NBER WORKING PAPER SERIES

NEW FACTS IN FINANCE

John H. Cochrane

Working Paper 7169
http://www.nber.org/papers/w7169

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1999

I thank Andrea Eisfeldt for research assistance and Gene Fama, Helen Koshy, and David Marshall for comments. My research is supported by the CRSP, Graduate School of Business, and by a grant from the National Science Foundation administered by the NBER. Data and updated copies of this paper can be found at http://www-gsb.uchicago.edu/fac/john.cochrane/ All opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

New Facts in Finance
John H. Cochrane
NBER Working Paper No. 7169
June 1999
JEL No. G00

ABSTRACT

The last 15 years have seen a revolution in the way financial economists understand the world around us. 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 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" have supplanted the CAPM to explain them. 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 premia in bond and foreign exchange markets as well as in stock markets. Once, we thought that mutual fund average returns were well explained by the CAPM. Now, we recognize ``value'' and other high return strategies in funds, and slight persistence in fund performance.

In this article, I survey these new facts. I show how they are related. 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.

John H. Cochrane
Graduate School of Business
University of Chicago
1101 E. 58th. St.
Chicago IL 60637
and NBER
john.cochrane@gsb.uchicago.edu

1 Overview

Up until the mid- 1980s, financial economists had a fairly confident view of the investment world. The view was built on three bedrocks:

1. The capital asset pricing model (CAPM) is a good measure of risk and thus a good explanation why some stocks, portfolios, strategies or funds (assets, generically) 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 together 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. In particular,

   (a) Stock returns are very close to unpredictable. 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 is nearly useless. Any apparent predictability is either a statistical artifact which will quickly vanish out of sample, or cannot be exploited after transactions costs.

   (b) 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 returns on long term bonds are any higher than those on short term bonds. Rather, it means that short term interest rates are expected to rise in the future, so you expect to earn about the same amount on short term or long term bonds at any horizon.

   (c) 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.

   (d) 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 old coin, over and over again.

3. Professional managers do not reliably outperform simple indices 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 good and bad luck. Managers who do well in one year are not more likely to do better than average the next year. The average actively-managed fund does about 1% *worse* than the market index. The more actively afund trades trading, the lower returns to investors.

Together, these views reflected a guiding principle that asset markets are, to a good approximation, *informationally efficient.* (Fama 1970, 1991.) This statement means that market prices already contain most information about fundamental value. Informational efficiency in turn derives from *competition.* The business of discovering information about the value of traded assets is extremely competitive, so there are no easy quick profits to be made, as there are not in every other well-established and competitive industry. The only way to earn large returns is by taking on additional risk.

These statements are not doctrinaire beliefs. Rather, they summarize the findings of a quarter-century of extensive and careful empirical work. However, every single one of them has now been extensively revised, by a new generation of empirical research. These findings need not overturn the cherished view that markets are reasonable competitive and therefore reasonably efficient. It does substantially enlarge our view of what activities provide rewards for holding risks, and it challenges our economic understanding of those risk premia.

Now, we know that:

1. There are assets, portfolios, funds, and strategies whose average returns cannot be explained by their beta, or tendency to move with the market as a whole. 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 additional risk factors in addition to movements in the market as a whole. (See Box 1.)

2. Returns are predictable. In particular,

   (a) Variables including the dividend/price ratio and term premium can in fact 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.

   (b) 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.

(c) 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.

(d) Stock market volatility does in fact change through time.

3. Some funds seem to outperform simple indices, even after controlling for risk through market betas. Fund returns are also slightly predictable: past winning funds seem to do better 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, this no longer seems to be the case. Multifactor controls explain most fund persistence: funds earn persistent returns by following fairly mechanical "styles," not by skill at stock selection.

Again, these views summarize a large body of empirical work. The strength and usefulness of many results are hotly debated, and I will review some of that debate below. In addition, *why* many of these new facts are there is still somewhat of an open issue. But the old world is gone.

2 The CAPM and Multifactor Models

2.1 The CAPM

The CAPM was a theory that proved stunningly successful in empirical work. Time after time, every strategy that seemed to give high average returns turned out to also have high betas, 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.

To give some sense of that empirical work, Figure 1 presents a typical evaluation of the Capital Asset Pricing Model. 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 spread along the vertical axis shows, there is a sizeable spread in average returns between large stocks (lower average return) and small stocks (higher average return), and also 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.

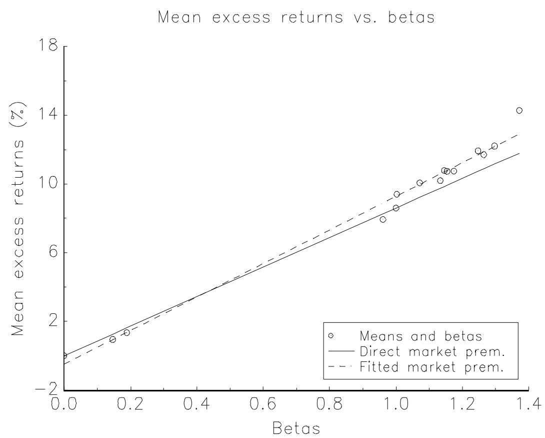

Figure 1: The CAPM. Average returns vs. betas on the NYSE value-weighted portfolio for 10 size-sorted stock portfolios, government bonds, and corporate bonds. Sample 1947-1996. The solid line draws the CAPM prediction by fitting the market proxy and treasury bill rates exactly (a time-series test). The dashed line draws the CAPM prediction by fitting an OLS cross-sectional regression to the displayed data points (the second-pass or cross-sectional test). The small firm portfolios are at the top right. The points far down and to the left are the government bond and treasury bill returns.

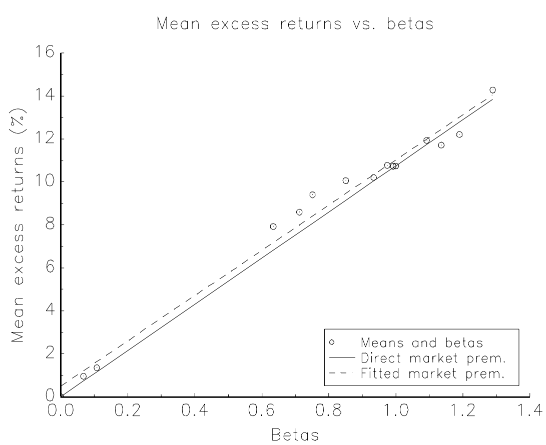

Figure 2: The CAPM, using the equally weighted NYSE as the "market portfolio." Otherwise, the figure is identical to Figure 1.

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, rather than forcing the line to go through the market and T-bill return – the dashed line – halves 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 vs. betas across the stock portfolios too shallow rather than too steep.

Box 1: The CAPM and Multifactor models.

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

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

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 via

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

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.

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

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

$$E\left(R_t^i - R_t^f\right) = \beta_{im}\lambda_m + \beta_{iA}\lambda_A + \beta_{iB}\lambda_B + ... \text{ for each } i. \tag{2}$$

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

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

2.2 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 (1971a,b) the theoretical possibility, indeed probability, that we should need *factors*, *state variables* or *sources of priced risk*, beyond movements in the market portfolio in order to explain why some average returns are higher than others.

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 the performance of his investment portfolio. While there are a few investors like that, for most of the rest 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. Very few 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 in recessions. 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 will bid up the price of that stock, or, equivalently they are willing to hold it at a lower average return. The "procyclical" stock will see its price fall, or, equivalently, 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? Most 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 like portfolios that pay off well.

[1] The 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.

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 rather than forgotten. People are unhappy if the market crashes.

2. Events such as recessions that drive investor's non-investment sources of income.

3. Variables such as the price/dividend ratio or slope of the yield curve that forecast stock or bond returns. These are called "state variables for changing investment opportunity sets."

4. Returns on other well-diversified portfolios.

The first three factors are rigorously derived 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. 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 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), you'd 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 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. These are called *factor-mimicking* portfolios.

It is vitally important that the extra 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, the average investor must be affected by it, so investors collectively bid up or down the price and expected return of assets that covary with the event rather than just transfer the risk without affecting equilibrium prices.

2.3 New Factors

The empirical search for additional, priced sources of risk has found quite a number. In general, empirical success varies inversely with theoretical purity.

2.3.1 Small and value/growth stocks.

"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 accountant's book value. 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 sensitivities 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 ratio[2] rather than size alone. As you can see, 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.

Figure 3: Average returns vs. market beta for 25 stock portfolios sorted on the basis of size and book/market ratio.

Figure 4 digs a little deeper to diagnose the problem. In the left panel, I connect portfolios that have different size within the same book/market category. As you can see, variation in size produces a variation in average returns that is positively related to variation in market betas, as we had in Figure 1. In the right panel, I connect

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

portfolios that have different *book/market* ratios within *size* categories. Variation in book/market ratio produces a variation in average return is *negatively* related to market beta. Because of this value effect, the CAPM is a disaster when confronted with these portfolios.

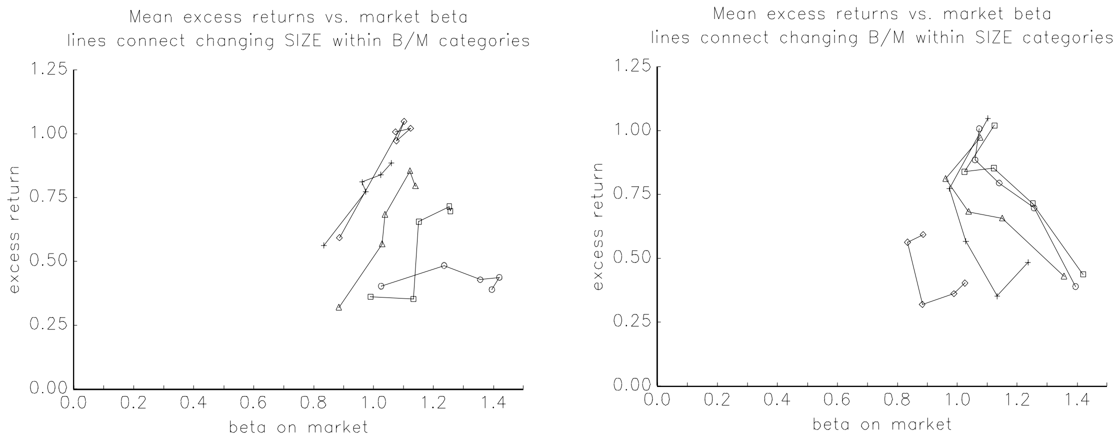

Figure 4. Average returns vs. market beta for 25 stock portfolios sorted on the basis of size and book/market ratio. The points are the same as Figure 3, with lines connecting portfolios as size varies within B/M categories (left) and as B/M varies within size categories (right).

To explain these facts, Fama and French 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° line if the model is correct. The points lie much closer to this prediction than they do in Figures 3 and 4. The worst fit is for the growth stocks (lowest line, left hand panel), for which there is little variation in average return despite large variation in size beta as one moves from small to large firms.

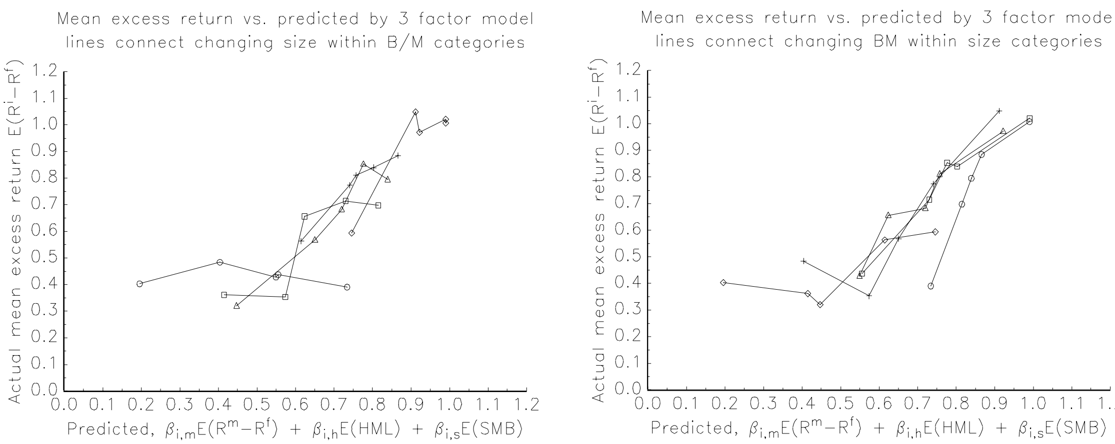

Figure 5. Average returns of 25 portfolios sorted on the basis of size and book/market, vs. predictions of the Fama-French 3 factor model. The predictions of the 3 factor model are derived by regressing each of the 25 portfolio returns R_t^i on the market portfolio R_t^m, and the two Fama French factor portfolios SMB_t (small minus big) and HML_t (high minus low book/market) (See equation (2) in Box 1).

2.3.2 What are the size and value factors?

We would like to understand the real, macroeconomic, aggregate, nondiversifiable risk that is proxied by the returns of the HML and SMB portfolios. Why 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" firm has a price that has been driven down due to financial distress. Stocks bought on the verge of bankruptcy have come back more often than not, which generates the high average returns of this strategy[3]. This observation suggests a natural interpretation of the value premium: If a credit crunch, liquidity crunch, flight to quality or similar financial event comes along, stocks in financial distress will do very badly, and this is just the sort of time at one particularly does not want to hear that one's stocks have become worthless! (One cannot count the "distress" of the individual firm as a "risk factor." Such distress

[3] The 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.

is idiosyncratic and can be diversified away. Only aggregate events that average investors care about can result in a risk premium.)

Heaton and Lucas' (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. Such an investor would therefore demand a substantial premium to hold value stocks, and would hold growth stocks despite a low premium.

Liew and Vassalou (1999) are an example of current attempts to link value and small firm returns to macroeconomic events. They find that in many countries counterparts to HML and SMB contain information above and beyond that in the market return for forecasting GDP growth. For example, they report a regression

$$GDP_{t\to t+4} = a + 0.065\ MKT_{t-4\to t} + 0.058\ HML_{t-4\to t} + \varepsilon_{t+4}$$

$GDP_{t\to t+4}$ denotes the next year's GDP growth and MKT, HML denote the previous year's return on the market index and HML portfolio. Thus, a 10% HML return raises the GDP forecast by 1/2 percentage point. Both coefficients are significant with t-statistics of 3.09 and 2.83 respectively.

The effects are still under investigation. Figure 6 plots the return on the HML and SMB portfolios, and a link between these returns and obvious macroeconomic events does not jump out at you. Both portfolios have essentially no correlation with the market return, though HML does seem to move opposite to large market declines. HML does go 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 following Ross (1976). If the returns of the 25 size and book/market portfolios could be *perfectly* replicated by the returns of the 3 factor portfolios – if the R^2 in the time-series regressions of the 25 portfolio on the three factors were 100% – 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% above the average return predicted by the Fama-French model, and its R^2 was 100%. Then, one could short a combination of the 3 factor portfolios, buy portfolio A, and earn a completely riskless profit. This logic is often used to argue that *high* R^2 should imply an *approximate* multifactor model. If the R^2 were only 95%, then an average return 5% above the factor model prediction would imply that the strategy long portfolio A and short a combination of the 3 factor portfolios would earn a very high average return with very little, though nonzero, risk – a very high Sharpe ratio.

In fact the R^2 of Fama and French's time-series regressions are all in the 90%-95% 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, and especially portfolios that do *not* have high R^2 values. Fama and French (1996) extend their analysis in this direction: They find that the SMB and HML portfolios comfortably explain strategies based on alternative price multiples (P/E, B/M), strategies based on 5 year sales growth (this is especially interesting since it is the only strategy that does not form portfolios based on price variables) and the tendency of 5 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; 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 sorts just haven't identified other sources of priced variation in stock returns. (Fama and French also find that HML and SMB do not explain "momentum," despite large R^2 values. I discuss this anomaly below.) The industry portfolios in Fama and French (1997) have lower R^2, and the model works less well.

A final worry is that the size and book/market premia seem to have diminished substantially in recent years. If this is permanent, it suggests that these opportunities were simply overlooked. The sharp decline in the SMB portfolio return right around 1980 when the small firm effect was first popularized is obvious. In the initial samples 1960-1990, the HML cumulative return starts about 1/2 below the market and ends up about 1/2 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 that of the market. (The actual numbers are 0.62 below the market and 0.77 above the market, so the HML cumulative return is $2^{0.62+0.77} = 2.6$ 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 1/2 relative to the market, meaning an investor in the market has increased his money 1 1/2 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)

On the other hand, average returns are hard to measure. There have been previous 10-20 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

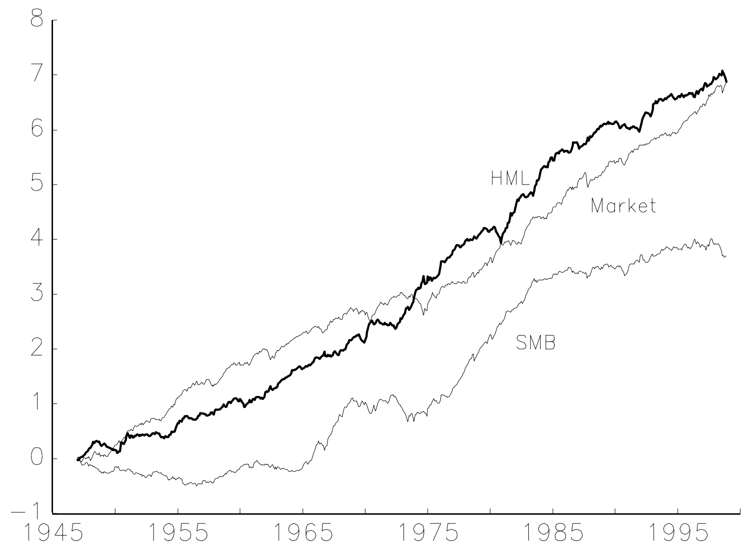

Figure 4: Cumulative returns on the market RMRF, SMB, and HML portfolios. The SMB return is formed by $R_t^{TB} + aSMB_t;\ a = \sigma(RMRF)/\sigma(SMB)$. In this way it is a return which can be cumulated rather than a zero-cost portfolio, and its standard deviation is equal to that of the market return. The market and HML are adjusted similarly. The vertical axis is the log base 2 of the cumulative return or value of one dollar invested at the beginning of the sample. Thus, each time a line increases by 1 unit, the value doubles.

some of the HML or SMB portfolio.

The argument over the status of size and book/market factors continues, but the important point is that it does so. Faced with the spectacular failure of the CAPM documented in Figures 3 and 4, one might have thought that any hope for a rational asset pricing theory was over. Now we are back where we were, examining small anomalies and arguing over refinements and interpretations of the theory. That is quite an accomplishment!

2.3.3 Macroeconomic factors

I have focused on the size and value factors since they provide the empirically most successful multifactor model, and have therefore attracted much industry as well as academic attention.

Several authors have used macroeconomic variables as factors in order to examine directly the above story that stock performance during bad macroeconomic times is indeed another determinant of 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. Cochrane (1996) looks at investment growth. All these authors find that average returns line up against betas calculated using these 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.

Though Merton's (1971a,b) theory says that variables which predict market returns should show up as factors which explain cross-sectional variation in average returns, I know of no tests that 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 dividend/price ratio at $t-1$, and they find 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, but may well help to explain why the latter work and which ones really do so.

3 Predictable returns

The view that risky asset returns are largely unpredictable, or that prices follow "random walks" was and remains immensely successful. It is simply stated, but its implications are many and subtle. (Malkiel 1990 is a classic and easily readable introduction.) It also remains widely ignored, and therefore is the source of lots of wasted trading activity.

Unpredictable returns means that if stocks went up yesterday, there is no exploitable tendency for them to decline today because of "profit taking," or for them to continue to rise today because of "momentum." "Technical factors" 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" fund managers are reportedly suffering large outflows because "value" 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, April 5 1999). Unpredictable returns means that this strategy will not do anything for one's portfolio over the long run beyond rack up trading costs. If funds are selling stocks, then "contrarian" individual investors must be buying them, but again unpredictable returns means that this strategy can't improve performance either. If you can't systematically make money, that means you can't systematically lose money either!

It was long thought that stock returns were completely unpredictable (more accurately, the excess return of stocks over short term interest rates). It now turns out that average returns on the market and individual securities *do* vary over time, and that stocks are *not* "random walks" (or even "martingales with drift") and hence returns are predictable. Much of that predictability comes at long horizons, and seems to be associated with business cycles and financial distress.

3.1 Market returns

Here is an example of market return predictability. (One of the best references for such predictability is Fama and French 1989. This example is taken from and discussed more fully in Cochrane 1997). "Low" prices – relative to dividends, book value, earnings, sales or other divisors – predict higher subsequent returns. As you can see in the R^2 values in Table 1, these are long-horizon effects: annual returns are only slightly predictable and month-to-month returns are still strikingly unpredictable, but returns at 5 year horizons seem very predictable.

Horizon k	b	$\sigma(b)$	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

Table 1. OLS regressions of excess returns (value weighted NYSE - treasury bill rate) on VW price-dividend ratio.

$$R^{VW}_{t\to t+k} - R^{TB}_{t\to t+k} = a + b(P_t/D_t) + \varepsilon_{t+k}.$$

$R_{t\to t+k}$ indicates the k year return. Standard errors in parenthesis use GMM to correct for heteroskedasticity and serial correlation. Sample 1947-1996. Source: Cochrane (1997).

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 1/3 degree per day in the springtime. 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

$$\begin{aligned} R_{t+1} - R_{t+1}^{TB} &= a + bx_t + \varepsilon_{t+1} && (3) \\ x_{t+1} &= c + \rho x_t + \delta_{t+1}. && (4) \end{aligned}$$

Small values of b and R^2 in (3) and a large coefficient ρ in (4) 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 act a bit like bonds. Any bond investor understands which pushes up prices is bad news for subsequent returns. Many stock investors see a string of good returns and become elated that we seem to be in a "bull market," inferring that future returns will be good as well. The regression reveals the opposite: a string of good returns that drives up stock prices is bad news for subsequent stock returns, as it is for bonds.

3.2 Momentum and reversal

Since a string of good returns gives a high price, it is not surprising that stocks that do well for a long time (and hence build up a high price) subsequently do poorly, and stocks that do poorly for a long time (and hence dwindle down to a low price, market value, or market to book ratio) subsequently do well. Table 2, taken from Fama and French 1996 reveals that this is in fact the case. (As usual, this table is the tip of an iceberg of research on these effects, starting with DeBont and Thaler 1985 and Jagadeesh and Titman 1993.)

Strategy	Period	Portfolio Formation Months	Average Return, 10-1 (Monthly %)
Reversal	6307-9312	60-13	-0.74
Momentum	6307-9312	12-2	+1.31
Reversal	3101-6302	60-13	-1.61
Momentum	3101-6302	12-2	+0.38

Table 2. Average monthly returns from reversal and momentum strategies. Each month, allocate all NYSE firms on CRSP 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 VI.

Reversal

The first row tracks the average monthly return from the "reversal" strategy. Each month, allocate all stocks to 10 portfolios based on performance in year -5 to year -1. Then, buy the best-performing portfolio and short the worst-performing portfolio. The first row of Table 2 shows that this strategy earns a hefty -0.74% monthly return[4]. Past long-term losers come back and past winners do badly. Fama and French verify that these portfolio returns are explained by their 3 factor model. Past winners move with value stocks, and so inherit the value stock premium.

Momentum

The second row of Table 2 tracks the average monthly return from a "momentum" strategy. Each month, allocate all stocks to 10 portfolios based on performance in the last *year*. Now, quite surprisingly, 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% monthly return.

At every moment there is a most-studied anomaly, and momentum is that anomaly right now. It is not explained by the Fama French 3 factor model. The past losers have low prices and tend to move with value stocks. Hence the model predicts 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) but is so obviously ad-hoc (i.e. an APT factor that will only explain returns of portfolios organized on the same characteristic as the factor) that nobody wants to add it.

Momentum, like the rash of other small return predictability results, may turn out not to really be there or not to be exploitable. Table 2 already presents some doubts. The third line shows 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 (1/4%) 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% or more at an annual basis. The average return of the best performing decile of a normal distribution is 1.76 standard deviations above

[4] Fama 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 market and factor portfolios that the standard deviations are roughly 1-2 times that of the market return, so that Sharpe ratios of these strategies are comparable to that of the market return in sample.

the mean[5], so the winning momentum portfolio typically went up about 80% in the previous year, and the typical losing portfolio went down about 60% per year. Only a small amount of continuation will give a 1% 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 $\sqrt{0.0025} \times 12\% = 0.6\%$. 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% 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 in the 0.01 range. Campbell, Lo and MacKinlay (1996), Table 2.4 provides an updated summary of index autocorrelations (the R^2 is the squared autocorrelation), part of which I repeat below as Table 3. Note especially the correlation of the equally weighted portfolio, which emphasizes small stocks[6].

Frequency	Portfolio	Correlation ρ_1
Daily	Value	0.18
	Equal	0.35
Monthly	Value	0.043
	Equal	0.17

Table 3. First order autocorrelation in CRSP value-weighted and equal-weighted index returns. Sample 1962-1994. Source: Campbell, Lo and MacKinlay (1996).

However, such small though statistically significant high frequency predictability, especially in small stock returns, has also since the 1960s always failed to yield exploitable profits after one accounts for transactions costs, thin trading, high short sale costs and other microstructure issues. If momentum is the same phenomenon, one may suspect that it also is not exploitable for the same reasons.

[5] We're looking for

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

where x is defined as the top 10th cutoff,

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

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

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

Momentum does require frequent trading. The portfolios in Table 2 are reformed every month. Annual winners and losers will not change that often, but the winning and losing portfolio must still be turned over at least once per year. Carhart (1996) calculates transactions costs and concludes that momentum is not exploitable after those costs are taken into account. Moskowitz and Grinblatt (1999) note that most of the apparent gains come from short positions in small, illiquid stocks. They also find that a large part of momentum profits come from the short positions taken November. Many investors sell losing stocks towards the end of December in order to establish tax losses. If one shorts small illiquid and losing stocks in November, one profits from the additional "selling pressure" in December. This is still an anomaly, but sounds a lot more like a small glitch rather than a central parable for risk and return in asset markets.

I emphasize in this as all of the changes in view that the glass is still at least half full. 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 we see. To get the 1% per month momentum return, you buy a portfolio that has typically gone up 80% in the last year, and short a portfolio that has typically gone down 60%. 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, the momentum profits may well be less than linear in the past year's returns, so mild momentum strategies may yield even less. And we have not quantified the substantial risk of momentum strategies.

3.3 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. 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 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. Equivalently, the expected rise in short term rates will imply small capital gains on the long-term bond, so investors can expect the same rate of return over one month or one year whether they hold long or short term bonds.

As with the CAPM and the view that stock returns are independent over time, the expectations model was the workhorse of empirical finance for a generation. And as with those other views, a new round of research has significantly modified the traditional view.

Maturity N	Avg. Return $E(hpr_{t+1}^{(N)})$	Std. error	Std. dev. $\sigma(hpr_{t+1}^{(N)})$
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

Table 4. Average continuously compounded one-year holding period returns on zero-coupon bonds of varying maturity. Annual data from CRSP 1953-1997. See box 2 for notation.

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 bond returns as maturity rises. The small increase in returns for long term bonds, equivalent to a slight average upward slope in the yield curve, is usually excused as a small "liquidity premium." Table 4 is again a tip of an iceberg of an illustrious career for the expectations hypothesis. Especially in times of great inflation and exchange rate instability, the expectations hypothesis does a very good first-order job.

However, one can ask a more subtle question. Perhaps there are *times* when long term bonds can be forecast to do better, and other times when short term bonds are expected to do better. If the times even out, the unconditional averages in Table 4 will show no pattern. Equivalently, we might want to check whether a forward rate that is *unusually high* forecasts an unusual *increase* in spot rates.

	Change in yields $y_{t+N}^{(1)} - y_t^{(1)} = a + b(f_t^{(N+1)} - y_t^{(1)}) + \varepsilon_{t+N}$					Holding period returns $hpr_{t+1}^{(N)} - y_t^{(1)} = a + b(f_t^{(N+1)} - y_t^{(1)}) + \varepsilon_{t+1}$				
N	a	$\sigma(a)$	b	$\sigma(b)$	$\overline{R}^2$	a	$\sigma(a)$	b	$\sigma(b)$	$\overline{R}^2$
1	0.1	0.3	-0.10	0.36	-0.02	-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	-0.3	0.5	0.77	0.31	0.11	-0.5	1.0	1.31	0.63	0.07

Table 5. Forecasts based on forward-spot spread. OLS regressions 1953-1997 annual data. Yields and returns in annual percentages. See box 3 for notation.

Table 5 gets at these issues, updating Fama and Bliss' (1986) classic regression tests. The left hand panel 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 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.

At a one-year horizon we see instead 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 4 years out, we see coefficients within one standard error of 1.0. Thus, the expectations hypothesis seems to do poorly at short (1 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 forecast one year bond returns. To check this fact, the right hand panel 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 forward-spot spread) should be able to tell you that this is a particularly good time for long bonds vs. 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 vs. bonds. As you can see, the coefficients in the right hand panel of Table 5 are all about 1.0. A high forward rate does not indicate that interest rates will be higher one year from now; it seems entirely to indicate that you will earn that much more holding long term bonds[7].

Of course, there is risk: the R^2 are all about 0.1-0.2, about the same values as the R^2 from the dividend/price 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 build with horizon as with the D/P regression, peaking in the 30% range (Fama and French 1989).

[7] The right hand panel is really not indpendent evidence, since the coefficients in the right and left hand panels of Table 5 are mechanically linked. For example 1.14 + (-0.14) = 1.0, and this holds as an accounting identity. Fama and Bliss call them "complementary regressions."

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 $y_t^{(N)} = -\frac{1}{N} p_t^{(N)}$. 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 you can contract today to borrow money $N-1$ years from now, and repay that money N years from now. Since you 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 you 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,

$$y_t^{(N)} = -\frac{1}{N} p_t^{(N)} = \frac{1}{N} \left(f_t^{(1)} + f_t^{(2)} + f_t^{(3)} + ... + f_t^{(N)} \right).$$

The expectations hypothesis states that the expected log or continuously compounded return should be the same for any bond strategy. This 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,

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

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 the difference between the expectations hypothesis and strict risk-neutrality.

3.4 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. While not a big problem for German government bonds, Russia and other 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%, US interest rates are 5%, but the Euro falls 5% 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 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 country or another, or in one currency or another[8].

As with the expectations hypothesis in the term structure, the expected depreciation view ruled for many years, and 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 (say) 5% in dollars to lend at (say) 20% in local currencies. This situation should lead one to suspect that traders expect a 15% devaluation, or a smaller chance of a larger devaluation. That is, in this case, exactly what happened. Many observers and policy analysts who ought to know better often attribute high nominal interest rates in troubled countries to "tight monetary policy" that is "strangling the economy" to "defend the currency." In fact, one's first order guess should be that such high nominal rates reflect a large probability of inflation and devaluation – loose monetary and fiscal policy – and that they correspond to much lower real rates.

Still, does a 5% interest rate differential correspond to an exactly 5% expected depreciation, or does some of it in fact still represent a high expected return from holding debt in that country's currency? Furthermore, while expected depreciation is clearly a large part of the story for high interest rates in countries that have constant high inflation or that may suffer spectacular depreciation of a pegged exchange rate, how does the story work for, say, the U.S. vs. Germany, where inflation rates diverge little, yet exchange rates fluctuate a surprisingly large amount?

Table 6 presents the facts, as summarized by Hodrick (2000) and Engel (1996).

[8] As with bonds, the expectations hypothesis is slighty 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.

The first row of Table 6 presents the average appreciation of the dollar against the indicated currency over the sample period. The dollar fell against DM, yen and Swiss Franc, but appreciated against the pound. The second row gives the average interest rate differential – the amount by which the foreign interest rate exceeds the US 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.

The second row shows roughly the right 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 shows that the difference between the first and second rows is not statistically different from zero. This fact is exactly analogous to the fact of Table 4 that the expectations hypothesis works well "on average" for US bonds and summarizes the tip of the iceberg of the long empirical success of the expectations hypothesis as applied to currencies.

As in the case of bonds, however, we can also 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 address this question, updating Fama's (1984) regression tests. The number here should be +1.0 in each case – an extra percentage point interest differential should correspond to one extra percentage point expected depreciation. As you can see, we have exactly the opposite pattern: a higher than usual interest rate abroad seems to lead, if anything to further *ap*preciation. This is the "forward discount puzzle," and takes its place alongside the forecastability of stock and bond returns. Of course it has produced a similar avalanche of academic work dissecting whether it is really there and if so, why. Engel (1996) and Lewis (1994) provide recent surveys.

The R^2 shown in Table 6 are quite low. However, like D/P, the interest differential is a slow-moving forecasting variable, so the return forecast R^2 build with horizon. Bekaert and Hodrick (1992) report that the R^2 rise to the 30-40% range at six month horizons and then decline again. Still, taking advantage of this predictability, like the bond strategies described above, is quite risky.

	DM	£	¥	SF
Mean appreciation	-1.8	3.6	-5.0	-3.0
Mean interest differential	-3.9	2.1	-3.7	-5.9
b, 1975-1989	-3.1	-2.0	-2.1	-2.6
R^2	.026	.033	.034	.033
b, 1976-1996	-0.7	-1.8	-2.4	-1.3

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

Table 6. 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 fifth rows give the coefficients and R^2 in a regression of exchange rate changes on the interest differential = forward premium,

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

where s = log spot exchange rate, f = forward rate, r^f = foreign interest rate, r^d = domestic interest rate. Source: Hodrick (1999) and Engel (1996).

I emphasize that 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.

4 Funds

One can study mutual fund returns as well as studying returns of artificially constructed stock and bond portfolios. Studying the returns of passively-managed funds that follow a specific strategy gives us a way to assess whether that strategy works in practice, after transactions costs and other realities are taken into account. Studying the returns of actively-managed funds tells us whether all the time, talent and effort put in to picking securities pays off.

My earlier depressing comments about actively-managed fund performance are also based not on doctrine but on lengthy empirical work of this sort, starting with Jensen's (1969) classic article. Studying funds is difficult because one must guard against *survivor bias*. Funds that do badly go out of business, so the average fund alive at any point in time has an artificially good track record. To address this issue, one must carefully construct databases that include all the dead funds. Carhart (1997) has done this, and I use his data.

To give some sense of the facts, Figure 7 examines Carhart's fund data. As with the stock portfolios in Figure 1, there is a definite correlation between beta and average return: funds that did well did so by taking 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 alive a few years, contributing to beta measurement error.) The average fund underperforms the line connecting treasury bills and the market index by 1.23% per year (average alpha).

Figure 7 is a bit surprising because of the wide dispersion in fund returns. Funds hold diversified portfolios, which should reduce their return variation. Yet the variation of fund returns in Figure 7 is almost as big as the variation in individual stock returns. Apparently, the vast majority of funds are *not* holding well-diversified portfolios on behalf of their clients, but rather loading up on specific bets.

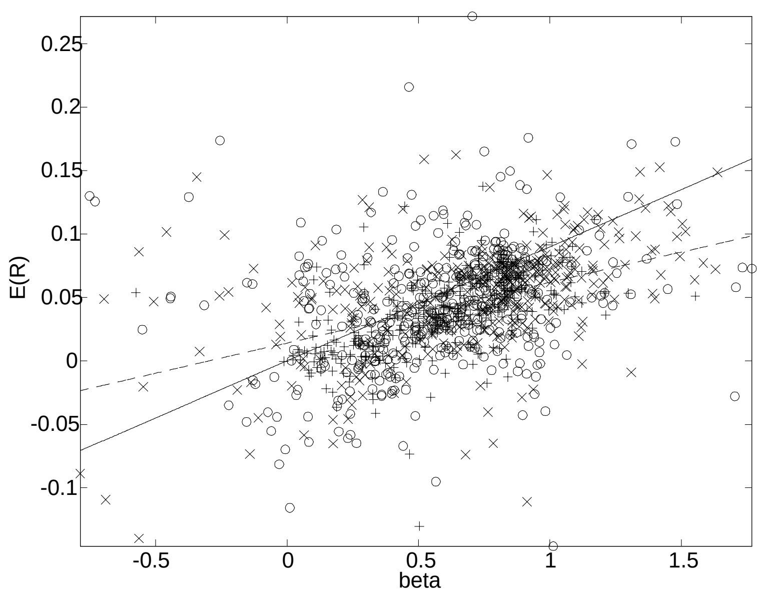

Figure 5: Average returns on mutual funds (over the treasury bill rate) vs. their market betas. The sample consists of all funds with average total net assets > \$25 million, > 25% allocation to stocks in the Carhart (1996) database. Data sample 1962-1996. The average excess return is computed as $E(R^i) - R^f = \alpha_i + \beta_i \times 9\%$. α_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. o,+ and x sort funds into thirds based on their regression coefficient 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 OLS regression of α_i on β_i. The solid line connects the treasury bill (β and excess return $=0$) and the market return ($\beta = 1$, excess return $= 9\%$).

Initially, the focus on the *average* fund seems confusing. 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. "What about fund xyz that has been doing great for all these years?" is a natural question, but we have no way of telling whether past performance of a particular fund is skill or luck.

One can be a bit more sophisticated than just looking at the average fund, of course. The most natural idea is to group funds based on past performance. Truly 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. Fund returns seemed to be random walks just like individual stock returns.

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 transactions 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 indices 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

Once we have seen the value, small firm and predictability effects, the fact 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 report that the (HML) portfolio alone gives nearly double the market Sharpe ratio – the same average return at half the standard deviation. One might well expect that funds should cluster around a line much higher than the market index.

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 "we cut our average returns in half in order to spare you exposure to non-market 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" or "asset allocation" 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%.

The strategies may be implementable if in fact very few funds were following them. That seems to be the conclusion of Figure 7. The funds are sorted by their HML beta, and the circles + and x denote funds with high, middle and low HML beta. One would hope to see the high-HML beta funds outperform the market line. But the cutoff for the top 1/3 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). Similarly, Lakonishok, Shleifer and Vishny's (1992) documentation of "value" funds' underperformance reveals that their market beta is almost precisely 1.0. It seems that "value" funds *said* they were following value, but were not really doing so, 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 by the confusing marketing) is increasing quickly. Before 1990, 14% of the funds had measured SMB betas greater than one, and 12% had HML betas greater than one. In the full sample, both numbers have *doubled* to 22% and 23%. This trend suggests that funds will, in the future, be much less well described by the market index. It is also very uncomfortable for the view that the value premium is an equilibrium risk premium, i.e. that everyone knew about the value returns but chose not to invest all along because they feared the risks of value strategies.

Persistence in fund returns.

The 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. Hendricks, Patel and Zeckhauser (1993) started the parade, and it has been the subject of a great deal of statistical analysis, especially focusing on survivor bias, since then.

Last year rank	Average Return	CAPM alpha	4-factor alpha
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%

Table 7. Portfolios of mutual funds formed on previous year's return. 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. 4 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 that 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).

Table 7, taken from Carhart (1997) is the state of the art in this question. The second column of Table 7 shows that a portfolio of the best-performing 1/30th of funds last year outperforms a portfolio of the worst-performing 1/30th of funds by 1% per month. 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 column verifies that mutual fund performance is, in fact, persistent.

Perhaps the funds that did well did so by taking 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. We have to understand fund returns with "multifactor models," if at all, just as we do for individual stock returns.

The fourth column presents alphas (intercepts, the part of average return not explained by the model) from a four factor model. The four factors include 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. Precisely, Carhart runs a regression of the fund portfolio return on the four factor portfolios, and this is the intercept or unexplained average return.

In general, one should object to the inclusion of so many factors and such ad-hoc factors. However, this is a "performance attribution" use of a multifactor model rather than a "economic explanation" use of such a model. We want to know whether fund performance, and persistence in fund performance in particular, is due to persistent stock-picking skill, or whether it results from mechanical strategies. If, for example, the persistence in fund performance is completely explained by momentum in stocks, then one can obtain the same results without hiring a professional manager and paying his fees by following a mechanical momentum strategy. Whether the momentum and other premia are really there (out of sample) or are economically justified is irrelevant to this question.

The alphas in the fourth column are almost all the same, and almost all about 1-2% per year negative. Thus, Carhart's four-factor performance attribution model explains that the persistence in fund performance is due to persistence in the underlying stocks, not due to persistent stock-picking skill. The old conclusion that actively-managed funds underperform mechanical indexing strategies by 1-2% per year is restored. There is some remaining puzzling persistence, but it is all in the large *negative* alphas of the bottom 10th to bottom 30th of performers. Somehow, the bottom funds manage to lose money year after year. Carhart also shows that the persistence of fund performance is due to momentum in the underlying stocks, rather than due to momentum funds. If a fund by good luck happened to pick stocks that go up last year, that portfolio will continue to go up a bit this year.

In sum, the advice against active management survives. However, we discover that passively managed "style" portfolios exist, and that they can earn returns that are anomalous with respect to the CAPM. Thus, an investor will have to consider how to pick among "style" funds as well as a market index.

5 Catastrophe insurance

A number of prominent funds have earned very good returns (and some, like LTCM, spectacular losses) by following "convergence trades" or similar sophisticated-sounding strategies. These strategies may also reflect high average returns as compensation for a non-market dimensions of risk. They have not been examined in the detail at which the value and small-cap strategies have been studied, so what I offer is a possible interpretation rather than a documented one, but one that is I hope plausible.

"Convergence trades" take strong positions in very similar securities that have small price differences. For example, a 29 1/2 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.) A "convergence trade" puts a strong short position on the "expensive" security and a strong long position on the "cheap" security. This is often (and by people who should know better) called an "arbitrage opportunity." However the securities are only similar, they are not identical. The spread between 29.5 and 30 year treasuries reflects the lower liquidity of 29.5 year treasuries and thus the difficulty of unloading them in rare financial panics. It is possible for this spread to widen. Nonetheless, panics are rare, and the average returns one earns in all the years when they do not happen may more than make up for the occasionally spectacular losses when they do.

Put options protect one from large price declines. The "volatility smile" in put option prices reflects the surprisingly high prices of such options, compared to the small probability of large market collapses (even when one calibrates the probability directly, rather than using the lognormal distribution of the Black-Scholes formula.) Thus, 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 well be quite good.

All of these strategies are best thought of as "catastrophe insurance" (Hsieh 1998). Most of the time they earn a small fee or "premium" in the insurance sense of the word. Once in a great while they lose a lot, and they lose a lot in times of financial panic or "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 correctly measure the small probabilities of large losses.

The difficulty in empirically examining such strategies, of course, is that rare events are rare, so estimating the true average return, including the rare events, is difficult. Many long samples will give a false sense of security because "the big one" that justifies the premium happened not to hit.

The value strategy and the profitable yield curve and foreign exchange strategies I surveyed above also smack 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 premia – the average return from holding stocks sensitive to HML or from following the bond and foreign exchange strategies – may be overstated in our data. We have had an unusually good 50 years, and rare but devastating financial panics have not happened. We have been lucky.

6 Summary and interpretation

While the list of new facts appears long, similar patters 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 knew 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 that one could do about them. A 51/49 bet is not very attractive, especially if there is any transactions cost at all. Also, the increased Sharpe ratio one can obtain by exploiting predictability is directly related to the forecast R^2, so tiny R^2, even if exploitable, did not seem like an important phenomenon. What is new is a greater understanding of how important the effects might be and of their economic interpretations.

For "price" effects, we now realize that R^2 rise with horizon when the forecasting variables are slow-moving. Hence small R^2 at high frequency can mean really substantial R^2, in the 30-50% range at longer horizons. Also the nature of these effects suggests the kinds of additional sources of price 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 yet that this amplification of the small predictability really does survive transactions costs.

6.1 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, 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 "5 year reversal" effect derives from the fact that 5 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 bonds 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 intepretatation 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.

An absent offsetting event.

In each case, an apparent difference in yield should give rise to an offsetting movement, but seems not to do so. Something *should* be predictable so that returns

[10] This 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, reduce subsequent average returns. High average returns perist, in equilibrium, when investors fear the increased risks of an asset and try to sell, lowering prices.

are *not* predictable, and it isn't.

As a specific example, Figure 8 provides a pictorial version of the results in Table 5. Suppose that the yield curve is upward sloping as in the left panel. What does this mean? If the expectations model were true, the forward rates plotted against maturity in the left hand panel would translate one-for-one to the forecast of future spot rates in the right hand panel, as plotted in the 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.

We can calculate the actual forecast of future spot rates from the estimates in Table 5, and these are given by the line market "Estimates" in Figure 8. You can see that 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 that 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 the right hand panel of Table 5 mean that a one percentage point increase in forward rate translates into a one percentage point increase in expected return. It seems that old fallacy of confusing bond *yields* with their *expected returns* also contains a grain of truth.

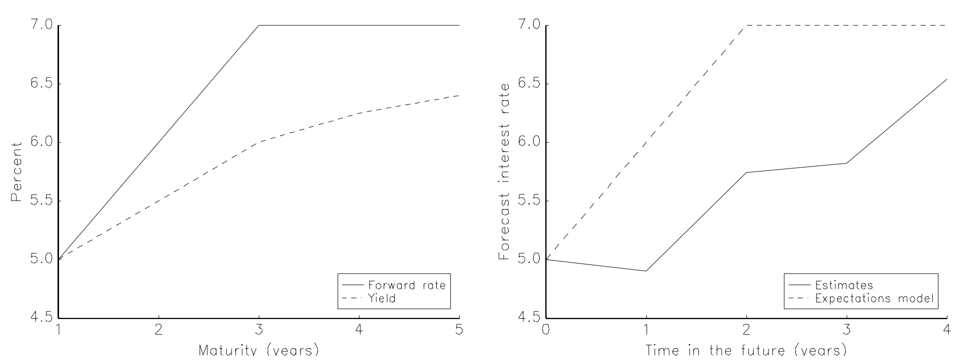

Figure 8. If the current yield curve is as plotted in the left hand panel, the right hand panel gives the forecast of future one year interest rates. The dashed line in the right hand panel gives the forecast from the expectations hypothesis, in which case forward rates today are the forecast of future spot rates. The solid line in the right hand panel is constructed from the estimates in Table 5.

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 one percentage point increase in the dividend yield translates into a one percentage point increase in return. It seems that the old fallacy of confusing increased dividend yield with increased total return does in fact 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 return 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 bottoms of recessions, and is inverted at the top of a boom. Return forecasts are high at the bottom of business cycles and low at the top of booms; "Value" and "small-cap" stocks are typically distressed. Formal quantitative and empirically successful economic models of the recession and distress premia are still in their infancy (I think Campbell and Cochrane 1999 is a good 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 value or 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 financial panic; a time when long-term bond prices and corporate bond prices are unusually low. This is a time when few people have the guts (the risk-tolerance) 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 market capitalization companies, dogs by any standards. These are companies with years of poor past returns, years of poor sales, companies on the edge of bankruptcy. 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, when 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 right now, and the stocks 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, and doing so you would have missed much of the runup in the Dow past the 6,000 point. In fact, value stocks have missed most of the recent market runup, driven by Blue-chips. This shouldn't worry you. A strategy that holds risks uncorrelated with the market must underperform the market close to 1/2 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 really *don't* really only care about long-run mean and variance; you care about doing well when the market is doing well. If you want to stay fully invested in stocks, maybe you too feel the time-varying aversion to or exposure to risk that drives the average investor to stay fully-invested right now despite low prospective returns.

This line of explanation for the foreign exchange puzzle is still a bit farther off, though there are recent attempts to make economic sense of the puzzle (See Engel's 1996 survey; Atkeson, Alvarez and Kehoe 1999 is a recent example.) At a verbal level, the strategy leads you 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 be 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, who are typically enjoying better times.

6.2 Return correlation

Momentum and its mirror in 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 noting that, with a slow-moving forecasting variable, the R^2 build over horizon. Momentum, however, is based on a fast-moving forecast variable – the last year's return. Therefore the R^2 decline rather than build 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 lovely to understand momentum as a reflection of slowly time-varying average expected returns or risk premia, 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 we need 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 that we have relied on to explain why *low* P/D, P/E, B/M, value and size forecast *higher* subsequent returns. Therefore, positive correlation of *expected* returns typically yields a negative correlation of *actual* or 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 far future (an impulse-response that starts positive but is then negative at long horizons.) Campbell Lo and MacKinaly 1996 p.264-267 give a nice quantitative exposition of these effects.

Momentum returns have also not yet been linked to business cycles or financial distress in even the informal way that I suggested for the price-based strategies. Thus, it still lacks much of 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.

6.3 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 $\sigma = 16\%$, (typical of the index), even $T = 25$ years means that one standard error is $16/5 = 3\%$ per year, and a two-standard error confidence interval runs plus or minus *six* 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 make measuring average returns statistically even harder.

In addition, viewed the right way, we really have very few data points with which to evaluate predictability. The term premium and interest rate differentials change sign only 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 US is dominated by the increase, through two recessions, to a peak in 1980 and then slow declines after that.

Many of the anomalous risk premia seem to be declining over time. We saw in Figure 6 the decline in HML and SMB premia, and the same may be true of the predictability effects as well. The last 3 years of high market returns have cut the estimated return predictability from the dividend-price ratio in *half.* This fact suggests an uncomfortable implication: that at least some of the premium the new strategies yielded in the past was due to the fact that they were simply overlooked.

Put it this way: was it really clear to the average investor in 1947 or 1963 (the beginning of the data samples) that stocks would earn 9% over bonds, and that the strategy of buying distressed small stocks would double even that return for the same level of risk? Would that investor have changed his portfolio with this knowledge, or would he have stayed pat, patiently explaining that these average returns are earned in exchange for risk that he was not prepared 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 premia measured in sample as true risk premia, 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 they are hard to measure, one is tempted to put less emphasis on the premia. 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 disquietingly suggests the "overlooked strategy" interpretation.As explained earlier, funds still cluster around the market line. It turns out that very few funds 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 premia were there all along, but such an argument is obviously strained.

7 Last words

In sum, it now seems that one can earn substantial premia for holding dimensions of risk unrelated to market movements, such as recession-related or distress related risk. One does this by following strategies such as "value" and "growth," by using the 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 premia is still a bit open to question, and the exact economic nature of the underlying risks is still debated, but we are unlikely to go back to the simple view of returns that are independent over time and the CAPM.

The next question is, *should* one change one's portfolio based on this information, and if so, to what extent? The sequel, "Portfolio Advice for a Multifactor World," addresses this question.

8 References

Atkeson, Andrew, Fernando Alvarez, and Patrick Kehoe, 1999, "Volatile Exchange Rates and the Forward Premium Anomaly: A Segmented Asset Market View," Working paper, University of Chicago.

Banz, R. W. 1981, "The Relationship Between Return and Market Value of Common Stocks," *Journal of Financial Economics*, 9, 3-18.

Bekaert, Geert and Robert J. Hodrick, 1992, "Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets," *Journal of Finance* 47, 467-509.

Berk, Jonathan, 1997, "Does Size Really Matter?" *Financial Analysts Journal*, September/October 1997, 12-18.

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,*

Campbell, John Y., Andrew W. Lo and A. Craig MacKinlay 1996, *The Econometrics of Financial Markets* Princeton NJ: Princeton University Press.

Carhart, Mark M. 1997, "On Persistence in Mutual Fund Performance," *Journal of Finance* LII 57-82.

Chen, Nai-Fu, Richard Roll and Steven Stephen A. Ross, 1986, "Economic Forces and the Stock Market," *Journal of Business* 59, 383-403.

Cochrane, John H., 1997, "Where is the Market Going? Uncertain Facts and Novel Theories," *Economic Perspectives* XXI: 6 (November/December 1997) Federal Reserve Bank of Chicago.

Cochrane, John H., 1996, "A Cross-Sectional Test of an Investment-Based Asset Pricing Model," *Journal of Political Economy* 104, 572-621.

DeBondt, Werner F.M. and Thaler, Richard H., 1985, "Does the Stock Market Overract?" *Journal of Finance* 40, 793-805.

Engel, Charles, 1996, "The Forward Discount Anomaly and the Risk Premium: a Survey of Recent Evidence," *Journal of Empirical Finance* 3, 123-192.

Fama, Eugene F., 1965, "The Behavior of Stock Market Prices, *Journal of Business* 38, 34-105.

Fama, Eugene F., 1970, "Efficient Capital Markets: A review of Theory and Empirical Work," *Journal of Finance* 25, 383-417.

Fama, Eugene F., 1984, "Forward and Spot Exchange Rates," *Journal of Monetary Economics* 14, 319-338.

Fama, Eugene F. 1991, "Efficient Markets II," *Journal of Finance* 46, 1575-1618.

Fama, Eugene F. and Robert R. Bliss, 1987, "The information in Long-Maturity Forward Rates," *American Economic Review*, 77, 680-92.

Fama, Eugene F. and Kenneth R. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics* 25, 23-49.

Fama, Eugene F. and Kenneth R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics* 33, 3-56.

Fama, Eugene F.,and Kenneth R. French, 1996, "Multifactor Explanations of Asset-Pricing Anomalies," *Journal of Finance* 47, 426-465.

Fama, Eugene F., Kenneth R. French, 1997, "Size and Book-to-Market Factors in Earnings and Returns,"*Journal of Finance* 50, 131-55.

Fama, Eugene F. and Kenneth R. French, 1997, "Industry Costs of Equity," *Journal of Financial Economics* 43, 153-193.

Heaton, John and Deborah Lucas, 1997 "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," Manuscript, Northwestern University.

Hendricks, Darryll, Jayendu Patel and Richard Zeckhauser, 1993, "Hot Hands in Mutual Funds: Short-Term Persistence of Performance," *Journal of Finance* 48, 93-130.

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

Hsieh, David and William Fung, 1999, "Hedge Fund Risk Management," Working paper, Duke University.

Jagannathan, Ravi and Zhenyu Wang, 1996 "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance* 51, 3-53.

Jagadeesh, Narasimham, and Sheridan Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance* 48, 65-91.

Jensen, Michael C., 1969, "The pricing of capital assets and evaluation of investment portfolios," *Journal of Business* 42, 167-247.

Lakonishok, Josef, Andrei Shleifer and Robert W. Vishny, 1992, "The Structure and Perfomrance of the Money Managament Industry," *Brookings Papers on Economic Activity: Microeconomics 1992*, 339-391

Lewis, Karen, K, 1995, in G.Grossamn and K. Rogoff, eds., "Puzzles in International Financial Markets," *Handbook of International Economics, Volume III*, Elsevier Science B.V, 1913-1971.

Liew, Jimmy, and Maria Vassalou, 1999, "Can Book-to-Market, Size and Momentum be Risk Factors that Predict Economic Growth?" Working paper, Columbia University.

Daniel, Kent, David Hirshleifer and Avanidhar Subrahmanyam, 1998, "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance* LIII, 1839-1885.

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

Markowitz, H., 1952, "Portfolio Selection," *Journal of Finance*, 7, 77-99.

MacKinlay, A. Craig 1995, "Multifactor Models Do Not Explain Deviations from the CAPM,"*Journal of Financial Economics* 38, 3-28.

Merton, Robert C., 1969 "Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case," *Review of Economics and Statistics* 51, 247-257.

Merton, Robert C., 1971a, "Optimum consumption and Portfolio rules in a Continuous Time Model," *Journal of Economic Theory* III 373-413.

Merton, Robert C., 1971b, "An Intertemporal Capital Asset Pricing Model," *Econometrica* 41, 867-87.

Moskowitz, Tobias and Mark Grinblatt, 1998, "Do industries explain momentum?" CRSP working paper 480, University of Chicago.

Moskowitz, Tobias and Mark Grinblatt, 1999, "Tax Loss Selling and Return Autocorrelation: New Evidence," Working Paper, University of Chicago.

New York Times, Mutual Funds Report; What's Killing the Value Managers? April 4, 1999, Sunday.

Reyfman, Alexander, 1997, "Labor Market Risk and Expected Asset Returns," Ph.D. Thesis, University of Chicago.

Ross, S. A., 1976, "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic theory* 13, 341-360.

Samuelson, Paul A., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics* 51, 239-246.

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

Lewis, Michael, 1999, "How the Eggheads Cracked," *New York Times Magazine* January 24.