

Portfolio advice for a multifactor world

John H. Cochrane

Introduction and summary

A companion article in this issue, “New facts in finance,” summarizes the revolution in how financial economists view the world. Briefly, there are strategies that result in high average returns without large *betas*, that is, with no strong tendency for the strategy’s returns to move up and down with the market as a whole. Multifactor models have supplanted the capital asset pricing model (CAPM) in describing these phenomena. Stock and bond returns, once thought to be independent over time, turn out to be predictable at long horizons. All of these phenomena seem to reflect a premium for holding macroeconomic risks associated with the business cycle and for holding assets that do poorly in times of financial distress. They also all reflect the information in prices—high prices lead to low returns and low prices lead to high returns.

The world of investment opportunities has also changed. Where once an investor faced a fairly straightforward choice between managed mutual funds, index funds, and relatively expensive trading on his own account, now he must choose among a bewildering variety of fund *styles* (such as value, growth, small cap, balanced, income, global, emerging market, and convergence), as well as more complex claims of active fund managers with customized styles and strategies, and electronic trading via the Internet. (Msn.com’s latest advertisement suggests that one should sign up in order to “check the hour’s hottest stocks.” Does a beleaguered investor really have to do that to earn a reasonable return?) The advertisements of investment advisory services make it seem important to tailor an investment portfolio from this bewildering set of choices to the particular circumstances, goals, and desires of each investor.

What should an investor do? An important current of academic research investigates how portfolio theory should adapt to our new view of the financial

world. In this article, I summarize this research and I distill its advice for investors. In particular, which of the bewildering new investment styles seem most promising? Should you attempt to time stock, bond, or foreign exchange markets, and if so how much? To what extent and how should an investment portfolio be tailored to your specific circumstances? Finally, what can we say about the future investment environment? What kind of products will be attractive to investors in the future, and how should public policy react to these financial innovations?

I start by reviewing the traditional academic portfolio advice, which follows from the traditional view that the CAPM is roughly correct and that returns are not predictable over time. In that view, all investors (who do not have special information) should split their money between risk-free bonds and a broad-based passively managed index fund that approximates the “market portfolio.” More risk-tolerant investors put more money into the stock fund, more risk averse investors put more money into the bond fund, and that is it.

The new academic portfolio advice reacts to the new facts. An investor should hold, in addition to the market portfolio and risk-free bonds, a number of passively managed “style” funds that capture the broad (nondiversifiable) risks common to large numbers

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

of investors. In addition to the overall level of risk aversion, his exposure to or aversion to the various additional risk factors matters as well. For example, an investor who owns a small steel company should shade his investments away from a steel industry portfolio, or cyclical stocks in general; a wealthy investor with no other business or labor income can afford to take on the “value” and other stocks that seem to offer a premium in return for potentially poor performance in times of financial distress. The stock market is a way of transferring risks; those exposed to risks can hedge them by proper investments, and those who are not exposed to risks can earn a premium by taking on risks that others do not wish to shoulder.

Since returns are somewhat predictable, investors can enhance their average returns by moving their assets around among broad categories of investments. However, the market timing signals are slow-moving, and I show that the uncertainty about the nature and strength of market timing effects dramatically reduces the optimal amount by which investors can profit from them.

I emphasize a cautionary fact: *The average investor must hold the market.* You should only vary from a passive market index if you are different from everyone else. It cannot be the case that every investor should tilt his portfolio toward “value” or other high-yield strategies. If everybody did it, the phenomenon would disappear. Thus, if such strategies will persist at all, it must be the case that for every investor who should take advantage of them, there is another investor who should take an unusually *low* position, sacrificing the good average returns for a reduction in risk. It cannot be the case that every investor should “market time,” buying when prices are low and selling when prices are high. If everyone did it, that phenomenon would also disappear. The phenomenon can only persist if, for every investor who should enhance returns by such market-timing, there is another investor who is so exposed to or averse to the time-varying risks that cause return predictability, that he *should* “buy high and sell low,” again earning a lower average return in exchange for avoiding risks.

We have only scratched the surface of asset markets’ usefulness for sharing risks. As often in economics, what appears from the outside to be greedy behavior is in fact socially useful.

The traditional view

The new portfolio theory really extends rather than overturns the traditional academic portfolio theory. Thus, it’s useful to start by reminding ourselves what the traditional portfolio theory is and why. The traditional academic portfolio theory, starting from

Markowitz (1952) and expounded in every finance textbook, remains one of the most useful and enduring bits of economics developed in the last 50 years.

Two-fund theorem

The traditional advice is to split your investments between a money-market fund and a broad-based, passively managed stock fund. That fund should concentrate on minimizing fees and transaction costs, period. It should avoid the temptation to actively manage its portfolio, trying to chase the latest hot stock or trying to foresee market movements. An index fund or other approximation to the *market portfolio* that passively holds a bit of every stock is ideal.

Figure 1 summarizes the analysis behind this advice. The straight line gives the *mean-variance frontier*—the portfolios that give the highest mean return for every level of volatility. Every investor should pick a portfolio on the mean-variance frontier. The upward-curved lines are *indifference curves* that represent investors’ preference for more mean return and less volatility. The indifference curve to the lower left represents a risk-averse investor, who chooses a portfolio with less mean return but also less volatility; the indifference curve to the upper right represents a more risk-tolerant investor who chooses a portfolio with more mean return but also higher risk.

This seems like a lot of person-specific portfolio formation. However, every portfolio on the mean-variance frontier can be formed as a combination of a risk-free money-market fund and the *market portfolio* of all risky assets. Therefore, every investor need only hold different proportions of these *two funds*.

Bad portfolio advice

The portfolio advice is not so remarkable for what it does say, which given the setup is fairly straightforward, as for what it does not say. Compared with common sense and much industry practice, it is radical advice.

One might have thought that investors willing to take on a little more risk in exchange for the promise of better returns should weight their portfolios to riskier stocks, or to value, growth, small-cap, or other riskier fund styles. Conversely, one might have thought that investors who are willing to forego some return for more safety should weight their portfolios to safer stocks, or to blue-chip, income, or other safer fund styles. Certainly, some professional advice in deciding which style is suited for an investor’s risk tolerance, if not a portfolio professionally tailored to each investor’s circumstances, seems only sensible and prudent. The advertisements that promise “we build the portfolio that’s right for *you*” cater to this natural and sensible-sounding idea.

Figure 1 proves that nothing of the sort is true. All stock portfolios lie on or inside the curved risky asset frontier. Hence, an investor who wants more return and is willing to take more risk than the market portfolio will do better by borrowing to invest in the market—including the large-cap, income, and otherwise safe stocks—than by holding a portfolio of riskier stocks. An investor who wants something less risky than the market portfolio will do better by splitting his investment between the market and a money-market fund than by holding only safe stocks, even though his stock portfolio will then contain some of the small-cap, value, or otherwise risky stocks. Everyone holds the same market portfolio; the only decision is how much of it to hold.

The two-fund theorem in principle still allows for a good deal of customized portfolio formation and active management if investors or managers have different *information* or *beliefs*. For example, if an investor knows that small-cap stocks are ready for a rebound, then the optimal (or tangency) portfolio that reflects this knowledge will be more heavily weighted toward small-cap stocks than the market portfolio held by the average investor. All the analysis of figure 1 holds, but this specially constructed *tangency portfolio* goes in the place indicated by the market portfolio in the figure. However, the empirical success of market efficiency, and the poor performance of professional managers relative to passive indexation, strongly suggests that these attempts will not pay off for most investors. For this reason, the standard advice is to hold passively managed funds that concentrate on minimizing transaction costs and fees, rather than a carefully constructed tangency portfolio that reflects an investor's or manager's special insights. However, a quantitative portfolio management industry tries hard to mix information or beliefs about the behavior

of different securities with the theory of figure 1 (for example, see Black and Litterman, 1991).


The two-fund theorem leaves open the possibility that the investor's *horizon* matters as well as his risk aversion. What could be more natural than the often repeated advice that a long-term investor can afford to ride out all the market's short-term volatility, while a short-term investor should avoid stocks because he may have to sell at the bottom rather than wait for the inevitable recovery after a price drop? The fallacy lies in the "inevitable" recovery. If returns are close to independent over time (like a coin flip), and prices are close to a random walk, a price drop makes it no more likely that prices will rise more in the future. This view implies that stocks are *not* safer in the long run, and the stock/bond allocation should be independent of investment horizon.

This proposition can be shown to be precisely true in several popular mathematical models of the portfolio decision. If returns are independent over time, then the mean and variance of continuously compounded returns rises in proportion to the horizon: The mean and variance of ten-year returns are ten times those of one-year returns, so the ratio of mean to variance is the same at all horizons. More elegantly, Merton (1969) and Samuelson (1969) show that an investor with a constant relative risk aversion utility who can continually rebalance his portfolio between stocks and bonds will always choose the same stock/bond proportion regardless of investment horizon, when returns are independent over time.

Taking the advice

This advice has had a sizable impact on portfolio practice. Before this advice was widely popularized in the early 1970s, the proposition that professional active management and stock selection could outperform blindly holding an index seemed self-evident, and passively managed index funds were practically unknown. They have exploded in size since then. The remaining actively managed funds clearly feel the need to defend active management in the face of the advice to hold passive index funds and the fact that active managers selected on any ex-ante basis underperform indexes ex-post.

The one input to the optimal portfolio advice is risk tolerance, and many providers of investment services have started thinking about how to measure risk tolerance using a series of questionnaires. This is the trickiest part of the conventional advice, in part since conventional



measures of risk tolerance often seem quite out of whack with risk aversion displayed in asset markets. (This is the *equity premium puzzle*; see Cochrane, 1997, for a review.) However, the basic question is whether one is more risk tolerant or less risk tolerant than the average investor. This question is fairly easy to conceptualize and can lead to a solid qualitative, if not quantitative, answer.

One might object to the logical inconsistency of providing portfolio advice based on a view of the world in which everyone is already following such advice. (This logic is what allowed me to identify the mean-variance frontier with the market portfolio.) However, this logic is only wrong if other investors are *systematically* wrong. If some investors hold too much of a certain stock, but others hold too little of it, market valuations are unaffected and the advice to hold the market portfolio is still valid.

New portfolio theory

Multiple factors: An *N-fund* theorem


Figure 2 shows how the simple two-fund theorem of figure 1 changes if there are multiple sources of priced risk. This section is a graphical version of Fama's (1996) analysis. Much of the theory comes from Merton (1969, 1971, 1973).

To keep the figure simple I consider just one additional factor. For concreteness, think of an additional recession factor. Now, investors care about three attributes of their portfolios: 1) They want higher

average returns; 2) they want lower standard deviations or overall risk; and 3) they are willing to accept a portfolio with a little lower mean return or a little higher standard deviation of return if the portfolio does not do poorly in recessions. In the context of figure 2, this means that investors are happier with portfolios that are higher up (more mean), more to the left (less standard deviation), and farther out (lower recession sensitivity). The indifference *curves* of figure 1 become indifference *surfaces*. Panel A of figure 2 shows one such surface curving upwards.

As with figure 1, we must next think about what is available. We can now calculate a frontier of portfolios based on their mean, variance, and recession sensitivity. This frontier is the *multifactor efficient frontier*. A typical investor then picks a point as shown in panel A of figure 2, which gives him the best possible portfolio—trading off mean, variance, and recession sensitivity—that is available. Investors want to hold multifactor efficient, rather than mean-variance efficient, portfolios. As the mean-variance frontier of figure 1 is a hyperbola, this frontier is a revolution of a hyperbola. The appendix summarizes the mathematics behind this figure.

Panel B of figure 2 adds a risk-free rate. As the mean-variance frontier of figure 1 was the minimal V shape emanating from the risk-free rate (R^f) that includes the hyperbolic risky frontier, now the multifactor efficient frontier is the minimum *cone* that includes the hyperbolic risky multifactor efficient frontier, as shown.



As every point of the mean-variance frontier of figure 1 can be reached by some combination of two funds—a risk-free rate and the market portfolio—now every point on the multifactor efficient frontier can be reached by some combination of *three* multifactor efficient funds. The most convenient set of portfolios is the risk-free rate (money-market security), the market portfolio (the risky portfolio held by the average investor), and one additional multifactor efficient portfolio on the tangency region as shown in panel B of figure 2. It is convenient to take this third portfolio to be a zero-cost, zero-beta portfolio, so that it isolates the extra dimension of risk.

Investors now may differ in their desire or ability to take on recession-related risk as well as in their tolerance for overall risk. Thus, some will want portfolios that are farther in and out, while others will want portfolios that are farther to the left and right. They can achieve these varied portfolios by different weights in the *three* multifactor efficient portfolios, or *three funds*.

Implications for mean-variance investors

The mean-variance frontier still exists—it is the projection of the cone shown in figure 2 on the mean-variance plane. As the figure shows, the average investor is willing to give up some mean or accept more variance in order to reduce the recession-sensitivity of his portfolio. The average investor must hold the market portfolio, so *the market return is no longer on the mean-variance frontier*.

Suppose, however, that *you* are concerned only with mean and variance—you are not exposed to the recession risk, or the risks associated with any other factor, and you only want to get the best possible mean return for given standard deviation. If so, you still want to solve the mean-variance problem of figure 1, and you still want a mean-variance efficient portfolio. The important implication of a multifactor world is that you, the mean-variance investor, should no longer hold the *market* portfolio.

You can still achieve a mean-variance efficient portfolio just as in figure 1 by a combination of a money market fund and a single *tangency portfolio*, lying on the upper portion of the curved risky-asset frontier. The tangency portfolio now takes stronger positions than the market portfolio in factors such as value or recession-sensitive stocks that the *average* investor fears.

Predictable returns

The fact that returns are in fact somewhat predictable modifies the standard portfolio advice in three ways. It introduces horizon effects, it allows market-timing strategies, and it introduces multiple factors

via hedging demands (if expected returns vary over time, investors may want to hold assets that protect them against this risk).

Horizon effects

Recall that when stock returns are independent over time (like coin flips), the allocation between stocks and bonds does not depend at all on the investment horizon, since mean returns (reward) and the variance of returns (risk) both increase in proportion to the investment horizon. But if returns are predictable, the mean and variance may no longer scale the same way with horizon. If a high return today implies a high return tomorrow—positive serial correlation—then the variance of returns will increase with horizon faster than does the mean return. In this case, stocks are worse in the long run. If a high return today implies a lower return tomorrow—negative serial correlation or *mean reversion*—then the variance of long-horizon returns is lower than the variance of one-period returns times the horizon. In this case, stocks are more attractive for the long run.¹ For example, if the second coin flip is always the opposite of the first coin flip, then two coin flips are much less risky than they would be if each flip were independent, and a “long-run coin flipper” is more likely to take the bet.

Which case is true? Overall, the evidence suggests that stock prices do tend to come back slowly and partially after a shock, so return variances at horizons of five years and longer are about one-half to two-thirds as large as short-horizon variances suggest. Direct measures of the serial correlation of stock returns, or equivalent direct measures of the mean and variance of long-horizon returns, depend a lot on the period studied and the econometric method. Multivariate methods give somewhat stronger evidence. Intuitively, the price/dividend (p/d) ratio does not explode. Hence, the long-run variance of prices must be the same as the long-run variance of dividends, and this extra piece of information helps to measure the long-run variance of returns. (Cochrane and Sbordone, 1986, and Cochrane, 1994, use this idea. See Campbell, Lo, and MacKinlay, 1997, for a summary of these issues and of the extensive literature.)

How big are the horizon effects? Barberis (1999) calculates optimal portfolios for different horizons when returns are predictable. Figure 3 presents some of his results.

We start with a very simple setup: The investor allocates his portfolio between stocks and bonds and then holds it without rebalancing for the indicated horizon. His objective is to maximize the expected utility of wealth at the indicated horizon. The flat line in figure 3 shows the standard result: If returns are not

predictable, then the allocation to stocks does not depend on horizon.

The top (black) line in figure 3 adds the effects of return predictability on the investment calculation. The optimal allocation to stocks increases sharply with horizon, from about 40 percent for a monthly horizon to 100 percent for a ten-year horizon. To quantify the effects of predictability, Barberis uses a simple model,

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

$$2) \quad x_{t+1} = c + \rho x_t + \delta_{t+1},$$

using the d/p ratio for the forecasting variable x . (Whether or not one includes returns in the right hand side makes little difference.) Barberis estimates significant mean-reversion: In Barberis's regressions, the implied standard deviation of ten-year returns is 23.7 percent, just more than half of the 45.2 percent value implied by the standard deviation of monthly returns. Stocks are indeed safer in the long run, and the greater allocation to stocks shown in figure 3 for a long-run investor reflects this fact.


Uncertainty about predictability

This calculation ignores the fact that we do not know how predictable returns really are. One could address this fact by calculating standard errors for portfolio computations; and such standard errors do indicate substantial uncertainty. However, standard error uncertainty is symmetric—returns might be more predictable than we think or they might be less predictable. This measure of uncertainty would say that we are just as likely to want an even greater long-run stock allocation as we are to shade the advice back to a constant allocation.

Intuitively, however, uncertainty about predictability should lead us to shade the advice back toward the standard advice. Standard errors do not capture the uncertainties behind this (good) intuition for at least two reasons.

First, the predictability captured in Barberis's regression of returns on dividend/price ratios certainly results to some extent from data-dredging. Thousands of series were examined by many authors, and we have settled on the one or two that seem to predict returns best in sample. The predictability will obviously be worse out of sample, and good portfolio advice should account for this bias. Standard errors take the set of forecasting variables and the functional form as given.

Second, the portfolio calculation assumes that the investor knows the return forecasting process perfectly. Standard errors only reflect the fact that we do not know the return forecasting process, so we



are unsure about what investors want to do.² What we would like to do is to solve a portfolio problem in which investors treat uncertainty about the forecastability of returns as part of the risk that they face, along with the risks represented by the error terms of the statistical model. Kandel and Stambaugh (1996) and Barberis (1999) tackle this important problem.

Figure 3 also gives Barberis's calculations of the effects of parameter uncertainty on the stock/bond allocation problem. (See box 1 for a description of how these calculations are made.) The lowest (dotted) line considers a simple case. The investor knows, correctly, that returns are independent over time (not predictable) but he is not sure about the mean return. Without parameter uncertainty, this situation gives rise to the constant stock allocation—the flat line. Adding parameter uncertainty *lowers* the allocation to stocks for long horizons; it declines from 34 percent to about 28 percent at a ten-year horizon.

The reason is simple. If the investor sees a few good years of returns after making the investment, this raises his estimate of the actual mean return and, thus, his estimate of the returns over the rest of the investment period. Conversely, a few bad years will lower his estimate of the mean return for remaining years. Thus, learning about parameters induces a

BOX 1

How to include model uncertainty in portfolio problems

A statistical model, such as equations 1 and 2, tells us the distribution of future returns once we know the parameters θ , $f(R_{t+1}|\theta, x_1, x_2, \dots, x_t)$, where x_t denotes all the data used (returns, d/p, etc.).

We would like to evaluate uncertainty by the distribution of returns conditional only on the history available to make guesses about the future, $f(R_{t+1}|x_1, x_2, \dots, x_t)$. We can use Bayesian analysis to evaluate this concept. If we can summarize the information about parameters given the historical data as $f(\theta|x_1, x_2, \dots, x_t)$, then we can find the distribution we want by

$$f(R_{t+1}|x_1, x_2, \dots, x_t) = \int f(R_{t+1}|\theta) f(\theta|x_1, x_2, \dots, x_t) d\theta.$$

In turn, we can construct $f(\theta|x_1, x_2, \dots, x_t)$, from a prior $f(\theta)$ and the likelihood function

$f(x_1, x_2, \dots, x_t|\theta)$ via the standard law for conditional probabilities,

$$f(\theta|x_1, x_2, \dots, x_t) = \frac{f(x_1, x_2, \dots, x_t|\theta) f(\theta)}{f(x_1, x_2, \dots, x_t)}$$

$$f(x_1, x_2, \dots, x_t) = \int f(x_1, x_2, \dots, x_t|\theta) f(\theta) d\theta.$$

Barberis (1999), Kandel and Stambaugh (1996), Brennan, Schwartz, and Lagnado (1997) use these rules to compute $f(R_{t+1}|x_1, x_2, \dots, x_t)$, and solve portfolio problems with this distribution over future returns.

positive correlation between early returns and later returns. Positive correlation makes long-horizon returns more than proportionally risky and reduces the optimal allocation to stocks.

The colored line in figure 3 shows the effects of parameter uncertainty on the investment problem, when we allow return predictability as well. As the figure shows, uncertainty about predictable returns cuts the increase in stock allocation from one to ten years *in half*. In addition to the positive correlation of returns due to learning about their mean mentioned above, uncertainty about the true amount of predictability adds to the risk (including parameter risk) of longer horizon returns.

Market timing

Market-timing strategies are the most obvious implication of return predictability. If there are times when expected returns are high and other times when they are low, you might well want to hold more stocks when expected returns are high, and fewer when expected returns are low. Of course, the crucial question is, how *much* market-timing should you engage in? Several authors have recently addressed this technically challenging question.

Much of the difficulty with return predictability (as with other dynamic portfolio questions) lies in computing the optimal strategy—exactly how should you adjust your portfolio as the return prediction signals change? Gallant, Hansen, and Tauchen (1990) show how to measure the potential benefits of market-timing without actually calculating the market-timing strategy.

The mean–standard deviation tradeoff or *Sharpe ratio*—the slope of the frontier graphed in figure 1—is a convenient summary of any strategy. If the risk-free rate is constant and known, *the square of the maximum unconditional Sharpe ratio is the average of the squared conditional Sharpe ratios*. (The appendix details the calculation.) Since we take an average of *squared* conditional Sharpe ratios, volatility in conditional Sharpe ratios—time-variation in expected returns or return volatility—is good for an investor who cares about the unconditional Sharpe ratio. By moving into stocks in times of high Sharpe ratio and moving out of the market in times of low Sharpe ratio, the investor does better than he would by buying and holding. Furthermore, *the best unconditional Sharpe ratio is directly related to the R² in the return-forecasting regression*.

The buy-and-hold Sharpe ratio has been about 0.5 on an annual basis in U.S. data—stocks have earned an average return of about 8 percent over Treasury bills, with a standard deviation of about 16 percent. Table 1 presents a calculation of the increased Sharpe ratio one should be able to achieve by market-timing, based on regressions of returns on dividend/price ratios. (I use the regression estimates from table 1 of “New facts in finance.”)

As table 1 indicates, market-timing should be a great benefit. Holding constant the portfolio volatility, market-timing should raise average returns by about two-fifths at an annual horizon and it should almost double average returns at a five-year horizon.

TABLE 1		
Maximum unconditional Sharpe ratios		
Horizon k (years)	R^2	Annualized Sharpe ratio
Buy & hold		0.50
1	0.17	0.71
2	0.26	0.72
3	0.38	0.78
5	0.59	0.95

Notes: Maximum unconditional Sharpe ratios available from market-timing based on regressions of value-weighted NYSE index returns on the dividend/price ratio. The table reports annualized Sharpe ratios corresponding to each R^2 .

The formula is $\frac{S^*}{\sqrt{k}} = \sqrt{0.5^2 + \frac{R^2}{k}} / \sqrt{1-R^2}$ and is derived in the appendix.

**Optimal market-timing:
An Euler equation approach**

Brandt (1999) presents a clever way to estimate a market-timing portfolio rule without solving a model. Where standard asset pricing models fix the consumption or wealth process and estimate preference parameters, Brandt fixes the preference parameters (as one does in a portfolio question) and estimates the portfolio decision, that is, he estimates the optimal consumption or wealth process.³ This calculation is very clever because it does not require one to specify a statistical model for the stock returns (like equations 1–2), and it does not require one to solve the economic model.

Figure 4 presents one of Brandt’s results. The figure shows the optimal allocation to stocks as a function of investment horizon and of the dividend/price ratio, which forecasts returns. There is a mild horizon effect, about in line with Barberis’s results of figure 3 without parameter uncertainty: Longer term investors hold more stocks. There is also a strong market-timing effect. The fraction of wealth invested in stocks varies by about 200 percentage points for all investors. For example, long-term investors vary from about 75 percent to 225 percent of wealth invested in stocks as the d/p ratio rises from 2.8 percent to 5.5 percent.


Optimal market-timing: A solution


Campbell and Vicera (1999) actually calculate a solution to the optimal market-timing question. They model investors

who desire lifetime consumption⁴ rather than portfolio returns at a fixed horizon. They model the time-variation in expected stock returns via equation 1 on d/p ratios. Their investors live only off invested wealth and have no labor income or labor income risk. Thus, these investors are poised to take advantage of business cycle related variation in expected returns.

As one might expect, the optimal investment strategy takes strong advantage of market-timing possibilities. Figure 5 reproduces Campbell and Vicera’s optimal allocation to stocks as a function of the expected return, forecast from d/p ratios via equation 1. A risk aversion coefficient of 4 implies that investors roughly want to be fully invested in stocks at the average expected excess return of 6 percent, so this is a sensible risk aversion value to consider. Then, as the d/p ratio ranges from minus two to plus one standard deviations from its mean, these investors range from –50 percent in stocks to 220 percent in stocks. This is aggressive market-timing indeed.


Figure 6 presents the calculation in a different way: It gives the optimal allocation to stocks over time, based on dividend/price ratio variation over time. The high d/p ratios of the 1950s suggest a strong stock position, and that strong position profits from the high returns of the late 1950s to early 1960s. The low d/p ratios of the 1960s suggest a much smaller position in stocks, and this smaller position avoids the bad returns of the 1970s. The high d/p ratios of the 1970s suggest strong stock positions again, which benefit from the good return of the 1980s; current





unprecedented high prices suggest the lowest stock positions ever. The optimal allocation to stocks again varies wildly, from 0 (now) to over 300 percent.

Campbell and Viceria's calculations are, if anything, conservative compared with others in this literature. Other calculations, using other utility functions, solution techniques, and calibrations of the forecasting process often produce even more aggressive market-timing strategies. For example, Brennan, Schwartz, and Lagnado (1997) make a similar calculation with two additional forecasting variables. They report market-timing strategies that essentially jump back and forth between constraints at 0 percent in stocks and 100 percent in stocks.



Campbell and Viceria also present achieved utility calculations that mirror the lesson of table 1: Failing to time the market seems to impose a large cost.

Doubts

One may be understandably reluctant to take on quite such strong market-timing positions as indicated by figures 5 and 6, or to believe table 1 that market timing can nearly double five-year Sharpe ratios. In particular, one might question advice that would have meant missing the dramatic runup in stock values of the late 1990s. Rather than a failure of nerve, perhaps such reluctance reveals that the calculations do not yet include important considerations and, therefore, overstate the desirable amount of market-timing and its benefits.

First, the unconditional Sharpe ratio as reported in table 1 for, say, five-year horizons answers the question, "Over very long periods, if an investor follows the best possible market-timing strategy and evaluates his portfolio based on five-year returns, what Sharpe ratio does he achieve?" It does *not* answer the question, "Given today's d/p, what is the best Sharpe ratio you can achieve for the next five years by following market-timing signals?" The latter question characterizes the return distribution conditional on today's d/p. It is harder to evaluate; it depends on the initial d/p level, and it is lower, especially for a slow-moving signal such as d/p.

To see the point, suppose that the d/p ratio is determined on day one, is constant thereafter, and indicates high or low returns in perpetuity. *Conditional* on the d/p ratio, one cannot time the market at all. But since the investor will invest less in stocks in the low-return state and more in the high-return state, he will *unconditionally* time the market (that is, adjust his portfolio based on day one information) and this gives him a better date-zero (unconditional) Sharpe ratio than he would obtain by fixing his allocation at date zero. This fact captures the intuition that there is a lot more money to be made from a 50 percent R^2 at a daily horizon than at a five-year horizon, where the calculations in table 1 are not affected by the persistence of the market-timing signal. Campbell and Viceria's (1999) utility calculations are also based on the unconditional distribution, so the optimal degree and benefit of market-timing might be less, conditional on the observed d/p ratio at the first date.

Second, there are good statistical reasons to think that the regressions overstate the predictability of returns.

1) Figure 6 emphasizes one reason: The d/p ratio signal has only crossed its mean four times in the 50 years of postwar history. You have to be very patient to profit from this trading rule. Also, we really have only four postwar data points on the phenomenon. 2) The dividend/price ratio was selected, in sample, among hundreds of potential forecasting variables. It has not worked well out of sample—the last two years of high market returns with low d/p ratios have cut the estimated predictability in half! 3) The model imposes a linear specification, where the actual predictability is undoubtedly better modeled by some unknown nonlinear function. In particular, the linear specification implies negative expected stock returns at many points in the sample, and one might not want to take this specification seriously for portfolio construction. 4) The d/p ratio is strongly autocorrelated, and estimates of this autocorrelation are subject to econometric problems. For this reason, long-horizon return properties inferred from a regression such as equations 1 and 2 are often more dramatic and apparently more precisely measured than direct long-horizon estimates.

The natural next step is to include this parameter uncertainty in the portfolio problem, as I did above for the case of independent returns. While this has not been done yet in a model with Campbell and Vicerá's (1999) level of realism (and for good reasons—Campbell and Vicerá's non-Bayesian solution is already a technical tour de force), Barberis (1999) makes such calculations in his simpler formulation. He uses a utility of terminal wealth and no intermediate trading, and he forces the allocation to stocks to be less than 100 percent.

Figure 7 presents Barberis's (1999) results.⁵ As the figure shows, uncertainty about the parameters of

the regression of returns on d/p almost eliminates the usefulness of market-timing.


Third, it is uncomfortable to note that fund returns still cluster around the (buy-and-hold) market Sharpe ratio (see figure 7 of "New facts in finance"). Here is a mechanical strategy that supposedly earns average returns twice those of the market with no increase in risk. If the strategy is real and implementable, one must argue that funds simply failed to follow it.

Market-timing, like value, does require patience and the willingness to stick with a portfolio that departs from the indexing crowd. For example, a market-timer following Campbell and Vicerá's rules in figures 5 and 6 would have missed most of the great runup in stocks of the last few years. Fund managers who did that are now unemployed. On the other hand, if an eventual crash comes, the market timer will look wise.

Finally, one's reluctance to take such strong market-timing advice reflects the inescapable fact that getting more return requires taking on more, or different, kinds of risk. A market-timer must buy at the bottom, when everyone else is in a panic; he must sell at the top (now) when everyone else is feeling flush. His portfolio will have a greater mean for a given level of variance over very long horizons, but it will do well and badly at very different times from everyone else's portfolio. He will often underperform a benchmark.

Hedging demands

Market-timing addresses whether you should *change* your allocation to stocks over time as a return signal rises or falls. Hedging demands address whether your *overall* allocation to stocks, or to specific portfolios, should be higher or lower as a result of



return predictability, in order to protect you against reinvestment risk.

A long-term bond is the simplest example. Suppose you want to minimize the risk of your portfolio ten years out. If you invest in apparently safe short-term risk-free assets like Treasury bills or a money-market fund, your ten-year return is in fact quite risky, since interest rates can fluctuate. You should hold a ten-year (real, discount) bond. Its price will fluctuate a lot as interest rates go up and down, but its value in ten years never changes.

Another way of looking at this situation is that, if interest rates decline, the price of the ten-year bond will skyrocket; it will skyrocket just enough so that, reinvested at the new lower rates, it provides the same ten-year return as it would have if interest rates had not changed. Changes in the ten-year bond value *hedge* the reinvestment risk of short-term bonds. If lots of investors want to secure the ten-year value of their portfolios, this will raise demand for ten-year bonds and lower their prices.

In general, the size and sign of a hedging demand depend on risk aversion and horizon and, thus, will be different for different investors. If the investor is quite risk averse—infinately so in my bond example—he wants to buy assets whose prices go up when expected returns decline. But an investor who is not so risk averse might want to buy assets whose prices go up when expected returns *rise*. If the investor is sitting around waiting for a good time to invest, and is willing to pounce on good (high expected return) investments, he would prefer to have a lot of money to invest when the good opportunity comes around. It turns out that the dividing line in the standard (CRRA) model is logarithmic utility or a risk aversion coefficient of 1—investors more risk averse than this want assets whose prices go up when expected returns decline, and vice versa. Most investors are undoubtedly more risk averse than this, but not necessarily all investors. Horizon matters as well. A short horizon investor cares nothing about reinvestment risk and, therefore, has zero hedging demand.

In addition, the relationship between price and expected returns is not so simple for stocks as for bonds and must be estimated statistically. The predictability evidence reviewed above suggests that high stock returns presage lower subsequent returns. High returns drive up price/dividend, price/earnings, and market/book ratios, all of which have been strong signals of lower subsequent returns. Therefore, stocks are a good hedge against their own reinvestment risk—they act like the long-term assets that they are. This consideration raises the attractiveness of stocks

for typical (risk aversion greater than 1) investors. Precisely, if the two-fund analysis of figure 1 suggests a certain split between stocks and short-term bonds for a given level of risk aversion and investment horizon, then return predictability, a long horizon, and typical risk aversion greater than 1 will result in a higher fraction devoted to stocks. Again, exactly how *much* more one should put into stocks in view of this consideration is a tough question.

(In this case, the hedging demand reduces to much the same logic as the horizon effects described above. The market portfolio is a good hedge against its own reinvestment risk, and so its long horizon variance is less than its short horizon variance would suggest. More generally, hedging demands can tilt a portfolio toward stocks whose returns better predict and, hence, better hedge the expected return on the market index, but this long-studied possibility from Merton [1971a, 1971b] has not yet been implemented in practice.)

Campbell and Vicerá's (1999) calculations address this hedging demand as well as market-timing demand, and figure 5 also illustrates the strength of the hedging demand for stocks. Campbell and Vicerá's investors want to hold almost 30 percent of their wealth in stocks even if the expected return of stocks is no greater than that of bonds. Absent the hedging motive, of course, the optimal allocation to stocks would be zero with no expected return premium. Almost a 2 percent *negative* stock return premium is necessary to dissuade Campbell and Vicerá's investors from holding stocks. At the average (roughly 6 percent) expected return, of the roughly 130 percent of wealth that the risk aversion 4 investors want to allocate to stocks, nearly half is due to hedging demand. Thus, hedging demands can importantly change the allocation to stocks.

However, hedging demand works in opposition to the usual effects of risk aversion. Usually, less risk averse people want to hold more stocks. However, less risk averse people have lower or even negative hedging demands, as explained above. It is possible that hedging demand exactly offsets risk aversion; everybody holds the same mean allocation to stocks. This turns out not to be the case for Campbell and Vicerá's numerical calibration; less risk averse people still allocate more to stocks on average, but the effect depends on the precise specification.

Choosing a risk-free rate

Figure 1 describes a portfolio composed of the market portfolio and the risk-free rate. But the risk-free rate is not as simple as it once was either. For a consumer or an institution⁶ with a one-year horizon,

one-year bonds are risk-free, while for one with a ten-year horizon, a ten year zero-coupon bond is risk-free. For a typical consumer, whose objective is lifetime consumption, an interest-only strip (or real level annuity) is in fact the risk-free rate, since it provides a riskless coupon that can be consumed at each date. Campbell and Vicera (1998) emphasize this point. Thus, the appropriate bond portfolio to mix with risky stocks in the logic of figure 1 is no longer so simple as a short-term money market fund.

Of course, these comments refer to *real* or indexed bonds, which are only starting to become easily available. When only nominal bonds are available, the closest approximation to a risk-free investment depends additionally on how much interest rate variability is due to real rates versus nominal rates. In the extreme case, if real interest rates are constant and nominal interest rates vary with inflation, then rolling over short-term nominal bonds carries less long-term real risk than holding long-term nominal bonds. In the past, inflation was much more variable than real interest rates in the U.S., so the fact that portfolio advice paid little attention to the appropriate risk-free rate may have made sense. We seem to be entering a period in which inflation is quite stable, so *real* interest rate fluctuations may dominate interest rate movements. In this case, longer term nominal bonds become more risk-free for long-term investors, and inflation-indexed bonds open up the issue in any case. Once again, new facts are opening up new challenges and opportunities for portfolio formation.

Notes of caution

The new portfolio theory can justify all sorts of interesting new investment approaches. However, there are several important qualifications that should temper one's enthusiasm and that shade portfolio advice back to the traditional view captured in figure 1.

The average investor holds the market

The portfolio theory that I have surveyed so far asks, given multiple factors or time-varying investment opportunities, How should an investor *who does not care* about these extra risks profit from them? This may result from intellectual habit, as the past great successes of portfolio theory addressed such investors, or it may come from experience in the money management industry, where distressingly few investors ask about additional sources of risk that multifactor models and predictable returns suggest should be a major concern.

Bear in mind, however, that *the average investor must hold the market portfolio*. Thus, *multiple factors and return predictability cannot have any portfolio*

implications for the average investor. In addition, for every investor who should follow a value strategy or time the market for the extra returns offered by those extra risks, *there must be an investor who should follow the exact opposite advice*. He should follow a growth strategy or sell stocks at the bottom and buy at the top, because he is unusually exposed to or averse to the risks of the value or market-timing strategies in his business or job. He knows that he pays a premium for not holding those risks, but he rationally chooses this course just as we all choose to pay a premium for home insurance.

Again, dividend/price, price/earnings, and book/market ratios forecast returns, if they do, *because* the average investor is unwilling to follow the value and market-timing strategies. If everyone tries to time the market or buy more value stocks, the premiums from these strategies will disappear and the CAPM, random walk view of the market will reemerge. Market-timing can only work if it involves buying stocks when nobody else wants them and selling them when everybody else wants them. Value and small-cap anomalies can only work if the average investor is leery about buying financially distressed and illiquid stocks. Portfolio advice to follow these strategies *must* fall on deaf ears for the average investor, and a large class of investors must want to head in exactly the other direction. If not, the premiums from these strategies will not persist.

One can see a social function in all this: *The stock market acts as a big insurance market*. By changing weights in, say, recession-sensitive stocks, people whose incomes are particularly hurt by recessions can purchase insurance against that loss from people whose incomes are not hurt by recessions. They pay a premium to do so, which is why investors are willing to take on the recession-related risk.

The quantitative portfolio advice is all aimed at the providers of insurance, which may make sense if the providers are large wealthy investors or institutions. But for each provider of insurance, there must be a purchaser, and his portfolio must take on the opposite characteristics.

Are the effects real or behavioral, and will they last?

So far, I have emphasized the view that the average returns from multifactor or market-timing strategies are earned as compensation for holding real, aggregate risks that the average investor is anxious not to hold. This view is still debated. Roughly half of the academic studies that document such strategies interpret them this way, while the other half interpret them as evidence that investors are systematically irrational. This half argues that a new "behavioral finance"

should eliminate the assumption of rational consumers and investors that has been at the core of all economics since Adam Smith, in order to explain these asset pricing anomalies.

For example, I have followed Fama and French's (1993, 1996) interpretation that the *value effect* exposes the investor to systematic risks associated with economywide financial distress. However, Lakonishok, Shleifer, and Vishny (1994) interpret the same facts as evidence for irrationality: Investors flock to popular stocks and away from unpopular stocks. The prices of the unpopular stocks are depressed, and their average returns are higher as the fad slowly fades. Fama and French point out that the behavioral view cannot easily account for the comovement of value stocks; the behavioral camp points out that the fundamental risk factor is still not determined.

Similarly, the predictability of stock returns over time is interpreted as waves of irrational exuberance and pessimism as often as it is interpreted as time-varying, business cycle related risk or risk aversion. Those who advocate an economic interpretation point to the association with business cycles (Fama and French, 1989) and to some success for explicit models of this association (Campbell and Cochrane, 1999); those who favor the irrational investors view point out that the rational models are as yet imperfect.

While this academic debate is entertaining, how does it affect a practical investor who is making a portfolio decision? At a basic level, it does not. If *you* are not exposed to the risk a certain investment represents, it does not matter why other investors shy away from holding it.

Analogously, to decide what to buy at the grocery store, you only have to know how you feel about various foods and what their prices are. You do not have to understand the economic determinants of food prices: You do not need to know whether a sale on tomatoes represents a "real" factor like good weather in tomato growing areas, or whether it represents an "irrational" fear or fad.

Will they last?

Investments do not come with average returns as clearly marked as grocery prices, however. Investors have to figure out whether an investment opportunity that did well in the past will continue to do well. This is one reason that it is important to understand whether average returns come from real or irrational aversion to risk.

If it is *real*, it is most likely to persist. If a high average return comes from exposure to risk, well understood and widely shared, that means all investors understand the opportunity but shrink from it. Even if

the opportunity is widely publicized, investors will not change their portfolio decisions, and the relatively high average return will remain.

On the other hand, if it is truly *irrational*, or a market inefficiency, it is least likely to persist. If a high average return strategy involves no extra exposure to real risks and is easy to implement (it does not incur large transaction costs), that means that the average investor will immediately want to invest when he hears of the opportunity. News travels quickly, investors react quickly, and such opportunities vanish quickly.

Recent work in behavioral finance tries to document a way that irrational phenomena can persist in the face of the above logic. If an asset-pricing anomaly corresponds to a fundamental, documented, deeply formed aspect of human psychology, then the average investor may *not* pounce on the strategy the minute he hears of it, and the phenomenon may last (DeBondt and Thaler, 1985, and Daniel, Hirshleifer, and Subrahmanyam, 1998). For example, many people systematically overestimate the probability that airplanes crash, and make wrong decisions resulting from this belief, such as choosing to drive instead. No amount of statistics changes this view. Most such people readily admit that a fear of flying is "irrational" but persist in it anyway. If an asset-pricing anomaly results from such a deep-seated misperception of risk, then it could in fact persist.

A final possibility is that the average return premiums are the result of *narrowly held risks*. This view is (so far) the least stressed in academic analysis. In my opinion, it may end up being the most important. It leads to a view that the premiums will be moderately persistent. Catastrophe-insurance enhanced bonds provide a good example of this effect. These bonds pay well in normal times, but either part of the principal or interest is pledged against a tranche of a property reinsurance contract. Thus, the bonds promise an average return of 10 percent to 20 percent (depending on one's view of the chance of hurricanes). However, the risk of hurricane damage is uncorrelated with anything else, and hence it is perfectly diversifiable. Therefore, catastrophe bonds are an attractive opportunity. Before the introduction of catastrophe bonds, there was no easy way for the average investor or fund to participate in property reinsurance. As more and more investors and funds hold these securities, the prices will rise and average returns will fall. Once the risks are widely shared, every investor (at least those not located in hurricane-prone areas) will hold a little bit of the risk and the high average returns will have vanished.

The essential ingredients for this story are that the risk is narrowly shared; the high average returns only disappear when the risk is widely shared (it cannot be arbitrated away by a few savvy investors); and an institutional change (the introduction, packaging, and marketing of catastrophe-linked bonds) is required before it all can happen.

This story gives a plausible interpretation of many of the anomalies I document above. Small-cap stocks were found in about 1979 to provide higher returns than was justified by their market (β) risk. Yet at that time, most funds did not invest in such stocks, and individual investors would have had a hard time forming a portfolio of small-cap stocks without losing all the benefits in the very illiquid markets for these stocks. The risks were narrowly held. After the popularization of the small-cap effect, many small-cap funds were started, and it is now easy for investors to hold such stocks. As the risk has been more widely shared, the average returns seem to have fallen.

The value effect may be amenable to a similar interpretation. Before about 1990, as I noted earlier, few funds actually followed the high-return strategy of buying really distressed stocks or shorting the popular growth stocks. It would be a difficult strategy for an individual investor to follow, requiring courage and frequent trading of small illiquid stocks. Now that the effect is clear, value funds have emerged that really do follow the strategy, and the average investor can easily include such an offering in his portfolio. The risk is becoming widely shared, and its average return seems to be falling.

Even average returns on the stock market as a whole (the equity premium) may follow the same story, since participation has increased a great deal through the invention of index funds, low-commission brokerages, and tax-sheltered retirement plans.

This story does not mean that the average returns corresponding to such risks will vanish. They will decline, however, until the markets have established an equilibrium, in which every investor has bought as much of the risk as he likes. In this story, one would expect a large return as investors discover each strategy and bid prices up to their equilibrium levels. This may account for some of the success of small and value stocks observed in the literature, as well as some of the stunning success of the overall market in recent years.

Inconsistent advice

Unfortunately, *the arguments that a factor will persist are all inconsistent with aggressive portfolio advice*. If the premium is *real*, an equilibrium reward

for holding risk, then the average investor knows about it but does not invest because the extra risk exactly counteracts the extra average return. If more than a minuscule fraction of investors are not already at their best allocations, then the market has not reached equilibrium and the premiums will change.

If the risk is *irrational*, then by the time you and I know about it, it's gone. An expected return corresponding to an irrational risk premium has the strongest portfolio implications—everyone should do it—but the shortest lifetime. Thus, this view is also inconsistent with the widespread usefulness of portfolio advice.

If the average return comes from a *behavioral* aversion to risk, it is just as inconsistent with widespread portfolio advice as if were real. We can not all be less behavioral than average, just as we can not all be less exposed to a risk than average. The whole argument for behavioral persistence is that the average investor would not change his portfolio, just as the average traveler would not quickly adjust his traveling behavior to fear the cab ride out to the airport more than the flight. Thus, the advice *must* be useless to the vast majority of investors. If most people, on seeing the strategy, can be persuaded to act differently and buy, then it is an irrational risk and will disappear. If it is real or behavioral and will persist, then this *necessarily* means that very few people will follow the portfolio advice.

If the average return comes from a *narrowly held* risk, one has to ask what institutional barriers keep investors from sharing this risk more widely. Simple portfolio advice may help a bit—most investors still do not appreciate the risk/return advantages of stocks overall, small-caps, value stocks, market-timing, and aggressive liquidity trades. But by and large, a risk like this needs packaging, securitizing, and marketing more than advice. Then there will be a period of high average returns to the early investors, followed by lower returns, but still commoditization of the product with fees for the intermediaries.

Economic logic

The issue of why the risk gives an average return premium is also important to decide whether the opportunity is really there. It is not that easy to establish the average returns of stocks and dynamic portfolio strategies. There are many statistical anomalies that vanish quickly out of sample. Figuring out *why* a strategy carries a high average return is one of the best ways to ensure that the high average return is really there in the first place. Anything that is going to work has a real economic function. A story such

as “I don’t care much about recessions; the average investor does; hence it makes good sense for me to buy extra amounts of recession sensitive stocks since I am selling insurance to the others at a premium” makes a strategy much more plausible than the output of some statistical black box.

Conclusion

Practical application of portfolio theory

How does an investor who is trying patiently to sort through the bewildering variety of investment opportunities use all the new portfolio theory? It’s best to follow a step by step procedure, starting with a little introspection.

1. *What is your overall risk tolerance?* As before, you must first figure out to what extent you are willing to trade off volatility for extra average returns, to determine an appropriate overall allocation to risky versus risk-free assets. While this question is hard to answer in the abstract, you only need to know whether you are more or less risk tolerant than the average investor. (Honestly, now—everyone wants to say they are a risk taker.) The overall market is about 60 percent stocks and 40 percent bonds, so average levels of risk aversion, whatever they are, wind up at this value.

2. *What is your horizon?* This question is first of all important for figuring out what is the relevant risk-free asset. Longer term investors can hold longer term bonds despite their poor one-year performance, especially in a low-inflation environment. Second, we have seen that stocks are somewhat safer for “long-run” investors.

3. *What are your risks?* Would you be willing to trade some average return in order to make sure that your portfolio does well in particular circumstances? For example, an investor who owns a small company would not want his investment portfolio to do poorly at the same time that his industry suffers a downturn, that there is a recession, or a credit crunch, or that the industries he sells to suffer a downturn. Thus, it makes good sense for him to avoid stocks in the same industry or downstream industries, or stocks that are particularly sensitive to recessions or credit crunches, or even to short them if possible. This strategy would make sense even if these stocks give high average returns, like the value portfolios. Similarly, he should avoid high yield bonds that will all do badly in a credit crunch. If the company will do poorly in response to increases in interest rates, oil prices or similar events, and if the company does not hedge these risks, then the investor should take positions in interest rate sensitive or oil-price sensitive securities to offset those

risks as well. We’re just extending the principles behind fire and casualty insurance to investment portfolios.

This logic extends beyond the kind of factors (size, book to market, and so on) that have attracted academic attention. It applies to any identifiable movement in asset portfolios. For example, industry portfolios are not badly explained by the CAPM, as they all seem to have about the same average return. Therefore, they do not show up in multifactor models. However, shorting your industry portfolio protects you against the risks of your occupation. In fact, factors that do *not* carry unusual risk premiums are even better opportunities than the priced factors that attract attention, since you buy insurance at zero premium. This was always true, even in the CAPM, unpredictable return view. I think that the experience with multifactor models just increases our awareness of how important this issue is.

4. *What are not your risks?* Next, figure out what risks you do not face, but that give rise to an average return premium in the market because most other investors do face these risks. For example, an investor who has no other source of income beyond his investment portfolio does not particularly care about recessions. Therefore, he should buy extra amounts of recession-sensitive stocks, value stocks, high yield bonds, etc., if these strategies carry a credible high average return. This action works just like selling insurance, in return for a premium. This is the type of investor for whom all the portfolio advice is well worked out.

In my opinion, too many investors think they are in this class. The extra factors and time-varying returns would not be there (and will quickly disappear in the future) if lots of people were willing and able to take them. The presence of multiple factors wakes us up to the possibility that we, like the average investor, may be exposed to extra risks, possibly without realizing it.

5. *Apply the logic of the multifactor-efficient frontier.* Figure 2 now summarizes the basic advice. After thinking through which risk factors are good to hold, and which ones you are already too exposed to; after thinking through what extra premiums you are likely to get for taking on extra risks, you can come to a sensible decision about which risks to take and which to hedge.

6. *Do not forget, the average investor holds the market.* If you’re pretty much average, all this thought will lead you right back to holding the market index. To rationalize anything but the market portfolio, you have to be different from the average investor in some identifiable way. The average investor sees some risk in value stocks that counteracts their attractive average returns. Maybe you should too! Right now the average

investor is feeling very wealthy and risk-tolerant, therefore stock prices have risen to unprecedented levels and expected stock returns look very low. It's tempting to sell, but perhaps you're feeling pretty wealthy and risk-tolerant as well.

7. *Of course, avoid taxes and snake oil.* The marketing of many securities and funds is not particularly clear on the nature of the risks. There is no reliable extra return without risk. The economic reasoning in this article should be useful to figure out exactly what type of risk a specific fund or strategy is exposed to, and then whether it is appropriate for you. The average actively managed fund still underperforms its style benchmark, and past performance has almost no information about future performance.

The most important piece in traditional portfolio advice applies as much as ever: *Avoid taxes and transaction costs.* The losses from churning a portfolio and paying needless short-term capital gain, inheritance, and other taxes are larger than any of the multifactor and predictability effects I have reviewed. Tax issues are much less fun but more important to the bottom line.

A big insurance market

It is tempting to think of asset markets like a racetrack, but they are in reality a big insurance market. Value funds seem to provide extra returns to their investors by buying distressed stocks on the edge of bankruptcy. Long-Term Capital Management was, it seems, providing catastrophe insurance by intermediating liquid assets that investors like into illiquid assets that were vulnerable to a liquidity crunch. Who better to provide catastrophe insurance than rich investors with no other labor income or other risk exposure? Once again, we are reminded that Adam Smith's invisible hand guides self-interested decisions to socially useful ends, often in mysterious ways.

However, asset markets could be better insurance markets. Both new and old portfolio advice implies that the typical investor should hold a stock position that is *short* his company, industry, or other easily hedgeable kinds of risk. Many managers and some senior employees must hold long positions in their own companies, for obvious incentive reasons. But there is no reason that this applies to union pension funds, for example. A little marketing and help from policy should make funds that hedge industry-specific risks to labor income much more attractive vehicles.

APPENDIX

Multifactor portfolio mathematics

This section summarizes algebra in Fama (1996). The big picture is that we still get a hyperbolic region since betas are linear functions of portfolio weights just like means.

The problem is, minimize the variance of a portfolio given a value for the portfolio mean and its beta on some factor. Let

$$w = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_N \end{bmatrix}; R = \begin{bmatrix} R^1 \\ R^2 \\ \vdots \\ R^N \end{bmatrix}; 1 = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}; \beta = \begin{bmatrix} \beta_{1,F} \\ \beta_{2,F} \\ \vdots \\ \beta_{N,F} \end{bmatrix}.$$

Then the portfolio return is

$$R^p = w'R;$$

the condition that the weights add up to 1 is

$$1 = 1'w.$$

The mean of the portfolio return is

$$E(R^p) = E(w'R) = w'E(R) = w'E.$$

The last equality just simplifies notation. The beta of the portfolio on the extra factor is

$$\beta^p = w'\beta.$$

The variance of the portfolio return is

$$\text{var}(R^p) = w'Vw,$$

where V is the variance-covariance matrix of returns. The problem is then

$$\min_w \frac{1}{2} w'Vw \text{ s.t. } w'E = \mu; w'1 = 1; w'\beta = \beta^p.$$

The Lagrangian is

$$\frac{1}{2} w'Vw - \lambda_0(w'E - \mu) - \lambda_1(w'1 - 1) - \lambda_2(w'\beta - \beta^p).$$

The first order conditions with respect to w give

$$w = V^{-1}(E\lambda_0 + 1\lambda_1 + \beta\lambda_2) = V^{-1}A\lambda$$

where

$$A = \begin{bmatrix} E & 1 & \beta \end{bmatrix}$$

$$\lambda = \begin{bmatrix} \lambda_0 & \lambda_1 & \lambda_2 \end{bmatrix}'$$

$$\delta = \begin{bmatrix} \mu & 1 & \beta^p \end{bmatrix}'.$$

Plugging this value of w into the constraint equations

$$A'w = \delta,$$

we get

$$AV^{-1}A\lambda = \delta$$

$$\lambda = (A'V^{-1}A)^{-1}\delta$$

$$w = V^{-1}A(A'V^{-1}A)^{-1}\delta.$$

The portfolio variance is then

$$\text{var}(R^p) = w'Vw = \delta'(A'V^{-1}A)^{-1}\delta.$$

Or, writing out the sum of the matrix notation,

$$\text{Var}(R^p) = \begin{bmatrix} \mu & 1 & \beta^p \end{bmatrix} (A'V^{-1}A)^{-1} \begin{bmatrix} \mu \\ 1 \\ \beta^p \end{bmatrix}.$$

The variance is a quadratic function of the mean return and of the desired beta on additional factors. That's why we draw cup-shaped frontiers. As with the mean-variance case, the multifactor efficient frontier is a revolution of a hyperbola. If V is a second moment matrix, to handle a risk-free rate,

$$\text{var}(R^p) = \delta'(A'V^{-1}A)^{-1}A'V^{-1}\Sigma V^{-1}A(A'V^{-1}A)^{-1}\delta,$$

where Σ now denotes the return variance-covariance matrix.

Finding the benefits of a market timing strategy without computing the strategy

I show first that the squared maximum unconditional Sharpe ratio is the average of the squared conditional Sharpe ratios when the riskfree rate is constant,

$$s^{*2} = E[s_t^2],$$

where

$$s^* = \max E(R - R^f) / \sigma(R - R^f)$$

denotes the unconditional Sharpe ratio, and

$$s_t = \max E_t(R - R^f) / \sigma_t(R - R^f)$$

denotes the conditional Sharpe ratio.

The technique exploits ideas from Gallant, Hansen, and Tauchen (1990). I exploit Hansen and Jagannathan's (1991) theorem that for any excess return Z and discount factor m such that $0 = E(mZ)$, we have

$$\frac{E(Z)}{\sigma(Z)} \leq \frac{\sigma(m)}{E(m)},$$

and equality is attained for some choice of m . Thus, the maximal unconditional Sharpe ratio is

$$\max \left(\frac{E}{\sigma} \right) = \frac{\sigma(m^*)}{E(m^*)},$$

where m^* solves

$$m^* = \arg \min_{\{m\}} \sigma(m) \text{ s.t.}$$

$$E_t(m_{t+1}Z_{t+1}) = 0; E_t(m_{t+1}) = 1/R_t^f.$$

Gallant, Hansen, and Tauchen show how to solve this problem in quite general situations. They phrase their result as a "lower bound on discount factor volatility" but given $E(Z)/\sigma(Z) \leq \sigma(m)/E(m)$, one can read the maximum slope of the unconditional mean-variance frontier (Sharpe ratio) available from market-timing portfolios. To keep the calculation transparent and simple, I specialize to the case of a constant and observed real risk-free rate $R^f = 1/E_t(m)$. Then, *the unconditional squared Sharpe ratio is the average of the conditional squared Sharpe ratios*,

$$\frac{\sigma^2(m)}{E(m)^2} = \frac{\sigma[E_t(m)^2] + E[\sigma_t^2(m)]}{E(m)^2} = E \left(\frac{\sigma_t^2(m)}{E_t(m)^2} \right).$$

Next, I show that when we forecast stock returns with a regression such as equation 1, and interest rates and the conditional variance of the error term are constant, then the best unconditional Sharpe ratio is related to the regression R^2 by

$$s^* = \frac{\sqrt{s_0^2 + R^2}}{\sqrt{1 - R^2}},$$

where $s_0 = E(R - R^f) / \sigma(R - R^f)$ denotes the unconditional buy-and-hold Sharpe ratio.

If the conditional Sharpe ratio is generated by a single asset (the market), and a linear model with constant error variance,

$$Z_{t+1} = EZ + b(x_t - Ex) + \varepsilon_{t+1},$$

then,

$$\left(\frac{\sigma_t(m)}{E_t(m)} \right)^2 = \left(\frac{E_t(Z)}{\sigma_t(Z)} \right)^2 = \left(\frac{EZ + b(x_t - Ex)}{\sigma_\varepsilon} \right)^2,$$

and

$$\begin{aligned} E\left(\frac{\sigma_t^2(m)}{E_t(m)^2} \right) &= E\left[\left(\frac{EZ + b(x_t - Ex)}{\sigma_\varepsilon} \right)^2 \right] \\ &= \frac{(E(Z))^2 + b^2 \sigma^2(x)}{\sigma_\varepsilon^2} \\ &= \frac{(E(Z))^2}{(1-R^2)\sigma^2(Z)} + \frac{b^2 \sigma^2(x)}{(1-R^2)\sigma^2(Z)} \\ &= \frac{1}{(1-R^2)} \frac{(EZ)^2}{\sigma^2(Z)} + \frac{R^2}{(1-R^2)} \\ &= \frac{1}{(1-R^2)} \left[\left(\frac{E(Z)}{\sigma(Z)} \right)^2 + R^2 \right]. \end{aligned}$$

The last line demonstrates $s^* = \sqrt{s_0^2 + R^2} / \sqrt{1-R^2}$. To obtain the annualized Sharpe ratios reported in table 1, I divide by the square root of horizon, since mean returns roughly scale with horizon and standard deviations roughly scale with the square root of horizon.

NOTES

¹To be precise, these statements refer to the conditional serial correlation of returns. It is possible for the conditional serial correlations to be non-zero, resulting in conditional variances that increase with horizon faster or slower than linearly, while the unconditional serial correlation of returns is zero. Conditional distributions drive portfolio decisions.

²This effort falls in a broader inquiry in economics. Once we recognize that people are unlikely to have much more data and experience than economists, we have to think about economic models in which people *learn* about the world they live in through time, rather than models in which people have so much history that they have learned all there is to know about the world. See Sargent (1993) for a review of learning in macroeconomics.

³The standard first-order condition for optimal consumption and portfolio choice is

$$3) \quad E[(c_{t+1})^{-\gamma} Z_{t+1}] = 0,$$

where c denotes consumption, Z denotes an excess return, and γ is a preference parameter. We usually take data on c and Z , estimate γ , and then test whether the condition actually does hold across assets. In a portfolio problem, however, we *know* the preference parameter γ , but we want to estimate the

portfolio. For example, in the simplest case of a one-period investment problem, consumption equals terminal wealth. Equation 3 then becomes

$$E[(\alpha R^f + (1 - \alpha) R_{t+1}^m)^{-\gamma} Z_{t+1}] = 0.$$

Brandt uses this condition to estimate the portfolio allocation α . He extends the technique to multiperiod problems and problems in which the allocation decision depends on a forecasting variable, that is, market-timing problems.

⁴That is, Campbell and Viceria model investors' objectives by a utility function, $\max E \sum_t \beta^t u(c_t)$ rather than a desire for wealth at some particular date, $\max Eu(W_T)$.

⁵I thank Nick Barberis for providing this figure. While it is not in Barberis (1999), it can be constructed from results given in that paper.

⁶Of course, in theory, institutions, as such, should not have preferences, as their stockholders or residual claimants can unwind any portfolio decisions they make. This is the famous Modigliani-Miller theorem. In practice, institutions often make portfolio decisions as if they were individuals, and people surveying portfolio advice will run into many such institutions.

REFERENCES

- Barberis, Nicholas**, 1999, "Investing for the long run when returns are predictable," *Journal of Finance*, forthcoming.
- Black, Fischer, and Robert Litterman**, 1991, "Global asset allocation with equities, bonds, and currencies," *Goldman Sachs Fixed Income Research*.
- Brandt, Michael W.**, 1999, "Estimating portfolio and consumption choice: A conditional Euler equations approach," *Journal of Finance*, October, forthcoming.
- Brennan, Michael J., Eduardo S. Schwartz, and Roland Lagnado**, 1997, "Strategic asset allocation," *Journal of Economic Dynamics and Control*, Vol. 21, No. 7, pp. 1377–1403.
- Campbell, John Y., and John H. Cochrane**, 1999, "By force of habit: A consumption-based explanation of aggregate stock market behavior," *Journal of Political Economy*, Vol. 107, No. 2, April, pp. 205–251.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay**, 1996, *The Econometrics of Financial Markets* Princeton, NJ: Princeton University Press.
- Campbell, John Y. and Luis M. Vicera**, 1999, "Consumption and portfolio decisions when expected returns are time varying," *Quarterly Journal of Economics*, forthcoming.
- _____, 1998, "Who should buy long term bonds?," Harvard University, manuscript.
- Cochrane, John H., and Argia M. Sbordone**, 1988, "Multivariate estimates of the permanent components in GNP and stock prices," *Journal of Economic Dynamics and Control*, Vol. 12, No. 2, pp. 255–296.
- Cochrane, John H.**, 1997, "Where is the market going? Uncertain facts and novel theories," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 21, No. 6, November/December, pp. 3–37.
- _____, 1994, "Permanent and transitory components of GNP and stock prices," *Quarterly Journal of Economics*, Vol. 109, February, pp. 241–266.
- Daniel, Kent, David Hirshleifer, and Anandhar Subrahmanyam**, 1998, "Investor psychology and security market under- and overreactions," *Journal of Finance*, Vol. 53, No. 6, pp. 1839–1885.
- DeBonds, Werner F. M., and Richard H. Thaler**, 1985, "Does the stock market overreact?," *Journal of Finance*, Vol. 40, No. 3, pp. 793–808.
- Fama, Eugene F.**, 1996, "Multifactor portfolio efficiency and multifactor asset pricing," *Journal of Financial and Quantitative Analysis*, Vol. 31, No. 4, December, pp. 441–465.

Fama, Eugene F., and Kenneth R. French, 1993, “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics*, Vol. 33, No. 1, February, pp. 3–56.

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

Gallant, A. Ronald, Lars Peter Hansen, and George Tauchen, 1990, “Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution,” *Journal of Econometrics*, Vol. 45, No. 1/2, pp. 141–179.

Kandel, Shmuel, and Robert F. Stambaugh, 1996, “On the predictability of stock returns: An asset allocation perspective,” *Journal of Finance*, Vol. 51, No. 2, June, pp. 385–424.

Kim, Tong Suk, and Edward Omberg, 1996, “Dynamic nonmyopic portfolio behavior,” *Review of Financial Studies*, Vol. 9, No. 1, Winter, pp. 141–161.

Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny, 1994, “Contrarian investment, extrapolation and risk,” *Journal of Finance*, Vol. 49, No. 5, December, pp. 1541–1578.

Markowitz, H., 1952, “Portfolio selection,” *Journal of Finance*, Vol. 7, No. 1, pp. 77–91.

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

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

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

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

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