not require using realized returns have a much higher payoff than any additional development of statistical tests that continue to rely on realized returns as a proxy for expected returns. I illustrate what I have in mind by examining the expected return on Government bonds, then I present some preliminary thoughts in the common stock area.

Government bonds are assets for which I believe we can obtain much better estimates of expected return from realized return. Government bonds have little asset-specific information that should affect their price.² Rather, the factors that affect the prices of Government bonds are in the form of aggregate economic information. There is wide consensus on which economic variables could affect the prices of Government bonds. Variables that are considered potentially important are common across the major players.³ Furthermore, the time when information about these variables is announced is known and fixed. Finally, the impact of the surprise component in announcements is rapidly incorporated into prices (see Balduzzi et al. (1997)). The combination of a common information set, known announcement time, and rapid reflection of information into price allows us to put together a data set unaffected by information surprises, and provides a unique opportunity to put together a set of reasonably accurate estimates of expected returns. This data set can then be used to examine some of the hypotheses about what affects expected return on Government bonds. In the first part, I explore some of these ideas. In the second part, I make more speculative comments in the common stock area. Before doing either, however, I would like to expand on my prior thoughts briefly.

I. An Overview

Before examining some applications of the ideas, it is useful to further explore the basic idea. We can think of returns as being decomposed into expected returns and unexpected returns. More formally,

$$R_t = E_{t-1}(R_t) + e_t, \quad (1)$$

where R_t is return in period t , $E_{t-1}(R_t)$ is expected return at t conditional on information available at $t - 1$, and e_t is unexpected return.

In the common stock area, unexpected return is viewed as coming from systematic factors or unique firm specific events. In the government bond area, unexpected returns result from surprises in the macroeconomic announcements. The hope in using realized return as a proxy for expected

² The exceptions are liquidity effects and tax effects that are asset-specific. However, these are very small in magnitude and relatively constant over time (see Elton and Green (1998)).

³ However, which of these variables actually affect prices is a serious research topic. See Balduzzi, Elton, and Green (1997) for a careful analysis of this topic.

return is that the unexpected returns are independent, so that as the observation interval increases they tend to a mean of zero. What I am arguing is that either there are information surprises that are so large or that a sequence of these surprises is correlated so that the cumulative effect is so large that they have a significant permanent effect on the realized mean. Furthermore, these surprises can dominate the estimate of mean returns and be sufficiently large that they are still a dominant influence as the observation interval increases. Thus, the model I have in mind is

$$R_t = E(R_{t-1}) + I_t + \epsilon_t, \quad (2)$$

where I_t is a significant information event.

I view I_t as mostly zero and occasionally a very large number. Thus I view the ϵ_t as a mixture of two distributions, one with standard properties and the other that more closely resembles a jump process. Let me discuss some examples. When I first entered the profession, those of us who looked at the efficient frontiers talked about the McDonalds Effect. Any data set that included McDonalds showed extremely large returns for very little risk. The use of McDonalds as an input to an efficient frontier produced portfolios that consisted almost exclusively of McDonalds and were simply not credible. What was going on? For the first few years, no one anticipated the size of the earnings that McDonald announced, and every time earnings were announced the stock soared dramatically. In the formality above, I_t was very large at earnings announcements and highly correlated. Estimates of the sensitivity were very low given the large shocks. Adding other observations would dampen the impact, but never eliminate it as the dominant feature of the data. As I discuss later, including McDonalds in asset pricing tests causes similar difficulties to using it as an input to the efficient frontier. Earnings surprises for companies in their early years is an obvious example of what I have in mind. There are a lot of other examples with more mature companies—Atlantic Richfield with North Shore Oil, Pfizer with the announcement of Viagra causing a large one-time return that would result in problems for asset pricing tests. Corporate restructuring, such as being acquired, should have a similar effect.

There are similar problems with market returns or factor returns. In the introduction, I mention that long-term government bonds earned less than T-bills over a 50-year period. There is a subsequent 5-year period during which the returns total 112 percent and include yearly returns of 40 percent, 31 percent, 25 percent, and 15 percent. This was a period in which inflation was brought under control and investors were continually surprised at the changes. Even averaging over 20 years, this period would still dominate the estimate of expected returns.⁴

⁴ For those who view return as stationary and see this as a sampling problem I note that the highest previous return was 16 percent. What are the odds of getting these returns with the historical distribution?

As a second example, the Japanese stock market had a rise of 20 percent per year from 1980 to 1990 as the Japanese economy continued to grow faster than people expected. A series of papers show how historical data could be used to select an international portfolio that dominates holding a U.S.-only portfolio since any use of historical data would overweight Japan and lead to high future returns. Likewise, international asset pricing tests have trouble if Japan is included given its large realized return and relative independence from the world market.

What are we to do? One answer is to try to remove the I_t 's by observing the announcement and adjusting for the surprise. Alternatively, we could develop econometric techniques for identifying the I_t without observing the announcement and then eliminating it. A different direction is to try to build expectations directly into our asset pricing models. The very old monograph by Meiselman (1962) is a very interesting attempt in this direction. Papers by Fama and Gibbons (1984) and Froot (1989) have some relationship to what I have in mind. Finally, we could test models by seeing if they provide a useful tool for decision making. Pastor and Stambaugh (1999) have made an attempt in this direction. To illustrate these ideas in more detail, it is useful to examine government bond returns.

II. The Data

This section describes the data set used in the empirical analysis: the GovPX bond price data and the MMS forecast survey data.

A. Price Data

The data set used to calculate expected returns is provided by GovPX. The data set we used contains bid and ask quotes, the price of the last trade as of 3:00 or 6:00 p.m. EST, and trading volume in the interdealer broker market for all Treasury bills, notes, and bonds. The data set covers the period from July 1, 1991 through December 31, 1997.⁵

The cash market for Treasury securities is much more active than the futures market. For example, during March to May 1993, dealer transactions in the futures market are only about 18 percent of the volume in the cash (Federal Reserve Bulletin (1993)). Likewise, within the cash market the majority of the trades are in the inner market—that is, trades among dealers. According to the same Federal Reserve Bulletin, approximately 62 percent of the March to May 1993 Treasury security transactions in the secondary market occurred within the inner market. Treasury dealers trade with one another mainly through intermediaries, called interdealer brokers. Six of the seven main interdealer brokers, representing approximately 75 percent of all quotes and a much higher percentage of the maturities we exam-

⁵ The time is 6:00 p.m. for data up to September 1996, and 3:00 p.m. for data beyond that date.

ine, provide price information to the firm GovPX.⁶ In turn, GovPX provides price, volume, and quote information to all of the Treasury bond dealers and to other traders through financial news providers, such as Bloomberg.

Dealers leave firm quotes with the brokers, and GovPX shows the best bid and ask for each bond along with the largest size the quote is good for. Thus, the posted quotes are also the prices at which actual trading takes place.⁷

B. Survey and Announcement Data

The data on economic announcements and expectations are from Money Market Services (MMS), a San Francisco-based corporation that has conducted telephone surveys since late 1977. The MMS data are the most commonly used expectational data in studies of economic announcements. Balduzzi, Elton, and Green (1998), Edison (1996), Hakkio and Pearce (1985), Ito and Roley (1991), Hardouvelis (1988), McQueen and Roley (1993), and Ulrich and Wachtel (1984) are some of the many studies that have used the MMS data.⁸ The MMS data show the median forecast, the magnitude of the actual announcement, and the time of the announcement. Thus, the error in the consensus forecast can be calculated as well as the exact time this error was known by the market.⁹

III. Getting Expected Return

In order to obtain estimates of expected return, we first need to construct a realized return series on zero coupon bonds of different maturities. Using zero coupon bonds means that each return is attributed uniquely to a specific maturity. We estimate prices of zero coupon bonds. Each day we estimated the term structure of spot rates by running a cubic spline as described in Elton and Green (1998). With the GovPX database, the average error in price is about 10¢ per \$100 and is random across time and maturity. The term structures are estimated using the last trade price for most bonds in the database.¹⁰ We then calculate the price of a zero coupon bond with maturities of six months, one year, 18 months, and so on up to five years. Daily returns are then computed by taking the log of the price relatives for adjacent days.¹¹

⁶ The exception is Cantor Fitzgerald, which deals almost exclusively at the long end. Thus, the percentage of all quotes present in the GovPX data for the range of maturities we are examining is much higher than 75 percent.

⁷ See Balduzzi, Elton, and Green (1997) for a more detailed description of the GovPX data and Elton and Green (1998) for a discussion of its accuracy relative to other sources.

⁸ See Balduzzi, Elton, and Green (1997) for a more detailed discussion of the properties of these data.

⁹ During the budget crisis at the end of 1996, surveys were not undertaken for two variables. Thus, we do not have errors on one date for two variables. This is an insignificant omission.

¹⁰ We exclude bonds that have not traded for a while, flower bonds, inflation linked and callable bonds, and any bond with a large pricing error.

¹¹ Obviously, after one day the bond has one less day to maturity and this is adjusted for.

Table I
Estimates of Expected Returns

This table shows the results of regressions of daily realized returns on important information surprises for that day. One regression is run for each maturity listed below over the full sample period (July 1991 to December 1997). Expected Return is the estimate of the intercept in this regression, expressed as an annualized percentage. Adjusted R^2 of the regression represents the impact of surprise removal from realized returns.

Maturity	Expected Return	Adjusted R^2
1 month	3.095%	0.004
6 months	3.575%	0.105
1 year	4.053%	0.154
1.5 years	4.395%	0.151
2 years	4.675%	0.147
2.5 years	5.070%	0.136
3 years	5.493%	0.132
3.5 years	5.808%	0.126
4 years	6.013%	0.122
4.5 years	6.312%	0.116
5 years	5.925%	0.081

Bond prices should adjust because of new information that enters the market or because of buying or selling pressure caused by inventory adjustment of financial intermediaries. Insofar as they reflect inventory adjustment, the price changes should be temporary as arbitrageurs take positions reflecting market beliefs about long-term prices. Thus inventory adjustments should introduce random noise, and the price changes—separate from the effect of announcement surprises—should be unbiased estimates of expected returns.

We purge the data of announcement effects by removing the estimated effect of the announcement surprises from the daily returns. The superior way to do this would be to remove the information effect by using return data in the 15 minutes around the announcement. Balduzzi et al. (1997) show that all information effects are reflected in the price within this time span. Unfortunately, we do not have intraday trade data for our full sample period. Our compromise is to use daily data and only the announcements that Balduzzi et al. have shown to be important. Although we place no reliance on the relative coefficients across announcements given that multiple announcements occur on the same day and that surprises across announcements are correlated, the effect on daily return should be reasonably well estimated. The variables used are surprises in the announcement of the Consumer Price Index, Durable Goods Orders, Housing Starts, Initial Jobless Claims, Non Farm Payrolls, Producers Price Index, Retail Sales, Capacity Utilization, Consumer Confidence, NAPM index, New Homes Sales, and M2 medians. The “Expected Return” column in Table I shows the realized return with the effect of announcements removed. The returns are percent-

age returns for bonds with maturities at six-month intervals up to five years over our full sample period. This is our best estimate of expected return using realized returns. Also included in Table I is the R^2 of the regression of realized returns on announcements.¹² The returns are annualized. Information surprises in economic announcements have little effect on the return for short maturity bonds. For bonds of maturities from one to five years, there is significantly more impact, and removing the effect of announcements has a real impact on the time series of returns.¹³ Most of the equilibrium models in the bond area are continuous time models so that it seems most appropriate to examine what happens to daily returns rather than returns over longer intervals. To examine the difference that the adjustment makes, I calculate on each announcement day the absolute difference between the adjusted and the raw returns and scale it by the mean absolute unadjusted return. For more than 75 percent of the days, the percentage difference is greater than 20 percent, for 40 percent of the days it is greater than 40 percent, and for 20 percent of the days, it is greater than 60 percent. As shown in Table I, the adjustment also affects the standard deviation of daily returns. Furthermore, when we remove the effect of information surprises, kurtosis—which is significant for the unadjusted data—is no longer significant.

IV. Tests in the Bond Area

Modern bond research views bond returns as composed of a return common to all bonds which reflects the return due to the passage of time and a risk premium for bearing term structure risk. Generally, empirical evidence and modern bond pricing theory suggest that the risk premium changes over time as a function of state variables, but we do some preliminary tests assuming a constant term premium.

A. A Constant Risk Premium

If the risk premium is constant, then the expected return should equal the one-month interest rate plus a term premium related to the maturity of the bond. If term premiums are defined per period rather than over the full life of the bond, then the difference between a bond's return (with the effect of surprises removed) and the riskless rate is simply the one-period term premium for that maturity bond. In equation form this is

$$R_{it} - R_{ft} = P_{it}, \quad (3)$$

¹² Balduzzi et al. (1997) found that several announcements accounted for more than 50 percent of the within-day volatility. We find a much lower R^2 . We are explaining daily volatility and many days do not have announcements.

¹³ There are many issues that I have ignored. These concern how well the median of the forecast measures the expected announcement, errors introduced because of imperfect measurement, and the best structural form for the effect of the announcement on return.

Table II
Estimates of Term Premiums

This table shows the estimates of term premiums, computed as the average over the sample period of the difference between the expected returns on the bonds of a given maturity and the riskless rate. The expected returns on bonds were constructed by removing the effect of information surprises from the realized returns. Returns on a constant maturity one-month bill are used for the riskless rate. The premiums are in annualized percentages.

Maturity	Premium
6 months	0.471%
1 year	0.937%
1.5 years	1.272%
2 years	1.547%
2.5 years	1.935%
3 years	2.352%
3.5 years	2.660%
4 years	2.855%
4.5 years	3.143%
5 years	2.795%

where R_{it} is the return on a bond of maturity i in period t with the effect of information surprises removed, P_{it} is the t th period term premium for a bond of maturity i , and R_{ft} is the riskless rate in period t .

If term premiums are constant over time then a simple average of the monthly excess returns should provide an estimate of the term premium and the i th term premium is $\bar{P}_i = \sum P_{it}/T$.

Table II shows the estimate of the term premium assuming a constant term premium over the period from July 1991 to December 1997. As the riskless rate we use the return on a constant maturity one-month bill.¹⁴ The estimates of the term premium are generally positive and increase in maturity.

One of the implications of the generalized expectations theory is that forward rates are expected future spot rates plus term premiums. It is useful to examine how term premiums estimated from forward rates compare to the term premiums just calculated.

B. Forward Rates and Risk Premiums

If we let F represent forward rates, then¹⁵

$$F_{t,t+1,t+J} = \sum_{K=t+1}^{t+J-1} \bar{R}_{J-K,K} + P_{J-t-1,t+1}, \quad (4)$$

¹⁴ A very similar method of estimating risk premiums is to regress excess return on surprises and use the intercept as the estimate of the risk premium. Since one does not expect correlation between the riskless rate and the surprises, this is equivalent to the procedure just described.

¹⁵ The definition of the term "premium" is slightly different here. In this case, it is the premium over the full period. Thus, these premiums should be averages of those shown in Table II.

Table III
Estimates of Term Premiums
Using Forward Rates

This table shows the estimates of term premiums computed as the difference between the forward rate and the expected future spot rate. The expected future spot rates are computed by summing the expected future returns, as in equation (2). The premiums are in annualized percentages.

Maturity	Premium
6 months	1.412%
1 year	1.581%
1.5 years	1.730%
2 years	1.890%
2.5 years	2.046%
3 years	2.183%
3.5 years	2.268%
4 years	2.288%
4.5 years	2.283%

where $\bar{F}_{t,t+1,t+J}$ is the forward rate from $t + 1$ to $t + J$ as of time t , and the overbar indicates expected value.

Rearranging, we have

$$P_{J-t-1,t+1} = \bar{F}_{t,t+1,t+J} - \sum_{K=t+1}^{t+J-1} R_{J-K,K}. \quad (5)$$

Estimates of annualized term premiums using forward rates are contained in Table III. Note that estimates of the term premium using forward rates have the same pattern of increasing with maturity as those obtained earlier and at least at the long end have the expected magnitude. One way to examine the reasonableness of the constant risk premiums is to see if they are related to factors that explain the structure of returns over time. To examine this, we initially need to examine what factors affect government bond returns.

C. Factor Analysis

To get an idea of the return structure, we factor analyze the variance/covariance matrix of returns. There is some history of this in bond research (see Elton, Gruber, and Michaely (1990) and Litterman and Scheinkman (1991)). Both papers factor analyze the return series on coupon-paying bonds. Since the return on a coupon bond is a weighted average of returns on zeros, using the return on coupon bonds induces correlation among bond returns even when there would be none if returns on pure discount bonds were used. Therefore, we perform a maximum likelihood factor analysis on the daily return series of the pure discount bonds. In this case we use the daily return series with the effect of surprises included. Since the biggest changes are

Table IV
Explanatory Power of the Factors

This table shows the adjusted R^2 in the regression of returns on factor portfolios for one- to four-factor solutions. The last column shows the explanatory power of a prespecified two-factor model where the two factors are the returns on the six-month and four-year bonds.

Maturity	One Factor	Two Factors	Three Factors	Four Factors	Prespecified Two Factors
1 month	0.012	0.406	0.450	0.551	0.352
6 months	0.472	1.000	1.000	1.000	1.000
1 year	0.831	0.908	0.971	0.987	0.860
1.5 years	0.905	0.916	0.981	0.994	0.852
2 years	0.962	0.965	0.984	0.989	0.888
2.5 years	0.953	0.955	0.944	0.999	0.871
3 years	0.987	0.985	0.977	0.977	0.907
3.5 years	0.970	0.966	0.976	0.969	0.942
4 years	0.942	0.940	0.976	0.988	1.000
4.5 years	0.917	0.916	0.951	0.970	0.959
5 years	0.599	0.599	0.604	0.616	0.608

likely to be associated with surprises, this gives us a better chance of uncovering the structure. After factor analyzing the series, we then regress each return series three months, six months, one year, etc. on the factors.

Table IV shows the adjusted R^2 of a regression of returns of the various maturity zeros on factor portfolios for one- to four-factor solutions. The first factor captures much of the return pattern for bonds that have a maturity of two years or more.¹⁶ Examining the composition of the first factor shows that it is one-third in the three-year bond and approximately 10 percent in each of the six bonds closest to it. Examining R^2 for the shorter maturity bonds shows that they seem to move independently of the longer maturity. When a second factor enters, the R^2 for the longer maturity bonds hardly changes. However, including two factors causes the R^2 on the short end to jump dramatically. Examining the composition of the second factor, it is exclusively the six-month bond. This suggests again that the returns on bonds of different maturities are affected by at least two factors. The last column shows the explanatory power for a prespecified two-factor model where the two factors are the return on the four-year and six-month bond. The factor analytical solution must better explain the return pattern in sample; but having recognized this, the prespecified two factors do almost as well. Whether a third factor is important is unclear. It leads to some improvement and some deterioration. Examining the composition of the factor shows that it is generally long the short maturity bonds and short the long maturities, with the heaviest weights being in the 1.5- and 4-year bonds. However, the pattern is not monotonic, suggesting that it is primarily picking up idiosyn-

¹⁶ The lower explanatory power on the five-year bond is caused by a poor fit of the spline for the long end on a few dates. The poor fit causes a high return on one date and a large negative return on the adjacent one and causes returns to be poorly aligned on a couple of dates.

Table V
Relationship between the Bond Factors and the Stock Factors

This table shows the results of regression of bond factors (extracted by factor analysis) on the Connor–Korajczyk stock factors (*ck1* to *ck5*). Models 1-*FM*, 2-*FM*, and 3-*FM* are one-, two-, and three-factor models, respectively. Numbers in parentheses are *t*-statistics.

Model	Factor	Constant	<i>ck1</i>	<i>ck2</i>	<i>ck3</i>	<i>ck4</i>	<i>ck5</i>	R^2
1- <i>FM</i>	<i>f1</i>	-2.82	259.52	38.61	326.38	-58.44	-199.01	0.011
		-(0.35)	(0.89)	(0.19)	(2.02)	-(0.35)	-(1.11)	
2- <i>FM</i>	<i>f1</i>	-1.19	73.11	-7.74	42.49	-8.91	-36.27	0.000
		-(0.70)	(1.19)	-(0.18)	(1.25)	-(0.25)	-(0.96)	
	<i>f2</i>	-2.91	263.99	38.89	332.02	-59.26	-204.02	0.010
		-(0.35)	(0.88)	(0.19)	(2.01)	-(0.34)	-(1.11)	
3- <i>FM</i>	<i>f1</i>	-1.19	73.11	-7.74	42.49	-8.91	-36.27	0.000
		-(0.70)	(1.19)	-(0.18)	(1.25)	-(0.25)	-(0.96)	
	<i>f2</i>	-2.96	256.09	32.37	320.83	-60.93	-203.33	0.008
		-(0.36)	(0.86)	(0.16)	(1.96)	-(0.36)	-(1.11)	
	<i>f3</i>	-13.55	-522.19	-685.94	-699.69	-119.39	-370.25	-0.028
		-(0.38)	-(0.40)	-(0.76)	-(0.96)	-(0.16)	-(0.46)	

cratic return. When we introduce four or more factors, the composition of the factor is heavily weighted in one maturity, indicating that two or three factors capture everything systematic.

What are these factors? One thing they are *not* is related to characteristics of the one-month rate. For example, regressing the three factors on volatility of the one-month rate shows no significance. Looking at the composition of the factors suggests they are a short rate and a long rate. Regressing the two factors on the four-year and six-month bond produces R^2 of 0.94 and 1.0 (since one factor is the six-month bond).

Theory suggests that what matters are the factors that are common across asset categories; after all, major investors hold all categories of assets. Asset pricing models in the bond area and the common stock area have developed almost completely independently. In the common stock area the asset pricing model is derived either from the efficiency of some portfolio or from an assumed return-generating process where the factors in the return-generating process are portfolios. In the bond area the fundamental building block is usually an interest rate series. It is unusual for anyone testing in one area to use results from the other asset class. The few attempts that exist have been primarily adding bonds to the S&P Index or an interest rate factor to a multifactor return-generating process when testing asset pricing models for common stocks.

To examine if there is a commonality between the bond and stock factors, we regress our factors extracted from bond returns on the Connor–Korajczyk factors derived using principal component analysis from common stocks.¹⁷ The results are shown in Table V. None of the factors that explain

¹⁷ Robert Korajczyk generously supplied the monthly factor returns.

government bond returns are related to any of the Connor–Korajczyk factors derived from the returns on common equities. This seeming independence very much simplifies the asset allocation decision.¹⁸

Standard asset pricing theory suggests that expected return should be related to the sensitivity to the factors. If the sensitivity is unchanged over time, then expected return should be linearly related to the factor loadings. To test this, we employ the standard maximum likelihood ratio test (see Campbell, Lo, and MacKinlay (1997)). This test involves comparing the determinant of the variance covariance matrix of the residuals when the intercept of the return-generating process is unconstrained and when the intercept is constrained to fit a linear asset pricing model. There is a difficulty with our sample (the Lindley (1957) problem). Since we have so many time series observations relative to the number of different maturities rejection even with no economic significance is highly likely. Therefore, not surprisingly, at standard levels of significance, we can reject the hypothesis that the constrained model is not different from the unconstrained. Examining the variance covariance matrix of the constrained and unconstrained solution and noting that they are identical to the seventh place, suggests that the Lindley problem is serious here. If this is the explanation, then we should observe that the constrained solution (the asset pricing model) yields estimates of expected return that are not different from our estimate in the unconstrained solution and this is the case. Comparing the expected return of the constrained and unconstrained model shows only one difference more than four percent and that difference is 17 percent. Thus, although there is a statistical difference, there is not an economic difference in the constrained and unconstrained solutions and the differences in expected return are explained by differences in sensitivity to the factors.

D. Changing Risk Premiums

The evidence on term premiums obtained from examining the relationship between futures rates and spot rates indicates that, at least for short maturity bonds, term premiums are changing.¹⁹ The evidence primarily consists of results from regressions of forward rates on one-period spots. These results consist of estimates of the slope less than one which remain unexplained by known biases in the estimates.

Modern bond pricing models imply that the risk premium should be changing as a function of maturity and as a function of changes in the level of state variables. Single-factor models are particularly easy to analyze.

¹⁸ A caution is in order. Bond indexes including corporate return series are related to stock returns. See Elton, Gruber, and Blake (1999). The reason for this is likely to be that the term premium on corporates is related to stock market factors. See Elton, Gruber, Agrawal, and Mann (1998).

¹⁹ See, for example, Backus et al. (1997) and Campbell and Shiller (1991). For a different opinion see Froot (1989).

In a single-factor bond pricing model the state variable is normally the short-term rate. In a single-factor model like that of Cox, Ingersoll, and Ross (1985), the sign of the term premium is not affected by maturity and the magnitude of the term premium is linear in the short rate. Multifactor models are more complex because the sign of the term premium can change as a function of maturity. In the two-factor model of Longstaff and Schwartz (1992), the magnitude of the term premium is a linear function in the short rate. However, the effect of increases in the short rate on the term premium is less clear. It must always be a positive function for short maturity bonds, but the effect of an increase is indeterminate for long maturity bonds. In the Longstaff and Schwartz model, the second state variable is the variance of the change in the short rate. Once again the term premium is linear in the variance, but the sign is indeterminate.

The other type of two-factor model examined is that where the state variables are the long and the short rate.²⁰ In this case the term premium is again linear in the state variables. In general, in all of these models the magnitude of the effect on the term premiums of changes in the state variables decreases with maturity due to mean reversion. Also, the variability of the estimated risk premium generally increases with maturity. Together these suggest the following regression:

$$R_{it} - R_{Ft} = c_0 + c_1 r_s + c_2 \sigma_s^2 + c_3 r_L + e_{it}, \quad (6)$$

where r_s is the short rate, r_L is the long rate, and σ_s^2 is the variance in the change in the short rate.

The expected results are

1. c_1 is nonzero and has the same sign across maturities and the change across maturities is monotonic (single-factor CIR).
2. c_1 is nonzero and is positive for short maturity bonds, and c_2 is nonzero (Longstaff and Schwartz (1992)).
3. c_1 and c_3 are nonzero (Brennan and Schwartz (1982)).
4. The variance of the residual increases with maturity.

We need to specify how the variables are calculated. Expected return is as discussed earlier. Like others, we use the one-month rate at the beginning of the month for the short rate. For the long-term bond, we select the spot rate on the five-year bond at the beginning of the month. Finally, for the estimate of the variance of the change in the one-month rate we follow Longstaff and Schwartz (1992) and estimate it using the following GARCH process:

$$r_{Ft+1} - r_{Ft} = \alpha_0 + \alpha_1 r_{Ft} + \alpha_2 V_t + \delta_{t+1},$$

²⁰ See, for example, Brennan and Schwartz (1982). The models can also be expressed as the short rate and a difference in the long and short rate. See Elton, Gruber, and Mei (1996) for an example.

Table VI
Sensitivities of Term Premiums to State Variables

This table shows the results of the regression: $R_{it} - R_{ft} = C_0 + C_1 r_s + C_2 \sigma_s^2 + C_3 r_L + e_{it}$, where r_s is the short rate, r_L is the long rate, and σ_s^2 is the variance of the short rate. Numbers in parentheses are the t -statistics.

Maturity	C_0	C_1	C_2	C_3	Adjusted R^2
1 month	0.003 (1.05)	-0.246 (-5.08)	11.256 (0.29)	-0.099 (-1.63)	0.428
6 months	-0.016 (-1.46)	-0.151 (-0.92)	82.109 (0.62)	0.205 (0.99)	-0.013
1 year	-0.059 (-2.14)	-0.147 (-0.36)	320.479 (0.97)	0.902 (1.75)	0.021
1.5 years	-0.110 (-2.37)	-0.190 (-0.27)	522.438 (0.94)	1.763 (2.02)	0.038
2 years	-0.156 (-2.40)	-0.246 (-0.25)	598.096 (0.77)	2.567 (2.10)	0.041
2.5 years	-0.196 (-2.38)	-0.351 (-0.29)	656.175 (0.67)	3.338 (2.16)	0.042
3 years	-0.242 (-2.43)	-0.545 (-0.37)	863.549 (0.72)	4.236 (2.26)	0.047
3.5 years	-0.287 (-2.44)	-0.755 (-0.43)	1054.509 (0.75)	5.111 (2.31)	0.050
4 years	-0.319 (-2.35)	-0.897 (-0.44)	1078.575 (0.66)	5.750 (2.26)	0.044
4.5 years	-0.351 (-2.37)	-0.896 (-0.40)	1167.716 (0.66)	6.301 (2.26)	0.045
5 years	-0.370 (-2.18)	-1.002 (-0.39)	820.403 (0.40)	6.716 (2.10)	0.031

where

$$\delta_{t+1} = N(0, V_t)$$

and

$$V_t = \beta_0 + \beta_1 R_{ft} + \beta_2 V_{t-1} + \beta_3 \delta_t^2. \quad (7)$$

The results of the regression are shown in Table VI. The one-month rate and the variance of the one-month rate are insignificant for all maturities. The long rate is significant for 1.5 to 4.5 years. These results suggest that if risk premiums do change, it is not due to changes in the characteristics of the one-month rate. There is some evidence suggesting that the long rate might serve as a state variable proxy.

As discussed above, most bond pricing models have focused on characteristics of the short rate as at least one of the state variables, and often in two-factor models, both state variables. I have always thought this was a misplaced emphasis. It seems to me that what permanently shifts the term structure are changes in the information set. Further, there is a general

belief among traders that the short-term end of the curve is anchored by Federal Reserve actions and long rates change because of changes in beliefs about the general state of the economy (which may predict future short rates). This is supported by the work in Balduzzi et al. (1997).

In this article we look at price reactions to major surprises associated with announcements of economic variables. We find that inclusion of the return on the short bond in a regression of a long bond return on surprises does not change the effect of surprises on long bond returns, while the inclusion of any intermediate or long bond captures the effect of information surprises on other long bonds. Thus the effect of information surprises on long bonds is captured by the return on a long bond, but not the return on a short bond. Likewise, inclusion of the return on a long bond does not capture the information effects on short bonds. It is hard to understand how the level or variability of the short rate captures sensitivity on the long end when the short rate does not seem to react to the same surprises and in the same way as the long rate.²¹

V. Asset Pricing Tests in the Common Stock Area

Asset pricing in the common stock area has universally involved using realized returns as a proxy for expected returns. Testing generally takes one of three forms. First is the time series testing where sensitivities and risk premiums are simultaneously estimated and the principal tests involve examining the change in explanatory power of the regression constrained to conform to the pricing model relative to the unconstrained regression. Second is the two-pass procedure where sensitivities are first estimated and then risk premiums are estimated and the principal test involves the reasonableness of the estimates of the premiums. Both of these tests can be done conditionally where sensitivities or premiums or both are allowed to be time varying. The third procedure is the test of the efficiency of the market portfolio. As shown in Campbell et al. (1997), the first and third procedures are essentially the same. The problem with using realized returns as a proxy for expected returns is prevalent in all three tests. The nature of this problem is easiest to understand with the third test. The test for the efficiency of the market portfolio involves the distance between the market portfolio and the efficient frontier.

Gibbons, Ross, and Shanken (1989) show that the test of the efficiency of the market portfolio depends on the Sharpe ratio. Knowing this helps us focus on how using realized returns as a proxy for expected returns impacts the tests. Consider a general multifactor model

$$R_{it} - R_{Ft} = \alpha_i + \sum \beta_{ij}^u I_{jt}^u + \sum \beta_{ij}^p I_{jt}^p + e_{it}, \quad (8)$$

²¹ In an earlier version of this paper, Balduzzi et al. (1997) found considerable movement in the short rate when the Federal Reserve Board made major changes with no corresponding reaction in long rates.

where α_i is the non-index-related return, R_{it} is the return on stock i in period t , R_{ft} is the return on the riskless rate in period t , β_{ij} is the sensitivity of stock i to factor j , I_{jt}^p is the excess return on a priced factor j , I_{jt}^u is the excess return on an unpriced factor j , and e_{it} is the random error term.

There are three ways that information surprises can affect the returns on security i :

1. an unanticipated firm-specific announcement where the impact is reflected in α_i
2. an unanticipated announcement that affects a priced index so that I_j^p is different than expected
3. an unanticipated announcement that affects an index in the return-generating process but is not priced, so that I_j^u is nonzero over the period being analyzed when its expected value is zero.

As we see shortly, unanticipated information surprises affect the tests of the reasonableness of the asset pricing model and tests for the number of factors. Each is discussed in turn.

A. Information Surprises and Tests of a Particular Asset Pricing Model

We know that if a multiindex model describes the variance covariance structure of returns and there are no unique returns on securities, then the efficient frontier can be constructed from a linear combination of the portfolios that replicate the indexes. Likewise, the market portfolio is a linear combination of the replicating portfolios and therefore the market is efficient. Whether the return on all factors is expected return or there are large information surprises that cause realized returns to be different from expected returns for these indexes does not affect the efficiency of the market portfolio. What does cause the market portfolio to be interior are information surprises about individual companies or missing indexes whether priced or not. Both of these cause a nonzero alpha on individual securities and cause the market to be interior.

What can be learned by realizing that the test involves comparison of Sharpe ratios? The presence of large informational surprises has two impacts on the Sharpe ratio. First, the larger the alpha, the greater the shift in the ratio and the more interior is the market. However, the increased alpha causes greater concentration and an increase in the denominator. Furthermore, the impact on the numerator will likely be much larger than the impact on the denominator. This follows because the likely effect of having securities with alphas is to have the portfolio weights be more concentrated in a few securities. Greater concentration has the same effect on portfolio variance as having fewer securities in the portfolio. However, the effect on variance as the number of securities in the portfolio increases is fairly modest for reasonable size portfolios. Thus, what is most important for changing the Sharpe ratio and rejecting the asset pricing model are a few large information surprises that have the impact of substantially increasing the nu-

erator with little impact on the denominator, rather than a lot of securities with small surprises that have much less impact on the numerator and essentially no impact on the denominator.

What do we learn from this discussion? If we are using individual company returns, we are likely to receive a substantial number of large surprises as the number of securities increases and always reject the asset pricing model with a high probability. What sort of information events might cause a substantial alpha for an individual security? Corporate restructuring as a result of being acquired is likely to be associated with a large alpha. Announcement of a successful new product or discovery such as a drug or oil field or announcement of a failure to develop a long-anticipated product is likely to be associated with a large alpha. Major earnings surprises would also likely be associated with a large alpha. In short, any announcement that would substantially change the market's perception of the future earning power of the firm would likely have a large alpha associated with it.

To get an idea of whether information surprises alone could lead to rejections of an asset pricing model using individual securities, I perform two simulation experiments. I assume a single-factor asset pricing model held exactly so that all alphas are zero absent information surprises. I then superimpose information surprises on the pricing model. The first information surprise I select is earnings. I calibrate the surprise to match the empirical findings of Bernard and Thomas (1990). If I assume a normal distribution, I find a mean of zero and a standard deviation of 0.38, which matches the effect of the quarterly earnings surprise they found. I assume that the earnings announcements are independent through time. For 10-year intervals, I calculate the distribution of surprises (how many positive and negative surprises). For simplicity, I assume that the negative and positive surprises cancel out, thus I use the net number of surprises in calculating alpha. If alpha is normally distributed with a mean of zero and standard deviation of 0.38 and surprises in announcements are independent, then the impact on alpha has a mean of zero and standard deviation 0.38 times the square root of the number of surprises that do not cancel out. As shown in Gibbons et al. (1989), the test statistic of whether the constrained model cannot be rejected is an F test with N degrees in the numerator and $T - N - 1$ in the denominator and is equal to

$$\frac{T - N - 1}{N} \left[1 - \frac{\bar{r}_M^2}{\sigma_M^2} \right]^{-1} \alpha' \Sigma^{-1} \alpha$$

and

$$\Sigma \alpha' \Sigma^{-1} \alpha = \frac{\bar{r}_P^2}{\sigma_P^2} - \frac{\bar{r}_M^2}{\sigma_M^2}, \quad (9)$$

where $\bar{r}$ is mean return, σ is standard deviation, the subscript M is market, the subscript P is the sample portfolio, T is the number of time periods, N is

Table VII
Frequency Distribution of Surprises

This table shows the empirical distribution of information surprises for 500 common stocks randomly picked from the universe of NYSE-listed stocks with returns data in the 10-year period from 1986 to 1995 taken from CRSP files. A return greater than one percent (in magnitude) and greater than three times the market return (in magnitude) is treated as an information surprise for that stock. All such surprises for a stock are cumulated over the 10-year period to find the total net surprise. The total net surprise over 10 years is divided by 10 to get the annualized surprise. The range shown is for a total of 500 annualized surprises.

Range	Frequency
–50% to –25%	2
–25% to 0%	179
0% to 25%	238
25% to 50%	64
50% to 75%	12
75% to 100%	2
100% to 125%	1
125% to 150%	1
150% to 175%	0
175% to 200%	1

the number of securities, α is a vector of alphas, and Σ is the variance covariance matrix.

I select an annual $\bar{r}_M$ equal to eight percent and σ_M^2 equal to 20 percent, roughly matching historical numbers.²²

The probability of rejecting the asset pricing model due to the information surprise given sample sizes of 10 to 100 securities varies from four percent to 18 percent. This is substantial. The second simulation assumes a discrete distribution of surprises. I select 500 stocks at random and compute the number of times daily return exceeds one percent, two percent, etc., or minus the same, conditional on this return being more than three times the market return. Since betas of -3 are not plausible, this seems to be a good measure of surprises, although it does not control for nonmarket factors. Then, for each firm, I sum the surprises over 10 years. The distribution of annualized 10-year surprises is shown in Table VII. Once again, I draw observations from this distribution for portfolios of different sizes and I compute rejection probabilities. The rejection due to information surprises varies from one to 95 percent as the sample size varies from 10 to 100. Thus, information surprises for individual securities lead to substantial rejection of the asset pricing model even when it is true.

It is difficult to remove information surprises from individual company returns by examining individual announcements for firms. Perhaps statistical techniques can be developed here. Grouping is another way to mitigate

²² For simplicity, I assume a market model holds and standard deviations on all assets are the same. I use a simplified version of Elton, Gruber, and Padberg (1976) to calculate security proportions.

the effect of information surprises on individual securities. There is, of course, a danger in grouping. Misspecified asset pricing models lead to large alphas and these should be correctly rejected. The grouping criteria involve attempting to average out information surprises and attempting to not average out non-zero alphas caused by a misspecified model. This is likely to require careful thought about alternative models. Obviously we need to be careful that we do not group on a criterion that implicitly groups on the basis of the sign and size of the information surprise. For example, if the information surprises are correlated across securities such as commodity price shocks, then the grouping has to be careful to diversify across firms that have different exposures to the shock. Likewise, using country returns and grouping by country is a problem because an information surprise involving market movements or exchange rate changes will likely impact all securities in a country. An information surprise that affects a sector, such as an oil price shock, can be eliminated by introducing an index capturing the shock to the return-generating process. This leads to a reduction in the nonzero alphas and if the information shock is substantial appears as a priced factor whether it is or is not.

B. Number of Priced Factors

Now let us consider a factor that is in the return-generating process but is not priced. Also assume there is a large information shock or a number of shocks in the same direction. If the security-specific returns are all zero, and the return-generating process includes all the factors, then the asset pricing model would not be rejected but the factor price on the unpriced factor would likely be nonzero and would be determined by the information surprise. The impact of leaving out an unpriced factor from the return-generating process is likely to be worse. Again, in any sample period, information surprises are likely to mean that the factor has a nonzero return and the intercept with the factor left out is equal to the product of the sensitivities times the information surprise. If some sensitivities are large, positive or negative, the result is large alphas. Thus the asset pricing model is rejected. A factor should be priced for one of two reasons: (1) if investors are mean variance maximizers and sensitivity to the factor ex ante affects mean returns and (2) because sensitivity to it satisfies a hedging demand either because of a changing consumption set (see Merton (1969)) or in a one-period context because investors care about sensitivity to it. Let us consider inflation as an example. Assume unanticipated inflation surprises do not meet a hedge demand. Then inflation will be priced if investors are mean variance maximizers in real returns because inflation will ex ante affect mean real returns. If investors are mean variance maximizers in nominal returns then inflation should not be priced even if it is part of the return-generating process. If it serves a hedge demand, then it also should be priced.²³

²³ This distinction and the assumptions underlying different asset pricing models that incorporate inflation are discussed in Elton, Gruber, and Rentzler (1983).

There should be many factors in the return-generating process that are useful in explaining covariance but that do not serve any hedge demands nor affect expected returns. It is likely that there will be information surprises in these factors and that the sample mean will be significantly nonzero in many periods. Thus it is difficult to separate out whether it is an unpriced factor or risk premiums are changing over time. However, this is an area where the techniques discussed with reference to bonds may well be relevant. Although it is likely to never be possible to control for information surprises for individual company events, it may well be possible to control for information surprises in factors if these factors are economic variables. One test for priced factors is to look at the excess return on the factor-replicating portfolios (see Balduzzi and Kallal (1997)). Consider again a factor like inflation. The realized return on the factor-replicating portfolios will be heavily influenced by the unanticipated announcements concerning inflation. However, these are identifiable and the return on days of information surprises could be eliminated or the effect of the surprise could be controlled for.

C. Implications

First, grouping should be part of any tests using current methodology. One reason for grouping is to eliminate company-specific surprises, and this purpose should have a dominant role in deciding on a grouping procedure. Second, traditional testing is likely to overstate the number of priced factors. Insights into which factors could be priced are likely to come as much from theory as from empirical results. Inflation is a useful example. Generally we believe that investors should look at real returns in making their investment decisions. Thus, there is a lot of theory that suggests inflation sensitivity is important in affecting investors hedging demands and should be priced. Now consider momentum. Accept for a moment the empirical evidence that momentum is related to realized returns. If there is any connection between momentum and changes in the opportunity set, I am not aware of it. Thus, momentum is the kind of factor that is likely to appear in the return-generating process and likely to appear priced in sample but for which there is no theory that would suggest that it should be priced and for which current testing procedures are unlikely to be helpful.

VI. Summary

When I first entered the profession, anyone using realized returns as expected returns made the argument that in the long run we should get what we expect. Even this weak defense is no longer used and researchers generally treat realized returns as expected returns in their tests without any qualifications. The purpose of this article is to convince the reader there is a distinction and that it is worth our collective efforts to think about alternative ways to estimate expected returns. I have tried to examine some alternatives. I do not pretend that I have the final answers here. However, I do hope that I have convinced you that it is a worthwhile search.

REFERENCES

- Backus, David, Silverio Foresi, Abon Mozumdar, and Liuren Wu, 1997, Predictable changes in yields and forward rates, Working paper, New York University.
- Balduzzi, Pierluigi, Edwin J. Elton, and T. Clifton Green, 1997, Economic news and the yield curve: Evidence from the U.S. Treasury market, Working paper, New York University.
- Balduzzi, Pierluigi, and Hedi Kallal, 1997, Risk premium and variance bounds, *Journal of Finance* 52, 1913–1950.
- Bernard, Victor L., and Jacob K. Thomas, 1990, Evidence that stock prices do not fully reflect the implications of current earnings for future earnings, *Journal of Accounting and Economics* 13, 305–340.
- Brennan, Michael J., and Eduardo Schwartz, 1982, An equilibrium model of bond pricing and a test of market efficiency, *Journal of Financial and Quantitative Analysis* 17, 301–329.
- Buser, Stephen A., Andrew G. Karolyi, and Anthony B. Sanders, 1996, Adjusted forward rates as predictors of future spot rates, *Journal of Fixed Income* 6, 29–42.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, N.J.).
- 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.
- Conner, Gregory, and Robert A. Korajczyk, 1988, Risk and return in an equilibrium APT, *Journal of Financial Economics* 21, 255–289.
- Cox, John C., Jonathan E. Ingersoll Jr., and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–403.
- Edison, H. J., 1996, The reaction of exchange rates and interest rates to news releases, International Finance discussion paper 570, Board of Governors of the Federal Reserve System.
- Elton, Edwin J., and T. Clifton Green, 1998, Tax and liquidity effects in pricing government bonds, *Journal of Finance* 53, 1533–1562.
- Elton, Edwin J., Martin J. Gruber, Deepak Agrawal, and Christopher Mann, 1998, Is there a risk premium in corporate bonds?, Working paper, New York University.
- Elton, Edwin J., Martin J. Gruber, and Christopher Blake, 1998, Common factors in mutual fund returns, *European Finance Review*, forthcoming.
- Elton, Edwin J., Martin J. Gruber, and Jianping Mei, 1996, Return generating process and the determinants of risk premiums, *Journal of Banking and Finance* 20, 1251–1269.
- Elton, Edwin J., Martin J. Gruber, and Roni Michaely, 1990, The structure of spot rates and immunization, *Journal of Finance* 45, 629–642.
- Elton, Edwin J., Martin J. Gruber, and Manfred Padberg, 1976, Simple criteria for optimal portfolio selection, *Journal of Finance* 31, 1341–1357.
- Elton, Edwin J., Martin J. Gruber, and Joel Rentzler, 1983, The arbitrage pricing model and returns on assets under uncertain inflation, *Journal of Finance* 38, 525–537.
- Fama, Eugene F., and Michael R. Gibbons, 1984, A comparison of inflation forecasts, *Journal of Monetary Economics* 13, 327–348.
- Fleming, Michael J., and Eli M. Remolona, 1996, Price formation and liquidity in the U.S. Treasuries market: Evidence from intra-day patterns around announcements, *Journal of Finance*, forthcoming.
- Froot, Kenneth, 1989, New hope for expectations hypothesis of the term structure of interest rates, *Journal of Finance* 44, 283–306.
- Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121–1152.
- Hakkio, Craig S., and Douglas K. Pearce, 1985, The reaction of exchange rates to economic news, *Economic Inquiry* 23, 621–636.
- Hardouvelis, Gikas A., 1988, Economic news, exchange rates and interest rates, *Journal of International Money and Finance* 7, 23–36.
- Ibbotson Associates, 1995, *Stocks, Bonds, Bills and Inflation, Yearbook* (Ibbotson Associates, Chicago, Ill.).
- Ito, Takatoshi, and V. Vance Roley, 1991, Intraday yen/dollar exchange rate movements: News or noise?, *Journal of International Financial Management, Institutions and Markets* 1, 1–31.

- Lindley, D., 1957, A statistical paradox, *Biometrika* 44, 187–192.
- Litterman, Robert, and Jose Scheinkman, 1991, Common factors affecting bond returns, *Journal of Fixed Income* 1, 54–61.
- Longstaff, Francis A., and Edward S. Schwartz, 1992, Interest rate volatility and the term structure: A two-factor general equilibrium model, *Journal of Finance* 47, 1259–1282.
- McQueen, Grant, and C. Vance Roley, 1993, Stock prices, news, and business conditions, *Review of Financial Studies* 6, 683–707.
- Meiselman, David, 1962, *The Term Structure of Interest Rates* (Prentice Hall, Englewood Cliffs, N.J.).
- Merton, Robert C., 1969, Lifetime portfolio selection under uncertainty: The continuous time case, *Review of Economics and Statistics* 51, 247–257.
- Pastor, Lubos, and Robert Stambaugh, 1999, Costs of equity capital and model mispricing, *Journal of Finance* 54, 67–123.
- Urich, Thomas, and Paul Wachtel, 1984, The effects of inflation and money supply announcements on interest rates, *Journal of Finance* 39, 1177–1188.