

New facts in finance

John H. Cochrane

Introduction and summary

The last 15 years have seen a revolution in the way financial economists understand the investment world. We once thought that stock and bond returns were essentially unpredictable. Now we recognize that stock and bond returns have a substantial predictable component at long horizons. We once thought that the capital asset pricing model (CAPM) provided a good description of why average returns on some stocks, portfolios, funds, or strategies were higher than others. Now we recognize that the average returns of many investment opportunities cannot be explained by the CAPM, and “multifactor models” are used in its place. We once thought that long-term interest rates reflected expectations of future short-term rates and that interest rate differentials across countries reflected expectations of exchange rate depreciation. Now, we see time-varying risk premiums in bond and foreign exchange markets as well as in stock markets. We once thought that mutual fund average returns were well explained by the CAPM. Now, we see that funds can earn average returns not explained by the CAPM, that is, unrelated to market risks, by following a variety of investment “styles.”

In this article, I survey these new facts, and I show how they are variations on a common theme. Each case uses price variables to infer market expectations of future returns; each case notices that an offsetting adjustment (to dividends, interest rates, or exchange rates) seems to be absent or sluggish. Each case suggests that financial markets offer rewards in the form of average returns for holding risks related to recessions and financial distress, in addition to the risks represented by overall market movements. In a companion article in this issue, “Portfolio advice for a multifactor world,” I survey and interpret recent advances in portfolio theory that address the question, What should an investor do about all these new facts?

First, a slightly more detailed overview of the facts then and now. Until the mid-1980s, financial

economists’ view of the investment world was based on three bedrocks:

1. The CAPM is a good measure of risk and thus a good explanation of the fact that some assets (stocks, portfolios, strategies, or mutual funds) earn higher average returns than others. The CAPM states that assets can only earn a high average return if they have a high “beta,” which measures the tendency of the individual asset to move up or down with the market as a whole. Beta drives average returns because beta measures how much adding a *bit* of the asset to a diversified portfolio increases the volatility of the *portfolio*. Investors care about portfolio returns, not about the behavior of specific assets.

2. Returns are unpredictable, like a coin flip. This is the *random walk* theory of stock prices. Though there are bull and bear markets; long sequences of good and bad *past* returns; the expected *future* return is always about the same. *Technical analysis* that tries to divine future returns from patterns of past returns and prices is nearly useless. Any apparent predictability is either a statistical artifact which will quickly vanish out of sample or cannot be exploited after transaction costs.

Bond returns are not predictable. This is the *expectations model* of the term structure. If long-term bond yields are higher than short-term yields—if the yield curve is upward sloping—this does not mean that you expect a higher return by holding long-term bonds rather than short-term bonds. Rather, it means

John H. Cochrane is the Sigmund E. Edelstone Professor of Finance in the Graduate School of Business at the University of Chicago, a consultant to the Federal Reserve Bank of Chicago, and a research associate at the National Bureau of Economic Research (NBER). The author thanks Andrea Eisfeldt for research assistance and David Marshall, John Campbell, and Robert Shiller for comments. The author’s research is supported by the Graduate School of Business and by a grant from the National Science Foundation, administered by the NBER.

that short-term interest rates are expected to rise in the future. Over one year, the rise in interest rates will limit the capital gain on long-term bonds, so they earn the same as the short-term bonds over the year. Over many years, the rise in short rates improves the rate of return from rolling over short-term bonds to equal that of holding the long-term bond. Thus, you expect to earn about the same amount on short-term or long-term bonds at any horizon.

Foreign exchange bets are not predictable. If a country has higher interest rates than are available in the U.S. for bonds of a similar risk class, its exchange rate is expected to depreciate. Then, after you convert your investment back to dollars, you expect to make the same amount of money holding foreign or domestic bonds.

In addition, stock market volatility does not change much through time. Not only are returns close to unpredictable, they are nearly identically distributed as well. Each day, the stock market return is like the result of flipping the same coin, over and over again.

3. Professional managers do not reliably outperform simple indexes and passive portfolios once one corrects for risk (beta). While some do better than the market in any given year, some do worse, and the outcomes look very much like luck. Funds that do well in one year are not more likely to do better than average the next year. The average actively managed fund performs about 1 percent *worse* than the market index. The more actively a fund trades, the lower the returns to investors.

Together, these views reflect a guiding principle that asset markets are, to a good approximation, *informationally efficient* (Fama, 1970, 1991). Market prices already contain most information about fundamental value and, because the business of discovering information about the value of traded assets is extremely competitive, there are no easy quick profits to be made, just as there are not in any other well-established and competitive industry. The only way to earn large returns is by taking on additional risk.

These views are not ideological or doctrinaire beliefs. Rather, they summarize the findings of a quarter century of careful empirical work. However, every one of them has now been extensively revised by a new generation of empirical research. The new findings need not overturn the cherished view that markets are reasonably competitive and, therefore, reasonably efficient. However, they do substantially enlarge our view of what activities provide rewards for holding risks, and they challenge our understanding of those risk premiums.

Now, we know that:

1. There are assets whose average returns can not be explained by their beta. Multifactor extensions of the CAPM dominate the description, performance attribution, and explanation of average returns. Multifactor models associate high average returns with a tendency to move with other risk factors in addition to movements in the market as a whole. (See box 1.)

2. Returns are predictable. In particular: Variables including the dividend/price (d/p) ratio and term premium can predict substantial amounts of stock return variation. This phenomenon occurs over business cycle and longer horizons. Daily, weekly, and monthly stock returns are still close to unpredictable, and technical systems for predicting such movements are still close to useless.

Bond returns are predictable. Though the expectations model works well in the long run, a steeply upward sloping yield curve means that expected returns on long-term bonds are higher than on short-term bonds for the next year. These predictions are not guarantees—there is still substantial risk—but the tendency is discernible.

Foreign exchange returns are predictable. If you put your money in a country whose interest rates are higher than usual relative to the U.S., you expect to earn more money even after converting back to dollars. Again, this prediction is not a guarantee—exchange rates do vary, and a lot, so the strategy is risky.

Volatility does change through time. Times of past volatility indicate future volatility. Volatility also is higher after large price drops. Bond market volatility is higher when interest rates are higher, and possibly when interest rate spreads are higher as well.

3. Some mutual funds seem to outperform simple indexes, even after controlling for risk through market betas. Fund returns are also slightly predictable: Past winning funds seem to do better than average in the future, and past losing funds seem to do worse than average in the future. For a while, this seemed to indicate that there is some persistent skill in active management. However, multifactor models explain most fund persistence: Funds earn persistent returns by following fairly mechanical *styles*, not by persistent skill at stock selection.

Again, these statements are not dogma, but a cautious summary of a large body of careful empirical work. The strength and usefulness of many results are hotly debated, as are the underlying reasons for many of these new facts. But the old world is gone.

BOX 1

The CAPM and multifactor models

The CAPM uses a *time-series* regression to measure beta, β , which quantifies an asset's or portfolio's tendency to move with the market as a whole,

$$R_t^i - R_t^f = a_i + \beta_{im}(R_t^m - R_t^f) + \varepsilon_t^i; \\ t = 1, 2 \dots T \text{ for each asset } i.$$

Then, the CAPM predicts that the expected excess return should be proportional to beta,

$$E(R_t^i - R_t^f) = \beta_{im}\lambda_m \text{ for each } i.$$

λ_m gives the “price of beta risk” or “market risk premium”—the amount by which expected returns must rise to compensate investors for higher beta. Since the model applies to the market return as well, we can measure λ_m via

$$\lambda_m = E(R_t^m - R_t^f).$$

Multifactor models extend this theory in a straightforward way. They use a time-series *multiple* regression to quantify an asset's tendency to move with multiple risk factors F^A , F^B , etc.

$$3) \quad R_t^i - R_t^f = a_i + \beta_{im}(R_t^m - R_t^f) + \beta_{iA}F_t^A + \beta_{iB}F_t^B \\ + \dots + \varepsilon_t^i; \quad t = 1, 2 \dots T \text{ for each asset } i.$$

Then, the multifactor model predicts that the expected excess return is proportional to the betas

$$4) \quad E(R_t^i - R_t^f) = \beta_{im}\lambda_m + \beta_{iA}\lambda_A + \beta_{iB}\lambda_B + \dots \\ \text{for each } i.$$

The residual or unexplained average return in either case is called an alpha,

$$\alpha_i \equiv E(R_t^i - R_t^f) - (\beta_{im}\lambda_m + \beta_{iA}\lambda_A + \beta_{iB}\lambda_B + \dots).$$

The CAPM and multifactor models

The CAPM

The CAPM proved stunningly successful in a quarter century of empirical work. Every strategy that seemed to give high average returns turned out to have a high beta, or a large tendency to move with the market. Strategies that one might have thought gave high average returns (such as holding very volatile stocks) turned out not to have high average returns when they did not have high betas.


Figure 1 presents a typical evaluation of the CAPM. I examine 10 portfolios of NYSE stocks sorted by size (total market capitalization), along with a portfolio of corporate bonds and long-term government bonds. As the vertical axis shows, there is a sizable spread in average returns between large stocks (lower average return) and small stocks (higher average return) and a large spread between stocks and bonds. The figure plots these average returns against market betas. You can see how the CAPM prediction fits: Portfolios with higher average returns have higher betas.

In fact, figure 1 captures one of the first significant *failures* of the CAPM. The smallest firms (the far right portfolio) seem to earn an average return a few percent too high given their betas. This is the celebrated “small-firm effect,” (Banz, 1981) and this deviation is statistically significant. Would that all failed economic theories worked so well! However, the plot shows that this effect is within the range that statisticians can argue about. Estimating the slope of the


line by fitting a cross-sectional regression (average return against beta), shown in the colored line, rather than forcing the line to go through the market and Treasury bill return, shown in the black line, halves

FIGURE 1

CAPM—Mean excess returns vs. beta, version 1



Notes: Average returns versus betas on the NYSE value-weighted portfolio for ten size-sorted stock portfolios, government bonds, and corporate bonds. Sample period 1947–96. The black line draws the CAPM prediction by fitting the market proxy and Treasury bill rates exactly (a time-series test) and the colored line draws the CAPM prediction by fitting an OLS cross-sectional regression to the displayed data points (a second-pass or cross-sectional test). The small-firm portfolios are at the top right. Moving down and to the left, one sees increasingly large-firm portfolios and the market index. The points far down and to the left are the government bond and Treasury bill returns.



the small-firm effect. Figure 2 uses the equally weighted portfolio as market proxy, and this change in specification eliminates the small-firm effect, making the line of average returns versus betas if anything too shallow rather than too steep.

Why we expect multiple factors

In retrospect, it is surprising that the CAPM worked so well for so long. The assumptions on which it is built are very stylized and simplified. Asset pricing theory recognized at least since Merton (1973, 1971) the theoretical possibility, indeed probability, that we should need *factors*, *state variables* or *sources of priced risk*, beyond movements in the market portfolio to explain why some average returns are higher than others. (See box 1 for details of the CAPM and multifactor models.)

Most importantly, *the average investor has a job*. The CAPM (together with the use of the NYSE portfolio as the market proxy) simplifies matters by assuming that the average investor only cares about the performance of his investment portfolio. While there are investors like that, for most of us eventual wealth comes both from investment and from earning a living. Importantly, events like recessions hurt the majority of investors. Those who don’t actually lose jobs get lower salaries or bonuses. A very limited number of people actually do better in a recession.

With this fact in mind, compare two stocks. They both have the same sensitivity to market movements. However, one of them does well in recessions, while the other does poorly. Clearly, most investors prefer the stock that does well in recessions, since its performance will cushion the blows to their other income.

If lots of people feel that way, they bid up the price of that stock, or, equivalently, they are willing to hold it at a lower average return. Conversely, the procyclical stock’s price will fall or it must offer a higher average return in order to get investors to hold it.

In sum, we should expect that procyclical stocks that do well in booms and worse in recessions will have to offer higher average returns than countercyclical stocks that do well in recessions, even if the stocks have the same market beta. We expect that *another dimension of risk*—covariation with recessions—will matter in determining average returns.¹

What kinds of additional factors should we look for? Generally, asset pricing theory specifies that assets will have to pay high average returns if they do poorly in “bad times”—times in which investors would particularly like their investments not to perform badly and are willing to sacrifice some expected return in order to ensure that this is so. Consumption (or, more generally, marginal utility) should provide the purest measure of bad times. Investors consume less when their income prospects are low or if they think future returns will be bad. Low consumption thus *reveals* that this is indeed a time at which investors would especially like portfolios not to do badly, and would be willing to pay to ensure that wish. Alas, efforts to relate asset returns to consumption data are not (yet) a great success. Therefore, empirically useful asset pricing models examine more direct measures of good times or bad times. Broad categories of such indicators are

1. The market return. The CAPM is usually included and extended. People are unhappy if the market crashes.
2. Events, such as recessions, that drive investors’ noninvestment sources of income.
3. Variables, such as the p/d ratio or slope of the yield curve, that forecast stock or bond returns (called “state variables for changing investment opportunity sets”).

4. Returns on other well-diversified portfolios.

One formally justifies the first three factors by stating assumptions under which each variable is related to average consumption. For example, 1) if the market as a whole declines, consumers lose wealth and will cut back on consumption; 2) if a recession leads people to lose their jobs, then they will cut back on consumption; and, 3) if you are saving for retirement, then news that interest rates and average stock returns have declined is bad news, which will cause you to lower consumption. This last point establishes a connection between predictability of returns and the presence of additional risk factors for understanding

the cross-section of average returns. As pointed out by Merton (1971), one would give up some average return to have a portfolio that did well when there was bad news about future market returns.

The fourth kind of factor—additional portfolio returns—is most easily defended as a proxy for any of the other three. The fitted value of a regression of any pricing factor on the set of all asset returns is a portfolio that carries exactly the same pricing information as the original factor—a *factor-mimicking* portfolio.

It is vital that the extra risk factors affect the *average* investor. If an event makes investor A worse off and investor B better off, then investor A buys assets that do well when the event happens and investor B sells them. They transfer the risk of the event, but the price or expected return of the asset is unaffected. For a factor to affect prices or expected returns, it must affect the average investor, so investors collectively bid up or down the price and expected return of assets that covary with the event rather than just transferring the risk without affecting equilibrium prices.


Inspired by this broad direction, empirical researchers have found quite a number of specific factors that seem to explain the variation in average returns across assets. In general, empirical success varies inversely with theoretical purity.

Small and value/growth stocks

The size and book to market factors advocated by Fama and French (1996) are one of the most popular additional risk factors.

Small-cap stocks have small market values (price times shares outstanding). Value (or high book/market) stocks have market values that are small relative to the value of assets on the company's books. Both categories of stocks have quite high average returns. Large and growth stocks are the opposite of small and value and seem to have unusually low average returns. (See Fama and French, 1993, for a review.) The idea that low prices lead to high average returns is natural.

High average returns are consistent with the CAPM, if these categories of stocks have high sensitivity to the market, high betas. However, small and especially value stocks seem to have abnormally high returns even after accounting for market beta. Conversely, growth stocks seem to do systematically worse than their CAPM betas suggest. Figure 3 shows this value-size puzzle. It is just like figure 1, except that the stocks are sorted into portfolios based on size and book/market ratio² rather than size alone. The highest portfolios have *three* times the average excess return of the lowest portfolios, and this variation has nothing at all to do with market betas.



In figure 4, I connect portfolios of different sizes within the same book/market category (panel A). Variation in *size* produces a variation in average returns that is positively related to variation in market betas, as shown in figure 1. In panel B, I connect portfolios that have different book/market ratios within size categories. Variation in book/market ratio produces a variation in average return that is *negatively* related to market beta. Because of this value effect, the CAPM is a disaster when confronted with these portfolios.

To explain these facts, Fama and French (1993, 1996) advocate a multifactor model with the market return, the return of small less big stocks (SMB), and the return of high book/market less low book/market stocks (HML) as three factors. They show that variation in average returns of the 25 size and book/market portfolios can be explained by varying loadings (betas) on the latter two factors.

Figure 5 illustrates Fama and French's results. As in figure 4, the vertical axis is the average returns of the 25 size and book/market portfolios. Now, the horizontal axis is the predicted values from the Fama–French three-factor model. The points should all lie on a 45 degree line if the model is correct. The points lie much closer to this prediction in figure 5 than in figures 3 and 4. The worst fit is for the growth stocks (lowest line, panel A), for which there is little variation in average return despite large variation in size beta as one moves from small to large firms.


What are the size and value factors?

One would like to understand the real, macroeconomic, aggregate, nondiversifiable risk that is proxied by the returns of the HML and SMB portfolios. Why


FIGURE 4

Mean excess returns vs. market beta, varying size and book/market ratio

A. Changing size within book/market category



B. Changing book/market within size category



Notes: Average returns versus market beta for 25 stock portfolios sorted on the basis of size and book/market ratio. The points are the same as figure 3. In panel A, lines connect portfolios as size varies within book/market categories; in panel B, lines connect portfolios as book/market ratio varies within size categories.

are investors so concerned about holding stocks that do badly at the times that the HML (value less growth) and SMB (small-cap less large-cap) portfolios do badly, even though the market does not fall? The answer to this question is not yet totally clear.


Fama and French (1995) note that the typical value stock has a price that has been driven down due to financial distress. The stocks of firms on the verge of bankruptcy have recovered more often than not, which generates the high average returns of this

strategy.³ This observation suggests a natural interpretation of the value premium: In the event of a credit crunch, liquidity crunch, or flight to quality, stocks in financial distress will do very badly, and this is precisely when investors least want to hear that their portfolio is losing money. (One cannot count the “distress” of the individual firm as a risk factor. Such distress is idiosyncratic and can be diversified away. Only aggregate events that average investors care about can result in a risk premium.)


FIGURE 5

Mean excess return vs. three-factor model predictions

A. Changing size within book/market category



B. Changing book/market within size category



Notes: Average returns versus market beta for 25 stock portfolios sorted on the basis of size and book/market ratio versus predictions of Fama–French three-factor model. The predictions are derived by regressing each of the 25 portfolio returns, R^i_t , on the market portfolio, R^m_t , and the two Fama–French factor portfolios, SMB_t (small minus big) and HML_t (high minus low book/market). (See equation 4 in box 1.)

Heaton and Lucas's (1997) results add to this story for the value effect. They note that the typical stockholder is the proprietor of a small, privately held business. Such an investor's income is, of course, particularly sensitive to the kinds of financial events that cause distress among small firms and distressed value firms. Therefore, this investor would demand a substantial premium to hold value stocks and would hold growth stocks despite a low premium.

Liew and Vassalou (1999), among others, link value and small-firm returns to macroeconomic events. They find that in many countries, counterparts to HML and SMB contain supplementary information to that contained in the market return for forecasting gross domestic product (GDP) growth. For example, they report a regression

$$GDP_{t \rightarrow t+4} = a + 0.065 MKT_{t-4 \rightarrow t} + 0.058 HML_{t-4 \rightarrow t} + \epsilon_{t+4}$$


where $GDP_{t \rightarrow t+4}$ denotes the following year's GDP growth and $MKT_{t-4 \rightarrow t}$ and $HML_{t-4 \rightarrow t}$ denote the previous year's return on the market index and HML portfolio. Thus, a 10 percent HML return raises the GDP forecast by 0.5 percentage points. (Both coefficients are significant with t-statistics of 3.09 and 2.83, respectively.)

The effects are still under investigation. Figure 6 plots the cumulative return on the HML and SMB portfolios; a link between these returns and obvious macroeconomic events does not jump out. Both portfolios have essentially no correlation with the market return, though HML does seem to move inversely with large market declines. HML goes down more than the market in some business cycles, but less in others.

On the other hand, one can ignore Fama and French's motivation and regard the model as an *arbitrage pricing* theory (APT) following Ross (1976). If the returns of the 25 size and book/market portfolios could be *perfectly* replicated by the returns of the three-factor portfolios—if the R^2 values in the time-series regressions of the 25 portfolio on the three factors were 100 percent—then the multifactor model would have to hold exactly, in order to preclude arbitrage opportunities. To see this, suppose that one of the 25 portfolios—call it portfolio A—gives an average return 5 percent above the average return predicted by the Fama–French model, and its R^2 is 100 percent. Then, one could short a combination of the three-factor portfolios, buy portfolio A, and earn a completely riskless profit. This logic is often used to argue that a *high* R^2 should imply an *approximate* multifactor model. If the R^2 were only 95 percent, then an average return 5 percent above the factor model prediction

FIGURE 6

Cumulative returns on market portfolios



Notes: Cumulative returns on the market RMRF, SMB, and HML portfolios. The SMB return is formed by $R_{SMB}^B + aSMB$; $a = \sigma(RMRF)/\sigma(SMB)$. In this way it is a return that can be cumulated rather than a zero-cost portfolio, and its standard deviation is equal to that of the market return. HML is adjusted similarly. The vertical axis is the log base 2 of the cumulative return or value of \$1 invested at the beginning of the sample period. Thus, each time a line increases by 1 unit, the value doubles.

would imply that the strategy long portfolio A and short a combination of the three-factor portfolios would earn a very high average return with very little, though not zero, risk—a very high Sharpe ratio.

In fact, the R^2 values of Fama and French's (1993) time-series regressions are all in the 90 percent to 95 percent range, so extremely high risk prices for the residuals would have to be invoked for the model *not* to fit well. Conversely, given the average returns from HML and SMB and the failure of the CAPM to explain those returns, there would be near-arbitrage opportunities if value and small stocks did not move together in the way described by the Fama–French model.

One way to assess whether the three factors proxy for real macroeconomic risks is by checking whether the multifactor model prices additional portfolios, especially portfolios whose ex-post returns are not well explained by the factors (portfolios that do not have high R^2 values in time-series regressions). Fama and French (1996) find that the SMB and HML portfolios comfortably explain strategies based on alternative price multiples (price/earnings, book/market), five-year sales growth (this is the only strategy that does not form portfolios based on price variables), and the tendency of five-year returns to reverse. All of these strategies are not explained by CAPM betas. However, they all also produce portfolios with high R^2 values in a time-series regression on the HML and SMB portfolios. This is good and bad news. It might

mean that the model is a good APT, and that the size and book/market characteristics describe the major sources of priced variation in all stocks. On the other hand, it might mean that these extra ways of constructing portfolios just haven't identified other sources of priced variation in stock returns. (Fama and French, 1996, also find that HML and SMB do not explain *momentum*, despite high R^2 values. I discuss this anomaly below.) The portfolios of stocks sorted by industry in Fama and French (1997) have lower R^2 values, and the model works less well.

A final concern is that the size and book/market premiums seem to have diminished substantially in recent years. The sharp decline in the SMB portfolio return around 1980 when the small-firm effect was first popularized is obvious in figure 6. In Fama and French's (1993) initial samples, 1960–90, the HML cumulative return starts about one-half (0.62) below the market and ends up about one-half (0.77) above the market. On the log scale of the figure, this corresponds to Fama and French's report that the HML average return is about double (precisely, $2^{0.62+0.77} = 2.6$ times) that of the market. However, over the entire sample of the plot, the HML portfolio starts and ends at the same place and so earns almost exactly the same as the market. From 1990 to now, the HML portfolio loses about one-half relative to the market, meaning an investor in the market has increased his money one and a half times as much as an HML investor. (The actual number is 0.77 so the market return is $2^{0.77} = 1.71$ times better than the HML return.)

Among other worries, if the average returns decline right after publication it suggests that the anomalies may simply have been overlooked by a large fraction of investors. As they move in, prices go up further, helping the apparent anomaly for a while. But once a large number of investors have moved in to include small and value stocks in their portfolios, the anomalous high average returns disappear.

However, average returns are hard to measure. There have been previous ten- to 20-year periods in which small stocks did very badly, for example the 1950s, and similar decade-long variations in the HML premium. Also, since SMB and HML have a beta of essentially zero on the market, *any* upward trend is a violation of the CAPM and says that investors can improve their overall mean–variance tradeoff by taking on some of the HML or SMB portfolio.

Macroeconomic factors

I focus on the size and value factors because they provide the most empirically successful multifactor model and have attracted much industry as well as

academic attention. Several authors have used macroeconomic variables as factors. This procedure examines directly whether stock performance during bad macroeconomic times determines average returns. Jagannathan and Wang (1996) and Reyfman (1997) use labor income; Chen, Roll, and Ross (1986) look at industrial production and inflation among other variables; and Cochrane (1996) looks at investment growth. All these authors find that average returns line up with betas calculated using the macroeconomic indicators. The factors are theoretically easier to motivate, but none explains the value and size portfolios as well as the (theoretically less solid, so far) size and value factors.

Merton's (1973, 1971) theory says that variables which predict market returns should show up as factors that explain cross-sectional variation in average returns. Campbell (1996) is the lone test I know of to directly address this question. Cochrane (1996) and Jagannathan and Wang (1996) perform related tests in that they include “scaled return” factors, for example, market return at t multiplied by d/p ratio at $t - 1$; they find that these factors are also important in understanding cross-sectional variation in average returns.

The next step is to link these more fundamentally determined factors with the empirically more successful value and small-firm factor portfolios. Because of measurement difficulties and selection biases, fundamentally determined macroeconomic factors will never approach the empirical performance of portfolio-based factors. However, they may help to explain which portfolio-based factors really work and why.

Predictable returns

The view that risky asset returns are largely unpredictable, or that prices follow “random walks,” remains immensely successful (Malkiel, 1990, is a classic and readable introduction). It is also widely ignored.

Unpredictable returns mean that if stocks went up yesterday, there is no exploitable tendency for them to decline today because of “profit taking” or to continue to rise today because of “momentum.” “Technical” signals, including analysis of past price movements trading volume, open interest, and so on are close to useless for forecasting short-term gains and losses. As I write, value funds are reportedly suffering large outflows because their stocks have done poorly in the last few months, leading fund investors to move money into blue-chip funds that have performed better (New York Times Company, 1999). Unpredictable returns mean that this strategy will not do anything for investors' portfolios over the long run except rack up trading costs. If funds are selling stocks, then

contrarian investors must be buying them, but unpredictable returns mean that this strategy can not improve performance either. If one can not systematically make money, one can not systematically lose money either.

As discussed in the introduction, researchers once believed that stock returns (more precisely, the excess returns on stocks over short-term interest rates) were completely unpredictable. It now turns out that average returns on the market and individual securities *do* vary over time and that stock returns *are* predictable. Alas for would-be technical traders, much of that predictability comes at long horizons and seems to be associated with business cycles and financial distress.

Market returns

Table 1 presents a regression that forecasts returns. Low prices—relative to dividends, book value, earnings, sales, or other divisors—predict higher subsequent returns. As the R^2 values in table 1 show, these are long-horizon effects: Annual returns are only slightly predictable and month-to-month returns are still strikingly unpredictable, but returns at five-year horizons seem very predictable. (Fama and French, 1989, is an excellent reference for this kind of regression).

The results at different horizons are reflections of a single underlying phenomenon. If daily returns are very slightly predictable by a slow-moving variable, that predictability adds up over long horizons. For example, you can predict that the temperature in Chicago will rise about one-third of a degree per day in spring. This forecast explains very little of the day to day variation in temperature, but tracks almost all of the rise in temperature from January to July. Thus, the R^2 rises with horizon.

Precisely, suppose that we forecast returns with a forecasting variable x , according to

$$1) \quad R_{t+1} - R_{t+1}^{TB} = a + bx_t + \varepsilon_{t+1}$$

$$2) \quad x_{t+1} = c + px_t + \delta_{t+1}.$$

Small values of b and R^2 in equation 1 and a large coefficient p in equation 2 imply mathematically that the long-horizon regression as in table 1 has a large regression coefficient b and large R^2 .

This regression has a powerful implication: Stocks are in many ways like bonds. Any bond investor understands that a string of good past returns that pushes the price up is bad news for subsequent returns. Many stock investors see a string of good past returns and become elated that we seem to be in a “bull

market,” concluding future stock returns will be good as well. The regression reveals the opposite: A string of good past returns which drives up stock prices is bad news for subsequent stock returns, as it is for bonds.

Long-horizon return predictability was first documented in the volatility tests of Shiller (1981) and LeRoy and Porter (1981). They found that stock prices vary far too much to be accounted for by changing expectations of subsequent cash flows; thus changing discount rates or expected returns must account for variation in stock prices. These volatility tests turn out to be almost identical to regressions such as those in table 1 (Cochrane, 1991).

Momentum and reversal

Since a string of good returns gives a high price, it is not surprising that individual stocks that do well for a long time (and reach a high price) subsequently do poorly, and stocks that do poorly for a long time (and reach a low price, market value, or market to book ratio) subsequently do well. Table 2, taken from Fama and French (1996) confirms this hunch. (Also, see DeBont and Thaler, 1985, and Jegadeesh and Titman, 1993.)

The first row in table 2 tracks the average monthly return from the *reversal* strategy. Each month, allocate all stocks to ten portfolios based on performance from year -5 to year -1 . Then, buy the best-performing portfolio and short the worst-performing portfolio. This strategy earns a hefty -0.74 percent monthly return.⁴ Past long-term losers come back and past winners do badly. Fama and French (1996) verify that these portfolio returns are explained by their three-factor model. Past winners move with value stocks,

TABLE 1

OLS regression of excess returns on price/dividend ratio

Horizon k	b	Standard error	R^2
1 year	-1.04	0.33	0.17
2 years	-2.04	0.66	0.26
3 years	-2.84	0.88	0.38
5 years	-6.22	1.24	0.59

Notes: OLS regressions of excess returns (value-weighted NYSE–Treasury bill rate) on value-weighted price/dividend ratio.

$$R_{t \rightarrow t+k}^{W} - R_{t \rightarrow t+k}^{TB} = a + b(P_t / D_t) + \varepsilon_{t+k}.$$

$R_{t \rightarrow t+k}$ indicates the k year return. Standard errors use GMM to correct for heteroskedasticity and serial correlation.

TABLE 2			
Average monthly returns, reversal and momentum strategies			
Strategy	Period	Portfolio formation (months)	Average return, 10–1 (monthly%)
Reversal	July 1963–Dec. 1993	60–13	–0.74
Momentum	July 1963–Dec. 1993	12–2	+1.31
Reversal	Jan. 1931–Feb. 1963	60–13	–1.61
Momentum	Jan. 1931–Feb. 1963	12–2	+0.38

Notes: Each month, allocate all NYSE firms to 10 portfolios based on their performance during the “portfolio formation months” interval. For example, 60–13 forms portfolios based on returns from 5 years ago to 1 year, 1 month ago. Then buy the best-performing decile portfolio and short the worst-performing decile portfolio.

Source: Fama and French (1996, table 6).

and so inherit the value stock premium. (To compare the strategies, the table always buys the winners and shorts the losers. In practice, of course, you buy the losers and short the winners to earn +0.71 percent monthly average return.)

The second row of table 2 tracks the average monthly return from a *momentum* strategy. Each month, allocate all stocks to ten portfolios based on performance in the last *year*. Now, the winners continue to win and the losers continue to lose, so that buying the winners and shorting the losers generates a positive 1.31 percent monthly return.

Momentum is not explained by the Fama–French (1996) three-factor model. The past losers have low prices and tend to move with value stocks. Hence, the model predicts that they should have high average returns, not low average returns.

Momentum stocks move together, as do value and small stocks, so a “momentum factor” works to “explain” momentum portfolio returns (Carhart, 1997). This step is so obviously ad hoc (that is, an APT factor that will only explain returns of portfolios organized on the same characteristic as the factor rather than a proxy for macroeconomic risk) that most people are uncomfortable adding it. We obviously do not want to add a new factor for every anomaly.

Is momentum really there, and if so, is it exploitable after transaction costs? One warning is that it does not seem stable over subsamples. The third and fourth lines in table 2 show that the momentum effect essentially disappears in the earlier data sample, while reversal is even stronger in that sample.

Momentum is really just a new way of looking at an old phenomenon, the small apparent predictability of monthly individual stock returns. A tiny regression R^2 for forecasting monthly returns of 0.0025 (0.25 percent) is more than adequate to generate the momentum

results of table 2. The key is the large standard deviation of individual stock returns, typically 40 percent or more on an annual basis. The average return of the best performing decile of a normal distribution is 1.76 standard deviations above the mean,⁵ so the winning momentum portfolio went up about 80 percent in the previous year and the typical losing portfolio went down about 60 percent. Only a small amount of continuation will give a 1 percent monthly return when multiplied by such large past returns. To be precise, the monthly individual stock standard deviation is about $40\% / \sqrt{12} \approx 12\%$. If the R^2 is 0.0025,

the standard deviation of the predictable part of returns is $\sqrt{0.0025} \times 12\% \approx 0.6\%$. Hence, the decile predicted to perform best will earn $1.76 \times 0.6\% \approx 1\%$ above the mean. Since the strategy buys the winners and shorts the losers, an R^2 of 0.0025 implies that one should earn a 2 percent monthly return by the momentum strategy.

We have known at least since Fama (1965) that monthly and higher frequency stock returns have slight, statistically significant predictability with R^2 about 0.01. Campbell, Lo, and MacKinlay (1997, table 2.4) provide an updated summary of index autocorrelations (the R^2 is the squared autocorrelation), part of which I show in table 3. Note the correlation of the equally weighted portfolio, which emphasizes small stocks.⁶

However, such small, though statistically significant, high-frequency predictability has thus far failed to yield exploitable profits after one takes into account transaction costs, thin trading of small stocks, and high short-sale costs. The momentum strategy for exploiting this correlation may not work in practice for the same reasons. Momentum does require frequent trading. The portfolios in table 2 are re-formed every

TABLE 3		
First-order autocorrelation, CRSP value- and equally weighted index returns		
Frequency	Portfolio	Correlation ρ_1
Daily	Value-weighted	0.18
	Equally weighted	0.35
Monthly	Value-weighted	0.043
	Equally weighted	0.17

Note: Sample 1962–94.
Source: Campbell, Lo, and MacKinlay (1997).

month. Annual winners and losers will not change that often, but the winning and losing portfolio must be turned over at least once per year. In a quantitative examination of this effect, Carhart (1997) concludes that momentum is not exploitable after transaction costs are taken into account. Moskowitz and Grinblatt (1999) note that most of the apparent gains from the momentum strategy come from short positions in small illiquid stocks. They also find that a large part of momentum profits come from short positions taken in November. Many investors sell losing stocks toward the end of December to establish tax losses. By shorting illiquid losing stocks in November, an investor can profit from the selling pressure in December. This is also an anomaly, but it seems like a glitch rather than a central principle of risk and return in asset markets.

Even if momentum and reversal are real and as strong as indicated by table 2, they do not justify much of the trading based on past results that many investors seem to do. To get the 1 percent per month

momentum return, one buys a portfolio that has typically gone up 80 percent in the last year, and shorts a portfolio that has typically gone down 60 percent. Trading between stocks and fund categories such as value and blue-chip with smaller past returns yields at best proportionally smaller results. Since much of the momentum return seems to come from shorting small illiquid stocks, mild momentum strategies may yield even less. And we have not quantified the substantial risk of momentum strategies.

Bonds

The venerable expectations model of the term structure specifies that long-term bond yields are equal to the average of expected future short-term bond yields (see box 2). For example, if long-term bond yields are higher than short-term bond yields—if the yield curve is upward sloping—this means that short-term rates are expected to rise in the future. The rise in future short-term rates means that investors can expect

BOX 2

Bond definitions and expectations hypothesis

Let $p_t^{(N)}$ denote the log of the N year discount bond price at time t . The N period continuously compounded yield is defined by $y_t^{(N)} = -\frac{1}{N} p_t^{(N)}$. The continuously compounded holding period return is the selling price less the buying price, $hpr_{t+1}^{(N)} = p_{t+1}^{(N-1)} - p_t^{(N)}$. The forward rate is the rate at which an investor can contract today to borrow money $N-1$ years from now, and repay that money N years from now. Since an investor can synthesize a forward contract from discount bonds, the forward rate is determined from discount bond prices by

$$f_t^{(N)} = p_t^{(N-1)} - p_t^{(N)}.$$

The “spot rate” refers, by contrast with a forward rate, to the yield on any bond for which the investor take immediate delivery. Forward rates are typically higher than spot rates when the yield curve rises, since the yield is the average of intervening forward rates,

$$y_t^{(N)} = \frac{1}{N} (f_t^{(1)} + f_t^{(2)} + f_t^{(3)} + \dots + f_t^{(N)}).$$

The expectations hypothesis states that the expected log or continuously compounded return should be the same for any bond strategy. This statement has three mathematically equivalent expressions:

1. The forward rate should equal the expected value of the future spot rate,

$$f_t^{(N)} = E_t(y_{t+N-1}^{(1)}).$$

2. The expected holding period return should be the same on bonds of any maturity

$$E_t(hpr_{t+1}^{(N)}) = E_t(hpr_{t+1}^{(M)}) = y_t^{(1)}.$$

3. The long-term bond yield should equal the average of the expected future short rates,

$$y_t^{(N)} = \frac{1}{N} E_t(y_t^{(1)} + y_{t+1}^{(1)} + \dots + y_{t+N-1}^{(1)}).$$

The expectations hypothesis is often amended to allow a constant risk premium of undetermined sign in these equations. Its violation is then often described as evidence for a “time-varying risk premium.”

The expectations hypothesis is not quite the same thing as risk-neutrality, because the expected log return is not equal to the log expected return. However, the issues here are larger than the difference between the expectations hypothesis and strict risk-neutrality.

TABLE 4			
Zero-coupon bond returns			
Maturity N	Average holding period return	Standard error	Standard deviation
1	5.83	0.42	2.83
2	6.15	0.54	3.65
3	6.40	0.69	4.66
4	6.40	0.85	5.71
5	6.36	0.98	6.58
Note: Continuously compounded one-year holding period returns on zero-coupon bonds of varying maturity. Annual data from CRSP 1953–97.			

the same rate of return whether they hold a long-term bond to maturity or roll over short-term bonds with initially low returns and subsequent higher returns.

As with the CAPM and the view that stock returns are independent over time, a new round of research has significantly modified this traditional view of bond markets.

Table 4 calculates the average return on bonds of different maturities. The expectations hypothesis seems to do pretty well. Average holding period returns do not seem very different across bond maturities, despite the increasing standard deviation of longer-maturity bond returns. The small increase in average returns for long-term bonds, equivalent to a slight average upward slope in the yield curve, is usually excused as a “liquidity premium.” Table 4 is just the tip of an iceberg of successes for the expectations model. Especially in times of significant inflation and exchange rate instability, the expectations hypothesis has done a very good first-order job of explaining the term structure of interest rates.

However, if there are times when long-term bonds are expected to do better and other times when short-term bonds are expected to do better, the unconditional averages in table 4 could still show no pattern. Similarly, one might want to check whether a forward rate that is *unusually high* forecasts an unusual *increase* in spot rates.

Table 5 updates Fama and Bliss’s (1987) classic regression tests of this idea. Panel A presents a regression of the change in yields on the forward-spot spread. (The forward-spot spread measures the slope of the yield curve.) The expectations hypothesis predicts a slope coefficient of 1.0, since the forward rate should equal the expected future spot rate. If, for example, forward rates are lower than expected future spot rates, traders can lock in a borrowing position with a forward contract and then lend at the higher spot rate when the time comes.

Instead, at a one-year horizon we find slope coefficients near zero and a negative adjusted R^2 . Forward rates one year out seem to have no predictive power whatsoever for changes in the spot rate one year from now. On the other hand, by four years out, we see slope coefficients within one standard error of 1.0. Thus, the expectations hypothesis seems to do poorly at short (one-year) horizons, but much better at longer horizons.

If the expectations hypothesis does not work at one-year horizons, then there is money to be made—one must be able to foresee years in which short-term bonds will return more than long-term bonds and vice versa, at least to some extent. To confirm this implication, panel B of table 5 runs regressions of the one-year excess return on long-term bonds on the forward-spot spread. Here, the expectations hypothesis predicts a coefficient of zero: No signal (including the

TABLE 5										
Forecasts based on forward-spotspread										
A. Change in yields						B. Holding period returns				
N	Intercept	Standard error, intercept	Slope	Standard error, slope	Adjusted R ²	Intercept	Standard error, intercept	Slope	Standard error, slope	Adjusted R ²
1	0.10	0.3	−0.10	0.36	−0.020	−0.1	0.3	1.10	0.36	0.16
2	−0.01	0.4	0.37	0.33	0.005	−0.5	0.5	1.46	0.44	0.19
3	−0.04	0.5	0.41	0.33	0.013	−0.4	0.8	1.30	0.54	0.10
4	−0.30	0.5	0.77	0.31	0.110	−0.5	1.0	1.31	0.63	0.07

Notes: OLS regressions, 1953–97 annual data. Panel A estimates the regression $y_{t+N}^{(1)} - y_t^{(1)} = a + b (f_t^{(N+1)} - y_t^{(1)}) + \varepsilon_{t+N}$ and panel B estimates the regression $hpr_{t+1}^{(N)} - y_t^{(1)} = a + b (f_t^{(N+1)} - y_t^{(1)}) + \varepsilon_{t+1}$, where $y_t^{(N)}$ denotes the N -year bond yield at date t ; $f_t^{(N)}$ denotes the N -period ahead forward rate; and $hpr_{t+1}^{(N)}$ denotes the one-year holding period return at date $t + 1$ on an N -year bond. Yields and returns in annual percentages.

forward-spot spread) should be able to tell you that this is a particularly good time for long bonds versus short bonds, as the random walk view of stock prices says that no signal should be able to tell you that this is a particularly good or bad day for stocks versus bonds. However, the coefficients in panel B are all about 1.0. A high forward rate does not indicate that interest rates will be higher one year from now; it seems to indicate that investors will earn that much more by holding long-term bonds.⁷

Of course, there is risk. The R^2 values are all 0.1–0.2, about the same values as the R^2 from the d/p regression at a one-year horizon, so this strategy will often go wrong. Still, 0.1–0.2 is not zero, so the strategy does pay off more often than not, in violation of the expectations hypothesis. Furthermore, the forward-spot spread is a slow-moving variable, typically reversing sign once per business cycle. Thus, the R^2 builds with horizon as with the d/p regression, peaking in the 30 percent range (Fama and French, 1989).

Foreign exchange

Suppose interest rates are higher in Germany than in the U.S. Does this mean that one can earn more money by investing in German bonds? There are several reasons that the answer might be no. First, of course, is default risk. Governments have defaulted on bonds in the past and may do so again. Second, and more important, is the risk of devaluation. If German interest rates are 10 percent and U.S. interest rates are 5 percent, but the euro falls 5 percent relative to the dollar during the year, you make no more money holding the German bonds despite their attractive interest rate. Since lots of investors are making this calculation, it is natural to conclude that an interest rate differential across countries on bonds of similar credit risk should reveal an expectation of currency devaluation. The logic is exactly the same as that of the expectations hypothesis in the term structure. Initially attractive yield or interest rate differentials should be met by an offsetting event so that you make no more money on average in one maturity or currency versus another.⁸

As with the expectations hypothesis in the term structure, the expected depreciation view still constitutes an important first-order understanding of interest rate differentials and exchange rates. For example, interest rates in east Asian currencies were very high on the eve of the recent currency tumbles, and many banks were

making tidy sums borrowing at 5 percent in dollars to lend at 20 percent in local currencies. This suggests that traders were anticipating a 15 percent devaluation, or a smaller chance of a larger devaluation, which is exactly what happened. Many observers attribute high nominal interest rates in troubled economies to “tight monetary policy” aimed at defending the currency. In reality, high nominal rates reflect a large probability of inflation and devaluation—loose monetary policy—and correspond to much lower real rates.

Still, does a 5 percent interest rate differential correspond to a 5 percent expected depreciation, or does some of it represent a high expected return from holding debt in that country’s currency? Furthermore, while expected depreciation is clearly a large part of the interest rate story in high-inflation economies, how does the story play out in economies like the U.S. and Germany, where inflation rates diverge little but exchange rates still fluctuate a large amount?

The first row of table 6 (from Hodrick, 2000, and Engel, 1996) shows the average appreciation of the dollar against the indicated currency over the sample period. The dollar fell against the deutschemark, yen, and Swiss franc, but appreciated against the pound sterling. The second row gives the average interest rate differential—the amount by which the foreign interest rate exceeds the U.S. interest rate.⁹ According to the expectations hypothesis, these two numbers should be equal—interest rates should be higher in countries whose currencies depreciate against the dollar.

TABLE 6

Forward discount puzzle

	Deutsche-mark	Pound sterling	Yen	Swiss franc
Mean appreciation	–1.8	3.6	–5.0	–3.0
Mean interest differential	–3.9	2.1	–3.7	–5.9
<i>b</i> , 1975–89	–3.1	–2.0	–2.1	–2.6
R^2	.026	.033	.034	.033
<i>b</i> , 1976–96	–0.7	–1.8	–2.4	–1.3
<i>b</i> , 10-year horizon	0.8	0.6	0.5	–

Notes: The first row gives the average appreciation of the dollar against the indicated currency, in percent per year. The second row gives the average interest differential—foreign interest rate less domestic interest rate, measured as the forward premium—the 30-day forward rate less the spot exchange rate. The third through sixth rows give the coefficients and R^2 in a regression of exchange rate changes on the interest differential,

$$s_{t+1} - s_t = a + b(r_t^f - r_t^d) + \varepsilon_{t+1},$$

where s = log spot exchange rate, r^f = foreign interest rate, and r^d = domestic interest rate.

Source: Hodrick (2000), Engel (1996), and Meredith and Chinn (1998).

The second row shows roughly the expected pattern. Countries with steady long-term inflation have steadily higher interest rates and steady depreciation. The numbers in the first and second rows are not exactly the same, but exchange rates are notoriously volatile so these averages are not well measured. Hodrick (2000) shows that the difference between the first and second rows is not statistically different from zero. This fact is analogous to the evidence in table 4 that the expectations hypothesis works well *on average* for U.S. bonds.

As in the case of bonds, however, we can ask whether times of *temporarily* higher or lower interest rate differentials correspond to times of above- and below-average depreciation as they should. The third and fifth rows of table 6 update Fama's (1984) regression tests. The number here should be +1.0 in each case—1 percentage point extra interest differential should correspond to 1 percentage point extra expected depreciation. On the contrary, as table 6 shows, a higher than usual interest rate abroad seems to lead to further *appreciation*. This is the *forward discount puzzle*. See Engel (1996) and Lewis (1995) for recent surveys of the avalanche of academic work investigating whether this puzzle is really there and why.

The R^2 values shown in table 6 are quite low. However, like d/p and the term spread, the interest differential is a slow-moving forecasting variable, so the return forecast R^2 builds with horizon. Bekaert and Hodrick (1992) report that the R^2 rises to the 30 percent to 40 percent range at six-month horizons and then declines. That's high, but not 100 percent; taking advantage of any predictability strategy is quite risky.

The puzzle does *not* say that one earns more by holding bonds from countries with higher interest rates than others. Average inflation, depreciation, and interest rate differentials line up as they should. The puzzle *does* say that one earns more by holding bonds from countries whose interest rates are *higher than usual* relative to U.S. interest rates (and vice versa). The fact that the “usual” rate of depreciation and interest differential changes through time will, of course, diminish the out-of-sample performance of these trading rules.

One might expect that exchange rate depreciation works better for long-run exchange rates, as the expectations hypothesis works better for long-run interest rate changes. The last row of table 6, taken from Meredith and Chinn (1998) verifies that this is so. Ten-year exchange rate changes are correctly forecast by the interest differentials of ten-year bonds.

Mutual funds

Studying the returns of funds that follow a specific strategy gives us a way to assess whether that strategy works in practice, after transaction costs and other trading realities are taken into account. Studying the returns of actively managed funds tells us whether the time, talent, and effort put into picking securities pays off. Most of the literature on evaluating fund performance is devoted to the latter question.

A large body of empirical work, starting with Jensen (1969), finds that actively managed funds, on average, underperform the market index. I use data from Carhart (1997), whose measures of fund performance account for *survivor bias*. Survivor bias arises because funds that do badly go out of business. Therefore, the average fund that is alive at any point in time has an artificially good track record.


As with the stock portfolios in figure 1, the fund data in figure 7 show a definite correlation between beta and average return: Funds that did well took on more market risks. A cross-sectional regression line is a bit flatter than the line drawn through the Treasury bill and market return, but this is a typical result of measurement error in the betas. (The data are annual, and many funds are only around for a few years, contributing to beta measurement error.) The average fund underperforms the line connecting Treasury bills and the market index by 1.23 percent per year (that is, the average alpha is -1.23 percent).

The wide dispersion in fund average returns in figure 7 is a bit surprising. Average returns vary across funds almost as much as they do across individual stocks. This fact implies that the majority of funds are *not* holding well-diversified portfolios that would reduce return variation, but rather are loading up on specific bets.

Initially, the fact that the *average* fund underperforms the market seems beside the point. Perhaps the average fund is bad, but we want to know whether the good funds are any good. The trouble is, we must somehow distinguish skill from luck. The only way to separate skill from luck is to group funds based on some ex-ante observable characteristic, and then examine the average performance of the group. Of course, skillful funds should have done better, on average, in the past, and should continue to do better in the future. Thus, if there is skill in stock picking, we should see some persistence in fund performance. However, a generation of empirical work found no persistence at all. Funds that did well in the past were no more likely to do well in the future.

FIGURE 7

Average returns of mutual funds vs. market betas



Notes: Average returns of mutual funds over the Treasury bill rate versus their market betas. Sample consists of all funds with average total net assets greater than \$25 million and more than 25 percent of their assets in stocks, in the Carhart (1996) database. Data sample 1962–96. The average excess return is computed as $E(R - R_f) = a_i + \beta_i \times 9\%$. a_i and β_i are computed from a time-series regression of fund annual excess returns on market annual excess returns over the life of the fund. The o, +, and x labels in the figure sort funds into thirds based on their regression coefficient h on the Fama–French value (HML) portfolio. The breakpoints are $h = -0.084, 0.34$. The dashed line gives the fit of a cross-sectional ordinary least squares regression of a_i on β_i ; The solid line connects the Treasury bill (β and excess return = 0) and the market return ($\beta = 1$, excess return = 9%).

Since the average fund underperforms the market, and fund returns are not predictable, we conclude that active management does not generate superior performance, especially after transaction costs and fees. This fact is surprising. Professionals in almost any field do better than amateurs. One would expect that a trained experienced professional who spends all day reading about markets and stocks should be able to outperform simple indexing strategies. Even if entry into the industry is so easy that the *average* fund does not outperform simple indexes one would expect a few stars to outperform year after year, as good teams win championship after championship. Alas, the contrary fact is the result of practically every investigation, and even the anomalous results document very small effects.

Funds and value

Given the value, small-firm, and predictability effects, the idea that funds cluster around the market line is quite surprising. All of these new facts imply inescapably that there are simple, mechanical strategies that can give a risk/reward ratio greater than that of buying and holding the market index. Fama and

French (1993) report that the HML portfolio alone gives nearly double the market Sharpe ratio—the same average return at half the standard deviation. Why don't funds cluster around a risk/reward line significantly above the market's?

Of course, we should not expect *all* funds to cluster around a higher risk/reward tradeoff. The average investor holds the market, and if funds are large enough, so must the average fund. Index funds, of course, will perform like the index. Still, the typical actively managed fund advertises high mean and, perhaps, low variance. No fund advertises cutting average returns in half to spare investors exposure to nonmarket sources of risk. Such funds, apparently aimed at mean–variance investors, should cluster around the highest risk/reward tradeoff available from mechanical strategies (and more, if active management does any good). Most troubling, funds who *say* they follow value strategies don't outperform the market either. For example, Lakonishok, Shleifer, and Vishny (1992, table 3) find that the average value fund underperforms the S&P500 by 1 percent just like all the others.

We can resolve this contradiction if we think that fund managers were simply unaware of the possibilities offered by our new facts, and so (despite the advertising) were not really following them. That seems to be the implication of figure 7, which sorts funds by their HML beta. One would expect the high-HML beta funds to outperform the market line. But the cutoff for the top one-third of funds is only a HML beta of 0.3, and even that may be high (many funds don't last long, so betas are poorly measured; the distribution of measured betas is wider than the actual distribution). Thus, the “value funds” were really not following the “value strategy” that earns the HML returns; if they were doing so they would have HML betas of 1.0. Similarly, Lakonishok, Shleifer, and Vishny's (1992) documentation of value funds' underperformance reveals that their market beta is close to 1.0. These results imply that value funds are not really following a value strategy, since their returns correlate with the market portfolio and not the value portfolio.

Interestingly, the number of value and small-cap funds (as revealed by their betas, not their marketing claims) is increasing quickly. Before 1990, 14 percent of funds had measured SMB betas greater than 1.0,

and 12 percent had HML betas greater than 1.0. In the full sample, both numbers have *doubled* to 22 percent and 23 percent. This trend suggests that funds will, in the future, be much less well described by the market index.

The view that funds were unaware of value strategies, and are now moving quickly to exploit them, can explain why most funds still earn near the market return, rather than the higher value return. However, this view contradicts the view that the value premium is an equilibrium risk premium, that is, that everyone knew about the value returns but chose not to invest all along because they feared the risks of value strategies. If it is not an equilibrium risk premium, it won't last long.

Persistence in fund returns

The fund counterpart to momentum in stock returns has been more extensively investigated than the value and size effects. Fund returns have also been found to be persistent. Since such persistence can be interpreted as evidence for persistent skill in picking stocks, it is not surprising that it has attracted a great deal of attention, starting with Hendricks, Patel, and Zeckhauser (1993).

Table 7, taken from Carhart (1997), shows that a portfolio of the best-performing one-thirtieth of funds last year outperforms a portfolio of the worst-performing one-thirtieth of funds by 1 percent per month (column 2). This is about the same size as the momentum effect in stocks, and similarly results from a small autocorrelation plus a large standard deviation in

individual fund returns. This result verifies that mutual fund performance is persistent.

Perhaps the funds that did well took on more market risks, raising their betas and, hence, average returns in the following year. The third column in table 7 shows that this is not the case. The cross-sectional variation in fund average returns has nothing to do with market betas. Just as in the case of individual stock returns, we have to understand fund returns with multifactor models, if at all.

The last column of table 7 presents alphas (intercepts, the part of average return not explained by the model) from a model with four factors—the market, the Fama–French HML and SMB factors, and a momentum factor, PR1YR, that is long NYSE stocks that did well in the last year and short NYSE stocks that did poorly in the last year. In general, one should object to the inclusion of so many factors and such ad-hoc factors. However, this is a *performance attribution* rather than an *economic explanation* use of a multifactor model. We want to know whether fund performance, and persistence in fund performance in particular, is due to persistent stock-picking skill or to mechanical strategies that investors could just as easily follow on their own, without paying the management costs associated with investing through a fund. For this purpose, it does not matter whether the “factors” represent true, underlying sources of macroeconomic risks.

The alphas in the last column of table 7 are almost all about 1 percent to 2 percent per year negative. Thus, Carhart's model explains that the persistence in fund

performance is due to persistence in the underlying stocks, not persistent stock-picking skill. These results support the old conclusion that actively managed funds underperform mechanical indexing strategies. There is some remaining puzzling persistence, but it is all in the large *negative* alphas of the bottom one-tenth to bottom one-thirtieth of performers, which lose money year after year. Carhart also shows that the persistence of fund performance is due to momentum in the underlying stocks, rather than momentum funds. If, by good luck, a fund happened to pick stocks that went up last year, the portfolio will continue to go up a bit this year.

In sum, the new research does nothing to dispel the disappointing view of active management. However, we discover that passively managed “style”

TABLE 7			
Portfolios of mutual funds formed on previous year's return			
Last year rank	Average return	CAPM alpha	4-factor alpha
	(----- percent -----)		
1/30	0.75	0.27	-0.11
1/10	0.68	0.22	-0.12
5/10	0.38	-0.05	-0.14
9/10	0.23	-0.21	-0.20
10/10	0.01	-0.45	-0.40
30/30	-0.25	-0.74	-0.64

Notes: Each year, mutual funds are sorted into portfolios based on the previous year's return. The rank column gives the rank of the selected portfolio. For example, 1/30 is the best performing portfolio when funds are divided into 30 categories. Average return gives the average monthly return in excess of the T-bill rate of this portfolio of funds for the following year. Four-factor alpha gives the average return less the predictions of a multifactor model that uses the market, the Fama–French HML and SMB portfolios, and portfolio PR1YR which is long NYSE stocks that did well in the last year and short NYSE stocks that did poorly in the last year. Source: Carhart (1997).

portfolios can earn returns that are not explained by the CAPM.

Catastrophe insurance

A number of prominent funds have earned very good returns (and others, spectacular losses) by following strategies such as *convergence trades* and implicit *put options*. These strategies may also reflect high average returns as compensation for nonmarket dimensions of risk. They have not been examined at the same level of detail as the value and small-cap strategies, so I offer a possible interpretation rather than a documented one.

Convergence trades take strong positions in very similar securities that have small price differences. For example, a 29.5-year Treasury bond typically trades at a slightly higher yield (lower price) than a 30-year Treasury bond. (This was the most famous bet placed by LTCM. See Lewis, 1999.) A convergence trade puts a strong short position on the expensive security and a strong long position on the cheap security. This strategy is often mislabeled an “arbitrage.” However, the securities are similar, not identical. The spread between 29.5- and 30-year Treasury bonds reflects the lower liquidity of the shorter maturity and the associated difficulty of selling it in a financial panic. It is possible for this spread to widen. Nonetheless, panics are rare, and the average returns in all the years when they do not happen may more than make up for the spectacular losses when they do.

Put options protect investors from large price declines. The *volatility smile* in put option prices reflects the surprisingly high prices of such options, compared with the small probability of large market collapses (even when one calibrates the probability directly, rather than using the log-normal distribution of the Black–Scholes formula). Writers of out-of-the-money puts collect a fee every month; in a rare market collapse they will pay out a huge sum, but if the probability of the collapse is small enough, the average returns may be quite good.

All of these strategies can be thought of as *catastrophe insurance* (Hsieh and Fung, 1999). Most of the time they earn a small premium. Once in a great while they lose a lot, and they lose a lot in times of financial catastrophe, when most investors are really anxious that the value of their investments not evaporate. Therefore, it is economically plausible that these strategies can earn positive average returns, even when we account for stock market risk via the CAPM and we correctly measure the small probabilities of large losses.

The difficulty in empirically estimating the true average return of such strategies, of course, is that

rare events are rare. Many long samples will give a false sense of security because “the big one” that justifies the premium happened not to hit.

The value, yield curve, and foreign exchange strategies I survey above also exhibit features of catastrophe insurance. Value stocks may earn high returns because distressed stocks will all go bankrupt in a financial panic. Buying bonds of countries with high interest rates leaves one open to the small chance of a large devaluation, and such devaluation is especially likely to happen in a global financial panic. Similarly, buying long-term bonds in the depth of a recession when the yield curve is upward sloping may expose one to a small risk of a large inflation.

If these interpretations bear out, they also suggest that the premiums—the average returns from holding stocks sensitive to HML or from following the bond and foreign exchange strategies—may be overstated in the data. The markets have had an unusually good 50 years, and devastating financial panics have not happened.

Implications of the new facts

While the list of new facts appears long, similar patterns show up in every case. Prices reveal slow-moving market expectations of subsequent returns, because potential offsetting events seem sluggish or absent. The patterns suggest that investors can earn substantial average returns by taking on the risks of recession and financial stress. In addition, there is a small positive autocorrelation of high-frequency returns.

The effects are not completely new. We have known since the 1960s that high-frequency returns are slightly predictable, with R^2 of 0.01 to 0.1 in daily to monthly returns. These effects were dismissed because there didn’t seem to be much one could do about them. A 51/49 bet is not very attractive, especially if there is any transaction cost. Also, the increased Sharpe ratio (mean excess return/standard deviation) from exploiting predictability is directly related to the forecast R^2 , so a tiny R^2 , even if exploitable, did not seem important. Now, we have a greater understanding of the potential importance of these effects and their economic interpretations.

For price effects, we now realize that the R^2 rises with horizon when the forecasting variables are slow-moving. Hence, a small R^2 at short horizons can mean a really substantial R^2 in the 30 percent to 50 percent range at longer horizons. Also, the nature of these effects suggests the kinds of additional sources of priced risk that theorists had anticipated for 20 years. For momentum effects, the ability to sort stocks and

funds into momentum-based portfolios means that very small predictability times portfolios with huge past returns gives important subsequent returns, though it is not totally clear that this amplification of the small predictability really does survive transaction costs.

Price-based forecasts

If expected returns rise, prices are driven down, since future dividends or other cash flows are discounted at a higher rate. A “low” price, then, can *reveal* a market expectation of a high expected or required return.¹⁰


Most of our results come from this effect. Low price/dividend, price/earnings, or price/book values signal times when the market as a whole will have high average returns. Low market value (price times shares) relative to book value signals securities or portfolios that earn high average returns. The “small-firm” effect derives from low prices—other measures of size such as number of employees or book value alone have no predictive power for returns (Berk, 1997). The “five-year reversal” effect derives from the fact that five years of poor returns lead to a low price. A high long-term bond yield means that the price of long-term bonds is “low,” and this seems to signal a time of good long-term bond returns. A high foreign interest rate means a low price on foreign bonds, and this seems to indicate good returns on the foreign bonds.

The most natural interpretation of all these effects is that the expected or required return—the risk premium—on individual securities as well as the market as a whole varies slowly over time. Thus we can track market expectations of returns by watching price/dividend, price/earnings, or book/market ratios.

Absent offsetting events

In each case, an apparent difference in yield should give rise to an offsetting movement, but does not seem to do so. Something *should* be predictable so that returns are *not* predictable, and it is not. Figure 8 provides a picture of the results in table 5. Suppose that the yield curve is upward sloping as in panel A. What does this mean? If the expectations model were true, the forward rates plotted against maturity would translate one for one to the forecast of future spot rates in panel B, as plotted in the black line marked “Expectations model.” A high long-term bond yield relative to short-term bond yields should not mean a higher expected long-term bond return. Subsequent short rates should rise, cutting off the one-period advantage of long-term bonds and raising the multi-year advantage of short-term bonds.

In figure 8, panel b, the colored line marked “Estimates” shows the actual forecast of future spot interest rates from the results in table 5. The essence of the phenomenon is *sluggish adjustment* of the short rates. The short rates do eventually rise to meet the forward rate forecasts, but not as quickly as the forward rates predict they should. Short-term yields *should* be forecastable so that returns are *not* forecastable. In fact, yields are almost unforecastable, so, mechanically, bond returns are. The roughly 1.0 coefficients in panel B of table 5 mean that a 1 percentage point increase in the forward rate translates into a 1 percentage point increase in expected return. It seems that old fallacy of confusing bond *yields* with their expected *returns* for the first year contains a grain of truth.



In the same way, a high dividend yield on a stock or portfolio should mean that dividends grow more slowly over time, or, for individual stocks, that the firm has taken on more market risk and will have a higher market beta. These tendencies seem to be completely absent. Dividend/price ratios do not seem to forecast dividend growth and, hence, (mechanically) they forecast returns. The one-year coefficient in table 1 is very close to 1.00, meaning that a 1 percentage point increase in the dividend yield translates into a 1 percentage point increase in return. It seems that the old fallacy of confusing increased dividend yield with increased total return does contain a grain of truth.

A high foreign interest rate relative to domestic interest rates should not mean a higher expected return. We should see, on average, an offsetting depreciation. But here, the coefficients are even larger than 1.0. An interest rate differential seems to predict a further *appreciation*. It seems that the old fallacy of confusing interest rate differentials across countries with expected returns, forgetting about depreciation, also contains a grain of truth.

Economic interpretation

The price-based predictability patterns suggest a premium for holding risks related to recession and economy-wide financial distress. Stock and bond predictability are linked: The term spread (forward-spot, or long yield–short yield) forecasts stock returns as well as bond returns (Fama and French, 1989). Furthermore, the term spread is one of the best variables for forecasting business cycles. It rises steeply at the bottom of recessions and is inverted at the top of a boom. Return forecasts are high at the bottom of a business cycle and low at the top of a boom. Value and small-cap stocks are typically distressed. Empirically successful economic models of the recession and distress premiums are still in their infancy (Campbell and Cochrane, 1999, is a start), but the story is at least plausible and the effects have been expected by theorists for a generation.

To make this point come to life, think concretely about what you have to do to take advantage of the predictability strategies. You have to buy stocks or long-term bonds at the bottom, when stock prices are low after a long and depressing bear market, in the bottom of a recession or the peak of a financial panic. This is a time when few people have the guts or the wallet to buy risky stocks or risky long-term bonds. Looking across stocks rather than over time, you have to invest in value or small-cap companies, with years of poor past returns, poor sales, or on the edge of bankruptcy. You have to buy stocks that everyone else thinks are dogs. Then, you have to sell stocks

and long-term bonds in good times, when stock prices are high relative to dividends, earnings, and other multiples and the yield curve is flat or inverted so that long-term bond prices are high. You have to sell the popular growth stocks, with good past returns, good sales, and earnings growth.

You have to sell now, and the stocks that you should sell are the blue-chips that everyone else seems to be buying. In fact, the market timing strategies said to sell long ago; if you did so, you would have missed much of the runup in the Dow past the 6,000 point. Value stocks too have missed most of the recent market runup. However, this shouldn't worry you—a strategy that holds risks uncorrelated with the market must underperform the market close to half of the time.

If this feels uncomfortable, what you're feeling is risk. If you're uncomfortable watching the market pass you by, perhaps you *don't* really only care about long-run mean and variance; you also care about doing well when the market is doing well. If you want to stay fully invested in stocks, perhaps you too feel the time-varying aversion to or exposure to risk that drives the average investor to stay fully invested despite low prospective returns.

This line of explanation for the foreign exchange puzzle is still a bit farther off (see Engel, 1996, for a survey; Atkeson, Alvarez, and Kehoe, 1999, offer a recent stab at an explanation). The strategy leads investors to invest in countries with high interest rates. High interest rates are often a sign of monetary instability or other economic trouble, and thus may mean that the investments are more exposed to the risks of global financial stress or a global recession than are investments in the bonds of countries with low interest rates, which are typically enjoying better times.

Return correlation

Momentum and persistent fund performance explained by a momentum factor are different from the price-based predictability results. In both cases, the underlying phenomenon is a small predictability of high-frequency returns. The price-based predictability strategies make this predictability important by showing that, with a slow-moving forecasting variable, the R^2 builds over horizon. Momentum, however, is based on a fast-moving forecast variable—the previous year's return. Therefore, the R^2 declines rather than building with horizon. Momentum makes the small predictability of high-frequency returns significant in a different way, by forming portfolios of extreme winners and losers. The large volatility of returns means that the extreme portfolios will have extreme past returns, so only a small continuation of past returns gives a large current return.

It would be appealing to understand momentum as a reflection of slowly time-varying average expected returns or risk premiums, like the price-based predictability strategies. If a stock's average return rises for a while, that should make returns higher both today and tomorrow. Thus, a portfolio of past winners will contain more than its share of stocks that performed well because their average returns were higher, along with stocks that performed well due to luck. The average return of such a portfolio should be higher tomorrow as well.

Unfortunately, this story has to posit a substantially different view of the underlying process for varying expected returns than is needed to explain everything else. The trouble is that a surprise increase in expected returns means that prices will fall, since dividends are now discounted at a greater rate. This is the phenomenon we have relied on to explain why *low* price/dividend, price/earnings, book/market, value, and size forecast *higher* subsequent returns. Therefore, *positive* correlation of *expected* returns typically yields a *negative* correlation of *realized* returns. To get a positive correlation of realized returns out of slow expected return variation, you have to imagine that an increase in average returns today is either highly correlated with a decrease in expected future dividend growth or with a decrease in expected returns in the distant future (an impulse response that starts positive but is negative at long horizons). Campbell, Lo, and MacKinlay (1997) provide a quantitative exposition of these effects.

Furthermore, momentum returns have not yet been linked to business cycles or financial distress in even the informal way that I suggested for price-based strategies. Thus, momentum still lacks a plausible economic interpretation. To me, this adds weight to the view that it isn't there, it isn't exploitable, or it represents a small illiquidity (tax-loss selling of small illiquid stocks) that will be quickly remedied once a few traders understand it.

Remaining doubts

The size of all these effects is still somewhat in question. It is always hard to measure average returns of risky strategies. The standard formula $\sigma / \sqrt{T}$ for the standard error of the mean, together with the high volatility σ of any strategy, means that one needs 25 years of data to even start to measure average returns. With $\sigma = 16$ percent (typical of the index), even $T = 25$ years means that one standard error is $16/5 \approx 3$ percent per year, and a two-standard error confidence interval runs plus or minus 6 percentage points. This is not much smaller than the average returns we are trying to measure. In addition, all of

these facts are highly influenced by the small probability of rare events, which makes measuring average returns even harder.

Finally, viewed the right way, we have very few data points with which to evaluate predictability. The term premium and interest rate differentials only change sign with the business cycle, and the dividend/price ratio only crosses its mean once every generation. The history of interest rates and inflation in the U.S. is dominated by the increase, through two recessions, to a peak in 1980 and then a slow decline after that.

Many of the anomalous risk premiums seem to be declining over time. Figure 6 shows the decline in the HML and SMB premiums, and the same may be true of the predictability effects. The last three years of high market returns have cut the estimated return predictability from the dividend/price ratio in *half*. This fact suggests that at least some of the premium the new strategies yielded in the past was due to the fact that they were simply overlooked.

Was it really clear to average investors in 1947 or 1963 (the beginning of the data samples) that stocks would earn 9 percent over bonds, and that the strategy of buying distressed small stocks would double even that return for the same level of risk? Would average investors have changed their portfolios with this knowledge? Or would they have stayed pat, explaining that these returns are earned as a reward for risk that they were not willing to take? Was it clear that buying stocks at the bottom in the mid-1970s would yield so much more than even that high average return? If we interpret the premiums measured in sample as true risk premiums, the answer must be yes. If the answer is no, then at least some part of the premium was luck and will disappear in the future.

Since the premiums are hard to measure, one is tempted to put less emphasis on them. However, they are crucial to our interpretation of the facts. The CAPM is perfectly consistent with the fact that there are additional sources of *common* variation. For example, it was long understood that stocks in the same industry move together; the fact that value or small stocks also move together need not cause a ripple. The surprise is that investors seem to earn an average return premium for holding these additional sources of common movement, whereas the CAPM predicts that (given beta) they should have no effect on a portfolio's average returns.

The behavior of funds also suggests the "overlooked strategy" interpretation. As explained earlier, fund returns still cluster around the market line. It turns out that very few fund returns actually followed the value or other return-enhancing strategies. However, the number of small, value, and related funds—funds

that actually do follow the strategies—has increased dramatically in recent years. It might be possible to explain this in a way consistent with the idea that investors knew the premiums were there all along, but such an argument is obviously strained.

Conclusion

In sum, it now seems that investors can earn a substantial premium for holding dimensions of risk unrelated to market movements, such as recession-related or distress-related risk. Investors earn these premiums by following strategies, such as value and

growth, market-timing possibilities generated by return predictability, dynamic bond and foreign exchange strategies, and maybe even a bit of momentum. The exact size of the premiums and the economic nature of the underlying risks is still a bit open to question, but researchers are unlikely to go back to the simple view that returns are independent over time and that the CAPM describes the cross section.

The next question is, What should investors do with this information? The article, “Portfolio advice for a multifactor world,” also in this issue, addresses that question.

NOTES

¹The market also tends to go down in recessions; however recessions can be unusually severe or mild for a given level of market return. What counts here is the severity of the recession for a given market return. Technically, we are considering betas in a multiple regression that includes both the market return and a measure of recessions. See box 1.

²I thank Gene Fama for providing me with these data.

³The rest of the paragraph is my interpretation, not Fama and French’s. They focus on the firm’s financial distress, while I focus on the systematic distress, since idiosyncratic distress cannot deliver a risk price.

⁴Fama and French do not provide direct measures of standard deviations for these portfolios. One can infer, however, from the betas, R^2 values, and standard deviation of the market and factor portfolios that the standard deviations are roughly one to two times that of the market return, so Sharpe ratios of these strategies are comparable to that of the market return in sample.

⁵We are looking for

$$E(r|r \geq x) = \frac{\int_x^\infty rf(r)dr}{\int_x^\infty f(r)dr},$$

where x is defined as the top one-tenth cutoff,

$$\int_x^\infty f(r)dr = \frac{1}{10}.$$

With a normal distribution, $x = 1.2816\sigma$ and $E(r|r \geq x) = 1.755\sigma$.

⁶The index autocorrelations suffer from some upward bias since some stocks do not trade every day. Individual stock autocorrelations are generally smaller, but are enough to account for the momentum effect.

⁷Panel B is really not independent evidence, since the coefficients in panels A and B of table 5 are mechanically linked. For example, $1.14 + (-0.14) = 1.0$, and this holds as an accounting identity. Fama and Bliss (1987) call them “complementary regressions.”

⁸As with bonds, the expectations hypothesis is slightly different from pure risk neutrality since the expectation of the log is not the log of the expectation. Again, the size of the phenomena we study swamps this distinction.

⁹The data are actually the spread between the forward exchange rate and the spot exchange rate, but this quantity must equal the interest rate differential in order to preclude arbitrage.

¹⁰This effect is initially counterintuitive. One might suppose that a higher average return would attract investors, raising prices. But the higher prices, for a given dividend stream, must reduce subsequent average returns. High average returns persist, in equilibrium, when investors fear the increased risks of an asset and try to sell, lowering prices.

REFERENCES

- Atkeson, Andrew, Fernando Alvarez, and Patrick Kehoe**, 1999, “Volatile exchange rates and the forward premium anomaly: A segmented asset market view,” University of Chicago, working paper.
- Banz, R. W.**, 1981, “The relationship between return and market value of common stocks,” *Journal of Financial Economics*, Vol. 9, No. 1, pp. 3–18.
- Bekaert, Geert, and Robert J. Hodrick**, 1992, “Characterizing predictable components in excess returns on equity and foreign exchange markets,” *Journal of Finance*, Vol. 47, No. 2, June, pp. 467–509.
- Berk, Jonathan**, 1997, “Does size really matter?,” *Financial Analysts Journal*, Vol. 53, September/October, pp. 12–18.
- Campbell, John Y.**, 1996, “Understanding risk and return,” *Journal of Political Economy*, Vol. 104, No. 2, April, pp. 298–345.

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*, Vol. 107, No. 2, April, pp. 205–251.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton, NJ: Princeton University Press.

Carhart, Mark M., 1997, "On persistence in mutual fund performance," *Journal of Finance*, Vol. 52, No. 1, March, pp. 57–82.

Chen, Nai-Fu, Richard Roll, and Stephen A. Ross, 1986, "Economic forces and the stock market," *Journal of Business*, Vol. 59, No. 3, July, pp. 383–403.

Cochrane, John H., 1997, "Where is the market going? Uncertain facts and novel theories," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 21, No. 6, November/December, pp. 3–37.

_____, 1996, "A cross-sectional test of an investment-based asset pricing model," *Journal of Political Economy*, Vol. 104, No. 3, June, pp. 572–621.

_____, 1991, "Volatility tests and efficient markets: Review essay," *Journal of Monetary Economics*, Vol. 27, No. 3, June, pp. 463–485.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, "Investor psychology and security market under- and overreactions," *Journal of Finance*, Vol. 3, No. 6, December, pp. 1839–1885.

DeBondt, Werner F. M., and Richard H. Thaler, 1985, "Does the stock market overreact?," *Journal of Finance*, Vol. 40, No. 3, pp. 793–805.

Engel, Charles, 1996, "The forward discount anomaly and the risk premium: A survey of recent evidence," *Journal of Empirical Finance*, Vol. 3, pp. 123–192.

Fama, Eugene F. 1991, "Efficient markets II," *Journal of Finance*, Vol. 46, No. 5, December, pp. 1575–1617.

_____, 1984, "Forward and spot exchange rates," *Journal of Monetary Economics*, Vol. 14, No. 3, November, pp. 319–338.

_____, 1970, "Efficient capital markets: A review of theory and empirical work," *Journal of Finance*, Vol. 25, No. 2, May, pp. 383–417.

_____, 1965, "The behavior of stock market prices," *Journal of Business*, Vol. 38, No. 1, pp. 34–105.

Fama, Eugene F., and Robert R. Bliss, 1987, "The information in long-maturity forward rates," *American Economic Review*, Vol. 77, No. 4, September, pp. 680–692.

Fama, Eugene F., and Kenneth R. French, 1997, "Industry costs of equity," *Journal of Financial Economics*, Vol. 43, No. 2, February, pp. 153–193.

_____, 1996, "Multifactor explanations of asset-pricing anomalies," *Journal of Finance*, Vol. 51, No. 1, March, pp. 55–84.

_____, 1995, "Size and book-to-market factors in earnings and returns," *Journal of Finance*, Vol. 50, No. 1, March, pp. 131–155.

_____, 1993, "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, Vol. 33, No. 1, February, pp. 3–56.

_____, 1989, "Business conditions and expected returns on stocks and bonds," *Journal of Financial Economics*, Vol. 25, No. 1, November, pp. 23–49.

Heaton, John, and Deborah Lucas, 1997, "Portfolio choice and asset prices: The importance of entrepreneurial risk," Northwestern University, manuscript.

Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser, 1993, "Hot hands in mutual funds: Short-term persistence of performance," *Journal of Finance*, Vol. 48, No. 1, March, pp. 93–130.

Hodrick, Robert, 2000, *International Financial Management*, Englewood Cliffs, NJ: Prentice-Hall, forthcoming.

Hsieh, David, and William Fung, 1999, "Hedge fund risk management," Duke University, working paper.

Jagannathan, Ravi, and Zhenyu Wang, 1996, "The conditional CAPM and the cross-section of expected returns," *Journal of Finance*, Vol. 51, No. 1, March, pp. 3–53.

Jegadeesh, Narasimham, and Sheridan Titman, 1993, "Returns to buying winners and selling losers: Implications for stock market efficiency," *Journal of Finance*, Vol. 48, No. 1, March, pp. 65–91.

Jensen, Michael C., 1969, "The pricing of capital assets and evaluation of investment portfolios," *Journal of Business*, Vol. 42, No. 2, April, pp. 167–247.

Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny, 1992, "The structure and performance of the money management industry," *Brookings Papers on Economic Activity: Microeconomics 1992*, Washington, DC, pp. 339–391.

LeRoy, Stephen F., and Richard D. Porter, 1981, "The present-value relation: Tests based on implied variance bounds," *Econometrica*, Vol. 49, No. 3, May, pp. 555–574.

Lewis, Karen, K., 1995, "Puzzles in international financial markets," in *Handbook of International Economics*, Vol. 3, G. Grossman and K. Rogoff (eds.), Amsterdam, New York, and Oxford: Elsevier Science B.V., pp. 1913–1971.

Lewis, Michael, 1999, "How the eggheads cracked," *New York Times Magazine*, January 24, pp. 24–42.

Liew, Jimmy, and Maria Vassalou, 1999, "Can book-to-market, size and momentum be risk factors that predict economic growth?," Columbia University, working paper.

MacKinlay, A. Craig, 1995, "Multifactor models do not explain deviations from the CAPM," *Journal of Financial Economics*, Vol. 38, No. 1, pp. 3–28.

Malkiel, Burton, 1990, *A Random Walk Down Wall Street*, New York: Norton.

Markowitz, H., 1952, "Portfolio selection," *Journal of Finance*, Vol. 7, No. 1, March, pp. 77–99.

Meredith, Guy, and Menzie D. Chinn, 1998, "Long-horizon uncovered interest rate parity," National Bureau of Economic Research, working paper, No. 6797.

Merton, Robert C., 1973, "An intertemporal capital asset pricing model," *Econometrica*, Vol. 41, No. 5, September, pp. 867–887.

_____, 1971, "Optimum consumption and portfolio rules in a continuous time model," *Journal of Economic Theory*, Vol. 3, No. 4, pp. 373–413.

_____, 1969, "Lifetime portfolio selection under uncertainty: The continuous time case," *Review of Economics and Statistics*, Vol. 51, No. 3, August, pp. 247–257.

Moskowitz, Tobias, and Mark Grinblatt, 1999, "Tax loss selling and return autocorrelation: New evidence," University of Chicago, working paper.

_____, 1998, "Do industries explain momentum?," University of Chicago, CRSP working paper, No. 480.

New York Times Company, 1999, "Mutual funds report: What's killing the value managers?," *New York Times*, April 4, Section 3, p. 29.

Reyffman, Alexander, 1997, "Labor market risk and expected asset returns," University of Chicago, Ph.D. thesis.

Ross, S. A., 1976, "The arbitrage theory of capital asset pricing," *Journal of Economic Theory*, Vol. 13, No. 3, December, pp. 341–360.

Samuelson, Paul A., 1969, "Lifetime portfolio selection by dynamic stochastic programming," *Review of Economics and Statistics*, Vol. 51, No. 3, August, pp. 239–246.

Sargent, Thomas J., 1993, *Bounded Rationality in Macroeconomics*, Oxford: Oxford University Press.

Shiller, Robert J., 1981, "Do prices move too much to be justified by subsequent changes in dividends?," *American Economic Review*, Vol. 71, No. 3, June, pp. 421–436.