

NBER WORKING PAPER SERIES

PORTFOLIO ADVICE FOR A
MULTIFACTOR WORLD

John H. Cochrane

Working Paper 7170
<http://www.nber.org/papers/w7170>

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.

© 1999 by John H. Cochrane. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Portfolio Advice for a Multifactor World
John H. Cochrane
NBER Working Paper No. 7170
June 1999
JEL No. G00

ABSTRACT

Asset returns, it turns out, do not follow the Capital Asset Pricing Model, and are somewhat predictable over time. I survey and interpret the large body of recent work that adapts traditional portfolio theory to answer, what should an investor do about these new facts in finance? I survey the extension of the famous “2 - fund” theorem to an “N-fund” theorem in which investors either hedge or assume the additional, non-market, sources of priced risk; I survey the burgeoning literature on time-varying portfolio rules and the Bayesian literature that advocates a great deal of caution. In a survey, I emphasize the risk-sharing nature of asset markets, I note the likelihood that many supposed anomalies will not last, and I emphasize the fact that the average investor must hold the market so portfolio decisions must be driven by differences between an investor and the average investor.

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

“New Facts in Finance” summarized the revolution in how financial economists view the world. Briefly, there are strategies that result in high average returns without large “betas,” or a tendency to move with the market as a whole. Multifactor models have supplanted the Capital Asset Pricing Model 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 funds, index funds, and relatively expensive trading on his own account, he now must choose among a bewildering variety of fund “styles,” including “value,” “growth,” “balanced,” “income,” “global,” “emerging market,” and “convergence,” as well as more complex claims of active fund managers with all sorts of customized styles and strategies, and the temptation to trade on his own 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? To answer this question, there is now an important current of academic research that investigates how portfolio theory should adapt to our new view of the world. In this paper, I summarize this research, and I try to distill its advice for investors. In particular, which of the bewildering new investment “styles” should an investor follow? Should he 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 the specific circumstances of an individual investor? 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?

2 The traditional view

Before surveying how new facts have *changed* portfolio advice, it’s useful to remind ourselves what the traditional portfolio advice *is*, and why. Also, as the new facts have really extended rather than overturned the old facts, the new advice has really just

extended traditional advice in important ways. The traditional academic portfolio theory, starting from Markowitz (1952) and expounded in every finance textbook, remains one of the most beautiful, surprising, useful and enduring bits of economics developed in the last 50 years.

Advice: a two-fund theorem

The traditional advice is to split your investment between a money-market fund and a broad-based, passively managed stock fund. That fund should concentrate on minimizing fees and transactions costs, period. It should avoid the temptation to actively manage its portfolio, trying to chase the latest hot stock. 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 indifference curves capture the fact that investors want portfolios with greater mean return and lower return variance – portfolios that are higher up and to the left – and that they are willing to accept more volatile portfolios if they get a higher average return.

The *mean-variance frontier* gives the minimum possible variance of a portfolio return for each level of mean portfolio return. It summarizes set of available portfolios. We construct this frontier in two steps, the curved mean-variance frontier of all risky assets, and then the straight frontier that includes a risk-free rate. Investors should all hold portfolios on the mean-variance frontier, as shown. Furthermore, every portfolio on the frontier can be formed as a combinations of the risk free asset and the market portfolio. Therefore every investor need only hold different proportions of these *two funds*.

2.1 Bad portfolio advice

The portfolio advice is not so remarkable for what it does say, which given the setup is fairly straightforward, as it is 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,” “large-cap,” “income,” “capital preservation,” or 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

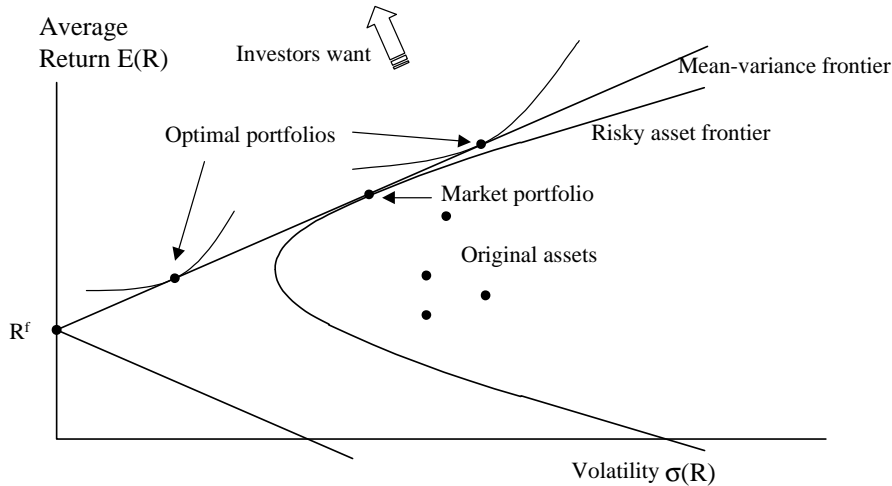


Figure 1: Mean-variance frontier, optimal portfolios and two-fund theorem.

prudent. The advertisements that promise “we listen” “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 he will 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 he will 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*. If you know (say) 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 will the market portfolio held by the average investor. All the analysis of Figure 1 goes through, 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 this reason, the standard advice is to hold passively-managed funds that concentrate on minimizing transactions costs and fees, rather than a carefully constructed “tangency portfolio” that reflects an investor’s or

manager's special insights. However, an 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)

Second, 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. Therefore, if returns are independent over time and stocks follow “random walks,” stocks are *not* safer in the long run, and the stock/bond allocation is 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 10 year returns are 10 times those of 1 year returns, so the ratio of mean to variance is the same at all horizons. More elegantly, Merton (1969) and Samuelson (1969) showed that an investor with a constant relative risk aversion utility and 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.

2.2 Taking the advice

This advice has had a sizeable impact on portfolio practice. Before the early 1970s, when this advice became standard in academia, 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 indices ex-post, where the proposition that professional active management and stock selection could outperform blindly holding an index seemed self-evident before 1970.

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 by 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 in this journal.) However, the basic question is whether

you are 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 at the logical inconsistency of providing portfolio advice based on a view of the world in which everyone is already following such advice. (This is what allowed us 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.

3 New Portfolio Theory

3.1 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. Most of the theory was first worked out by Merton 1969, 1971a, 1971b.)




Figure 2. Portfolio theory in a multifactor world. The left hand panel shows an indifference surface and optimal portfolio in the case with no riskfree rate. The dot marks the optimal portfolio where the indifference sheet touches the multifactor efficient frontier. The right hand panel shows the set of multifactor efficient portfolios with a riskfree rate. The two cone-shaped surfaces intersect on the marked line. The two dots are the market portfolio and an additional multifactor-efficient portfolio; all multifactor-efficient portfolios on the outer cone can be reached by combinations of the risk free rate, the market, and the extra multifactor-efficient portfolio.

To keep the figure simple I consider one additional factor, and 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 want portfolios that do not tend to go down in recessions. Investors 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 also farther out (lower recession sensitivity). The indifference *curves* of Figure 1 become indifference *surfaces*. The left hand panel of Figure 2 shows one such surface curving upwards.

As with Figure 1, we 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 the left hand panel 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.

The right hand panel of Figure 2 adds a risk free rate. As the mean-variance frontier of Figure 1 was the minimal V emanating from the risk free rate 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 the right hand panel of Figure 2. (It is especially convenient to take this third portfolio to be a zero-cost portfolio such as HML.)

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*.

3.1.1 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 you can see from the figure, the average investor is willing to trade some mean or 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 an investor who is only concerned only with mean and variance – one who is not exposed to the recession risk, or the risks associated with any other factor, and only wants to get the best possible mean return for given standard deviation. If so, you still want to solve the mean-variance problem of Figure 2. Thus, the important implication of a multifactor world for you is that *you should no longer hold the market portfolio.*

A mean-variance efficient portfolio can still be achieved 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.

3.2 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.

3.2.1 Horizon effects

Recall that when stock returns are independent over time (like coin flips), then 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) 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 run¹. 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 long horizons of 5 years and more are about 1/2 to 2/3 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 time-period studied and the econometric method. Multivariate methods give somewhat stronger evidence. Intuitively, the price/dividend 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. (I used this idea in Cochrane and Sbordone 1986, Cochrane 1994. Campbell Lo and MacKinlay 1996 have a nice summary of these issues and 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 reminds us of the standard result: If returns are not predictable, then the allocation to stocks does not depend on horizon.

The top, dashed 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% allocation to stocks for a monthly horizon to 100% allocation to stocks at a 10 year horizon. To quantify the effects of predictability, we need a mathematical model of predictability, and Barberis takes the simplest model,

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

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

using the dividend price ratio for the forecasting variable x . (Whether one includes returns or not in the right hand side makes little difference.) Barberis estimates significant mean-reversion: In Barberis’ regressions, the implied standard deviation of 10 year returns is 23.7%, just more than half of the 45.2% value implied by the standard deviation of monthly returns. Stocks are indeed safer in the long run, and the greater allocation to stocks for a long-run investor reflects this fact.

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.




Figure 3: Percent allocation to stocks for different investment horizons. The investor maximizes the utility of terminal wealth via a buy-and-hold investment in stocks vs. bonds. The investor has constant relative risk aversion utility with risk aversion coefficient 10. Top line (dashed): the calculation includes predictable returns modeled by a regression (1). Second line from top (solid): the calculation includes predictable returns and the effects of parameter uncertainty. Flat line (dot-dash): the calculation assumes independent unpredictable returns, and no parameter uncertainty. Bottom line (dots): the calculation assumes independent unpredictable returns, but adds parameter uncertainty. All distributions are conditional on a d/p ratio equal to its historical mean. Source: Barberis (1999).

Uncertainty in predictability

This calculation ignores the fact that we really don't know just 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 towards the standard advice. Standard errors don't capture the uncertainties behind this (good) intuition.

First, the predictability results we have settled on certainly result to some extent

from data-dredging. Thousands of series were examined, 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, while keeping intact the list of forecasting variables and functional forms intact, the portfolio calculation assumes that the *investor* knows the return forecasting process perfectly; the standard errors only reflect the fact that *we* don't know the return forecasting process, so we are unsure about what the investor wants to do. But investors are likely to be just as unsure as we are about the exact extent of return predictability.² What we would like to do is to solve a portfolio problem in which the 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 (1998) tackle this important problem.

Figure 3 also gives Barberis' calculations of the effects of parameter uncertainty on the stock/bond allocation problem. The lowest, dotted line considers a simple case. The investor knows, correctly, that returns are independent over time (not predictable), but the investor isn't 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% to about 28% at a 10 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 raises his estimate of the returns over the rest of the investment period. Conversely, a few bad years lowers his estimate of the mean return for remaining years. Thus, learning about parameters induces a 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 solid 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 1 to 10 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.

Even this calculation assumes that the investor knows a lot of things perfectly,

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.

so the actual amount one can optimally gain by exploiting long-horizon mean reversion is likely to be substantially lower. The dividend/price ratio was selected, in sample, among hundreds of potential forecasting variables. 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 advice. The dividend-price ratio is strongly autocorrelated, and estimates of this autocorrelation are subject to “unit root” econometric problems. For this reason, long-horizon return properties inferred from a VAR are often more dramatic and apparently more precisely measured than direct long-horizon estimates.

Box: How to include model uncertainty into portfolio problems.

A statistical model, such as the regression (1) tells us the distribution of future returns once we know the parameters θ , $f(R_{t+1}|\theta)$. We’d really like to evaluate uncertainty by the distribution of returns conditional only on the return history available to make guesses about the future, $f(R_{t+1}|x_1, x_2 \dots x_t)$ where x_t denotes all the data used (returns, d/p, etc.). 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 of returns 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 to 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.

3.2.2 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,

investors might well want to hold more stocks when expected returns are high, and fewer when expected returns are low. Exactly how *much* market timing one should do is of course the crucial question. This is a technically challenging question which several authors have recently addressed.

The Benefits of Market Timing

Much of the difficulty with return predictability (as with other dynamic portfolio questions) lies in computing the optimal strategy – exactly how should one adjust one’s portfolio as the return prediction signals change on each date? Gallant, Hansen and Tauchen (1990) show a clever way 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 riskfree rate is constant and known, it turns out that *the square of the maximum unconditional Sharpe ratio is the average of the squared conditional Sharpe ratios*. (The Appendix derives this result and 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 R^2 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% over Treasury bills, with a standard deviation of about 16% . 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.”)

Horizon k (years)	R^2	Annualized Sharpe ratio
buy&hold		0.5
1	0.17	0.71
2	0.26	0.72
3	0.38	0.78
5	0.59	0.95

Table 1. 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 $s^*/\sqrt{k} = \sqrt{0.5^2 + R^2/k}/\sqrt{1 - R^2}$ and is derived in the Appendix.

As indicated by the table, market-timing should be a great benefit. Holding constant the portfolio volatility, market timing should raise average returns by about 2/5 at an annual horizon, and almost double average returns at a 5 year horizon.

Optimal Market Timing: An Euler Equation Approach

Brandt (1999) presents a very 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, i.e. estimates the optimal consumption or wealth process³. This calculation is very clever because it does not require one to specify a statistical model for the stock returns (for example, a VAR of returns and d/p), 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. We see a mild horizon effect, about in line with Barberis’ results of Figure 3 without parameter uncertainty: longer term investors hold more stocks. We also see 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% to 225% of wealth invested in stocks,

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, which is the right way to pose the problem. They model the time-variation in expected stock returns via the simple vector autoregres-

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

$$E \left((c_{t+1})^{-\gamma} Z_{t+1} \right) = 0 \tag{3}$$

where c denotes consumption, Z denotes an excess return, and γ is a preference parameter. We usually take data on c, 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 $c_{t+1} = W_{t+1} = W_t \times (\alpha R^f + (1 - \alpha)R_{t+1}^m)$. Equation (3) then becomes

$$E \left[(\alpha R^f + (1 - \alpha)R_{t+1}^m)^{-\gamma} Z_{t+1} \right] = 0. \tag{4}$$

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, i.e. market-timing problems.

⁴I.e. Campbell and Vicera model investor’s objectives by a utility function $\max E \sum_t \beta^t u(c_t)$ rather than a desire for wealth at some particular date $\max Eu(W_T)$.

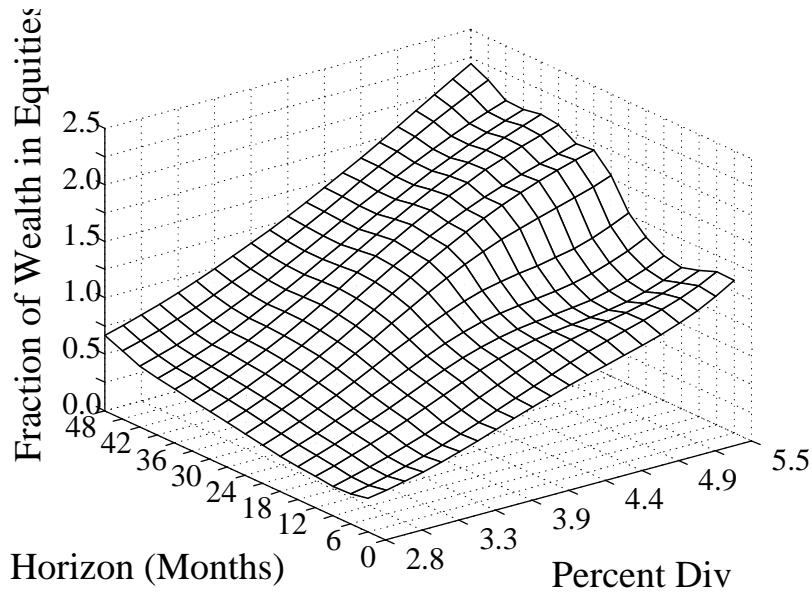


Figure 4: Optimal allocation to stocks as a function of horizon and dividend yield. Source: Brandt (1999).

sion (1) on dividend price ratios. Their investors live only off invested wealth, and have no labor income or labor income risk. Thus their 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 Viceria's optimal allocation to stocks as a function of the expected return, forecast from d/p ratios via (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%, 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% in stocks to 220% 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 dividend/price ratios of the 1950s suggest a strong stock position, and that strong position profits from the high returns of the late 50s to early 60s. 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%.




Figure 5: Optimal allocation to stocks as a function of the expected return implied from a regression that forecasts stock returns from dividend/price ratios. The line extends from a d/p ratio two standard deviations above its mean (low expected returns) to one standard deviation below its mean (high expected returns). Source: Campbell and Vicera (1999).




Figure 6: Optimal allocation to stocks over time, based on dividend/price ratio. Source: Campbell and Vicera (1999).

Campbell and Vicera’s calculations are if anything conservative compared to others in this literature. Other calculations, using other utility functions, solution techniques and calibrations of the forecastability process often produce even more aggres-

sive market timing strategies. For example, Brennan Schwartz and Lagnado (1997) make a similar calculation with two additional forecasting variables. They report (their figure 6) market timing strategies that essentially jump back and forth between constraints at 0% in stocks and 100% in stocks.

Campbell and Vicera present achieved utility calculations that mirror the lesson of Table 1: failing to market-time imposes 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 the near doubling of 5 year Sharpe ratios that the strong market timing strategies give in Table 1. In particular, one might have trouble taking advice that would have pointed one to miss most of the dramatic runup in stock values of the late 90's! Perhaps, rather than a failure of nerve, one's 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) 5 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 5 year returns, what Sharpe ratio does he achieve.” This is *not* “Given today's d/p, what is the best Sharpe ratio you can achieve for the next 5 years by following market-timing signals?” The latter question characterizes the return distribution conditional on today's d/p; the former does not. The answer to the latter question is harder to evaluate, it depends on the initial signal (bad right now, since p/d ratios are at all time highs), and is lower, especially for a slow-moving signal such as d/p .

To see the point, consider an extreme case that the d/p ratio is determined in day 1, is constant thereafter, and indicates high or low returns in perpetuity. *Conditional* on the d/p ratio, one cannot market-time 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* market time (i.e. adjust his portfolio based on day 1 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 ones' intuition that there is a lot more money to be made from a 50% R^2 at a daily horizon than at a 5 year horizon, where the calculations in Table 1 are not affected by the persistence of the market timing signal. Campbell and Vicera's 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. 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 a very patient investor to profit from this trading rule! Also, we really have only 4 postwar data points on the phenomenon. For this reason the magnitude and nature of d/p predictability are still debated.

The natural next step is to include this uncertainty in the portfolio problem, as we did above for the case of independent returns. While this has not been done yet in a model with Campbell and Vicera’s level of realism (and for good reasons – Campbell and Vicera’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%.

Figure 7 presents Barberis’ results⁵. As the figure shows, uncertainty about the parameters of the regression of returns on d/p almost eliminates the usefulness of market timing.




Figure 7: Percentage allocation to stocks as a function of dividend/price ratio. The solid line ignores parameter uncertainty, as in Campbell and Vicera (1999). The dashed line includes parameter uncertainty as in Barberis (1999). T is the data sample in months, A gives the risk aversion coefficient.

Third, it is uncomfortable to note that funds 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 requires patience

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.

and the willingness to suffer a portfolio that departs from the indexing crowd. For example, a market-timer following Campbell and Vicera'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.

Fourth, and finally, one's reluctance to take such strong market timing advice reflects the inescapable fact that you can't get more return without 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 and 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.

3.2.3 Hedging demands

Hedging demands are a bit more subtle. If expected returns vary over time, investors may want to hold assets that protect them against this risk.

The easiest example is a long-term bond. Imagine an investor who wants to minimize the risk of his portfolio ten years out. If he invests in apparently safe short term riskless assets like treasury bills or a money market fund, his 10 year return is in fact quite risky, since interest rates can fluctuate. This investor should hold a 10 year (real, discount) bond. Its price will fluctuate a lot as interest rates go up and down, but its value in 10 years never changes.

Another way of looking at this situation is that, if interest rates decline, the price of the 10 year bond will skyrocket; it will skyrocket just enough so that, reinvested at the new lower rates, it provides the same 10 year return as it would have if interest rates had not changed. Changes in the 10 year bond value *hedge* the reinvestment risk of short term bonds. If lots of investors want to secure the 10 year value of their portfolios, this will raise demand for 10 year bonds and lower their prices.

A hedging demand is different from market timing. Market timing addresses whether you should *change* your allocation to stocks over time as a signal of subsequent returns rises or falls. Hedging demands address whether your *overall* allocation to stocks, or to specific portfolios, should be higher or lower, independent of the level of a signal, in order to protect you against reinvestment risk.

In general, the size and sign of a hedging demand depends on risk aversion and horizon, and thus will be different for different investors. If the investor is quite risk averse – infinitely 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 that he has 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 – consumers 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 P/D, P/E, 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 one) investors. Precisely, if the two-fund analysis of Figure 1 suggested a certain split between stocks and short term bonds for a given level of risk aversion and investment horizon, return predictability, a long horizon and typical risk aversion greater than one will result in a higher fraction devoted to stocks. Again, exactly how *much* more one should put into stocks via 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 towards 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 Vicera’s 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 Vicera’s investors want to hold almost 30% 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% *negative* stock return premium is necessary to dissuade Campbell and Vicera’s investors from holding stocks. At the average (roughly 6%) expected return, of the roughly 130% 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 offset, 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.

3.3 Choosing a riskfree rate

The simple analysis of Figure 1 describes a portfolio composed of the market portfolio and the riskfree rate. But the “riskfree” 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 “riskfree,” while for one with a 10 year horizon, a ten year zero-coupon bond is “riskfree.” For a typical consumer, whose objective is lifetime consumption, an interest-only strip (or real level annuity) is in fact the “riskfree rate,” since it provides a riskless coupon that can be consumed at each date. Campbell and Vicerá (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 “riskfree” investment depends additionally on how much interest rate variability is due to real rates vs. 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 “riskfree” 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.

4 Warnings

The advice of the new portfolio theory is alluring and can justify all sorts of interesting new portfolio approaches. However, there are several important qualifications that should temper one’s enthusiasm, and that shade portfolio advice back to the

⁶Of course, in theory, “institutions” as such shouldn’t 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 purveying portfolio advice will run in to many such institutions.

traditional view captured in Figure 1.

4.1 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.

We must remember, however, that *the average investor must hold the market portfolio*. Thus, *multiple factors and predictability cannot have any portfolio implications for the average investor*. In addition, for every investor who should follow a value-tilt or market-time, taking advantage of the extra returns that exposure to those risks brings, *there must be an investor who should follow the exact opposite advice*. He should follow a growth tilt or sell stocks at the bottom and buy at the top, because he is unusually exposed to or averse to the risks of those strategies in his business or job. He knows that he pays a premium for not holding those risks, but rationally chooses to do so just as you pay a premium for home insurance. If not, if everyone tried to market-time or buy value stocks, the phenomena would disappear and the CAPM, random walk world would reemerge.

Again, P/D, P/E, etc. forecast returns, if they do, *because* the average investor is unwilling to follow the market timing strategy and buy more stocks when prices are low. The strategy can only work if it involves buying stocks when nobody else wants them; after crashes, in the depths of recessions, in times of financial panic; and selling them in booms (like now) when everyone else seems to put more and more into stocks despite very high valuations. “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 strategies can’t work.

You 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 what gets the investors 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 there must be a purchaser of insurance, and his portfolio must take on the opposite characteristics. The advice to “be a provider of insurance” cannot hold for everyone.

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

I have emphasized a view that the average returns from multifactor or market-timing strategies are earned because in following them, an investor takes on a real risk that others are anxious not to hold. This proposition is still debated for many strategies that seem to give high average returns. Roughly half of the academic studies that document such strategies interpret them as I have, while the other half interpret them as evidence that investors are systematically irrational.

For example, I have followed Fama and French’s interpretation of the “value effect” that it exposes the investor to systematic risks associated with economywide financial distress. However, authors such as 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 for example); irrational interpreters point to the fact that the models are as yet imperfect.

While this debate is entertaining, how does it affect a practical person making a portfolio decision? At a most basic level *we don’t have to take a stand on whether an average return is due to real risk or irrationality in making portfolio decisions*. If it’s there, and if *you* are not exposed to the risk it represents, it does not matter whether everyone else is shying away from holding it, and thus driving up its return, due to a real risk factor or due to some irrational fear.

A portfolio problem is economically identical to the shopping problem – what should you buy at the grocery store? To answer that question, you only have to know how *you* feel about various foods and what their prices are. You don’t need to know whether a sale on tomatoes represents a “real” factor like good weather in tomato growing areas, or whether it represents “irrational” or “inefficient” outcome, a sudden fad that people don’t buy tomatoes anymore.

Will they last?

You do have to figure out whether an investment opportunity that did well in the past will continue to do well in the future, and the question whether a high average return comes from a real or irrational aversion to risk does bear on this question.

1. *Real*. If it is real, it is most likely to persist. If a high average return comes exposure to risk, well understood and widely shared, that means all investors understand the opportunity but shrink from it. A horde of investment advisers advertising the opportunity will change nobody's portfolio decisions, and hence the average returns will remain.

2. *Irrational*. 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 risk, and is easy to implement (does not require large transactions costs) that means that the average investor immediately wants to invest when he hears of the opportunity. News travels quickly, investors quickly invest, and such opportunities will vanish quickly.

3. *Behavioral*. Recent work in "behavioral" finance, I think, tries to document a way that "irrational" phenomena can persist in the face of the above logic. If an inefficiency corresponds to a fundamental, documented, deeply formed aspect of psychology, then the average investor will *not* immediately pounce on the strategy the minute he hears of it, and the phenomenon may last. DeBondt and Thaler (1985), Daniel, Hirshleifer and Subrahmanyam (1998) emphasize this view.

For example, most people systematically overestimate the probability of airplanes crashing, and make systematic wrong decisions resulting from this belief, such as choosing to drive instead. No amount of statistics changes this view, in fact most 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 perceptions of risk, then it could in fact persist.

4. *Institutional; narrowly held risks*. The most likely analysis of high average return opportunities, in my opinion, is the least stressed in academic analysis, and leads to a view of moderate persistence.

The clearest example are the catastrophe insurance enhanced bonds. 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 10-20% (depending on one's view of the chance of hurricanes); they lose money when there is a hurricane, but the risk of hurricane damage is uncorrelated with anything else. Currently, these risks are narrowly held. Before the introduction of catastrophe bonds, there was no easy way for the average investor or fund could participate in this attractive opportunity. 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.

Notice the ingredients: the risk is narrowly shared; the high average returns only disappear when the risk is widely shared (it cannot be “arbitraged 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.

The form of the institutional change required thought and experimentation. Catastrophe options were introduced on the CBOT long before the catastrophe bonds, but many funds and investors do not participate in options markets. While this kind of security would be useful if the anomaly were the kind that would be eliminated by savvy traders (options are very good devices for letting people trade on information), this anomaly requires a wide sharing of the risk.

This story gives a plausible interpretation of many of the anomalies I documented above. Small-cap stocks were found in about 1979 to provide higher returns than the 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. These allow investors to easily hold such stocks, and most fund families now feature a small-cap style fund. 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 we have seen, few funds actually followed the strategy that gives high returns 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 a value tilt in his 401 (k) plan. The risk is becoming widely shared, and its average return may fall as well.

Even 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 we have established an equilibrium as described in #1; one 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 for the 10 years or so just preceding academic publication, and for some of the stunning success of the overall market in recent years.

Inconsistent advice

Unfortunately, *the arguments that a factor will persist are inconsistent with aggressive portfolio advice.*

1. If the risk is *real*, the portfolio advice cannot apply but to half of a minuscule fraction of investors. The average investor is exposed to the risk; he knows about it but chooses to hold the market portfolio anyway because the extra return exactly compensates for the extra risk. For every investor who should want more of the risk, there must be another who wants less. Thus, the advice to load up on small-cap, value, or market-timing or other high return strategies can only apply to at most half of the investors. Additionally, if more than a minuscule fraction of investors (on either side) are not already at their best allocations already, then the market has not reached equilibrium and the premia will change.

2. 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.

3. 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't all be less behavioral than average, just as we all can't be less exposed to a risk than average. The whole argument for behavioral persistence is that the average investor would not change his portfolio, because the risk runs into fundamental ways that humans perceive risks – just as the average traveler does not quickly adjust his traveling behavior to fear the cab ride out to the airport more than the flight.

The advice *must* again 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's an irrational risk (type 2) and will disappear. If it is behavioral and will persist, then this *necessarily* comes from the fact that very few people will follow the portfolio advice! It's good salesmanship to assure each client that he is one of the few truly smart people who can take advantage of a strategy that everyone else is too dumb to follow, but we're not here to study the flattery of salespeople.

4. If the average return comes from a *narrowly held* risk, we have to ask what institutional barriers kept investors from sharing this risk more widely. Simple portfolio advice may help a bit – most investors did not, and many still do not appreciate the risk/return advantages of stocks in general, small firms, value firms, market timing strategies 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.

What's real – economic logic

The issue of why the risk gives you average return premium is also important to decide whether the opportunity is really there. Alas, stocks (and dynamic portfolio strategies) don't carry "average return" labels quite as clearly as grocery price labels. There are many statistical anomalies that vanish quickly out of sample. The vast majority of technical trading rules, seasonal anomalies, and many too-good-to-be true strategies spit out of black boxes fall in this category. 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.

5 Conclusion

Now, what do all these theories and facts mean for the investor, trying patiently to sort through the bewildering variety of available fund styles, claims for active management techniques, or temptations of minute by minute internet trading?

First, you have to figure out who you are. We start with, as before:

1. What is your overall risk tolerance?

As always, an investor must first figure out to what extent he is willing to trade off volatility for extra average returns. Investors with lower risk tolerance will want portfolios more weighted to riskfree assets, and investors with greater risk tolerance will want to load up more on risky but rewarding opportunities. 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) The overall market is about 60% stocks and 40% 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 the relevant "riskfree" asset is, as above. Longer term investors can hold longer term bonds despite their poor one-year performance, especially in a low-inflation environment.

Second, the bottom line of the above analysis is that stocks are somewhat safer for "long run" investors. They are not arbitrage opportunities – some people infer that stocks can *never* underperform bonds at a 30 or 50 year horizon, and this is not and cannot be true – but they are maybe 1/3 less volatile than the old model of

independent returns suggests.

3. What are your risks?

This is the first really new question. Are there times or states of the world in which you would particularly dislike poor performance of your investment portfolio, so much so that you would be willing to trade some average return in order to make sure that the portfolio does well in these particular times?

For example, an investor who owns a small company might properly be especially concerned that his investment portfolio does not 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 would make sense even if these stocks do “better than they should,” like the value portfolios. It would make sense for this investor to avoid an “arbitrage” hedge fund that borrows heavily to invest in low-grade bonds. Even if such a strategy did well on average it will lose money in a credit crunch or recession, just as the investor’s business is also suffering. 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 position in interest-rate sensitive or oil-price sensitive securities to offset those risks as well.

These portfolio choices work just like buying insurance. Every investor would like his “asset” portfolio to do particularly well in the event that his house burns down; and is willing to buy an investment with a poor financial return – home insurance, on which one loses money on average – in order to arrange this. We’re just extending the principle to financial markets.

Institutions also should think about what are the risks; what are the “bad events” for the institution, and what can be done to ensure that the investment portfolio does not do particularly badly in those events. (This advice is subject to the usual warning about why institutions should be making portfolio or hedging decisions at all.)

Thinking about one’s risk exposure is a new exercise. Here are some ways to phrase the questions. Are you particularly concerned about your portfolio going down,

1. When the market goes down?
2. When the market goes down a lot?
3. In recessions?

4. In credit crunch / flight to quality / times when corporate and liquidity spreads widen?
5. At times when “value” stocks do badly?
6. At times when small stocks do badly?
7. If different regions of the world do badly?
8. Times of low/high inflation, interest rates?
9. Do you own property in Florida?

The first question measures your concern about market beta of course. The second kind of question gets at whether you should be paying or collecting “put premia” or other kinds of financial catastrophe insurance premia. The third through 6th questions get at the appropriateness of these kind of styles. The seventh question opens the door to international diversification or its opposite, in order to hedge international risks. Bond risk factors matter too, as in the eighth question. The last is not facetious. If you own a lot of property in a hurricane-prone zone though, you should be selling, not buying, catastrophe-linked bonds.

Note that this logic extends not only to the kind of “factors” that have attracted academic attention but it applies to any identifiable movement in asset portfolios. For example, industry portfolios do not display much variation in (unconditional) average returns, and hence don’t show up much in new factor models. However, shorting your industry portfolio protects you against the risks of your occupation. In fact, such “factors” are even better opportunities than the “priced” factors, since you can buy insurance at zero premium. This was always true; 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, an investor may figure out what risks he does *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 if a recession is on or not. Therefore, he should buy extra amounts of recession-sensitive stocks, “value” stocks, or invest in a “convergence” type hedge fund that concentrates on illiquid assets, if these strategies carry a credible high average return. This action works just like selling insurance, in return for a premium.

Since this advice is rather new, too many investors (in my opinion at least) 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 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, and after thinking through what extra premia 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, even if hedging them costs a substantial premium.

6. Don't forget, the average investor holds the market.

If you're pretty much average, you ignore everything and still hold the market index. To the average investor, "value" stocks have a high average return, but they contain some sort of risk that he is afraid of. Maybe you should be 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 if you're average, you too are feeling wealthy and risk-tolerant. Where else are you going to put the money? To rationalize anything but the market portfolio, you have to be different than the average investor in some identifiable way.

7. And, of course, avoid all the snake oil.

Unfortunately, the marketing of many securities and funds is not particularly clear on the nature of the risks. It always sounds better to sell a fund as a profit opportunity with *no* risk. However, there is no return without risk. The economic reasoning in this article should be useful to read behind the slick marketing and 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 is still underperforming its style benchmark, and even the most aggressive estimates of persistence in fund performance are tiny – past performance has almost no information about future performance. Multifactor logic suggests passively managed style portfolios, but active management and black boxes are just as suspect as ever.

The most important piece in this portfolio advice is hidden between the lines: *avoid taxes and transactions costs*. The losses from churning a portfolio and paying needless short term capital gain, inheritance and other taxes are large. Never take a short term capital gain; use the provision that capital gains are forgiven in inheritance, delay gains, buy stocks that repurchase rather than pay dividends, watch out for funds with capital gains and so on. Tax issues are much less fun but more important to the bottom line.

Policy

An important policy theme sounds throughout this article, and it is worth remembering it in closing.

1. Asset markets are a big insurance market.

Though they often sure don't look like it! "Value funds" seem to provide extra returns to their investors by buying distressed stocks on the edge of bankruptcy. LTCM was, it seems, providing "catastrophe insurance" by intermediating liquid assets that investors want to hold into illiquid assets that were vulnerable to a liquidity crunch. Who better to provide catastrophe insurance than wildly rich investors with no other labor income or other risk exposure? Once again, we are reminded that Adam Smith's invisible hand guides apparently greedy decisions to socially useful ends, often in mysterious ways.

2. Asset market could be better insurance markets.

Surveying my portfolio advice, I am struck that the average 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 should make funds that hedge industry specific risks to labor income much more attractive vehicles.

6 References

- Barberis, Nicholas, 1999, "Investing for the Long Run when Returns are Predictable" Forthcoming, *Journal of Finance*.
- 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," Forthcoming *Journal of Finance* (October 1999)
- Brennan, Michael J., Eduardo S. Schwartz and Roland Lagnado, 1997, "Strategic Asset Allocation," *Journal of Economic Dynamics and Control* 21, 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*, 107, 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, 1998, "Who Should Buy Long Term Bonds?" Manuscript, Harvard University.
- Campbell, John Y. and Luis M. Vicera, 1999, "Consumption and Portfolio Decisions when Expected returns are Time Varying," forthcoming *Quarterly Journal of Economics*.
- 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*, 12, 255-296.
- 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.
- Daniel, Kent, David Hirshleifer and Ananidhar Subrahmanyam 1998, "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance* LIII 1839-1885.
- DeBondt, Werner F. M. and Richard H. Thaler 1985, "Does the Stock Market Overreact? *Journal of Finance* 40, 793-808
- Fama, Eugene F. 1996, "Multifactor Portfolio Efficiency and Multifactor Asset Pricing," *Journal of Financial and Quantitative Analysis*, 31, 441-465.

- Fama, Eugene F. and Kenneth R. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics* 25, 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*, 45, 141-79.
- Kandel, Shmuel and Robert Stambaugh, 1996, "On the Predictability of Stock returns: An Asset Allocation Perspective," *Journal of Finance* 51, 385-424.
- Kim, T.S. and Omberg, E., 1996, "Dynamic Nonmyopic Portfolio Behavior," *Review of Financial Studies* 9, 141-161.
- Lakonishok, Josef, Andrei Shleifer and Robert W. Vishny, 1994, "Contrarian Investment, Extrapolation and Risk," *Journal of Finance* 49, 1541-1578..
- Markowitz, H. 1952, "Portfolio Selection," *Journal of Finance*, 7, 77-99.
- 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.
- 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.

7 Appendix

7.1 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 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

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

Then the portfolio return is

$$R^p = \mathbf{w}'\mathbf{R};$$

the condition that the weights add up to 1 is

$$\mathbf{1} = \mathbf{1}'\mathbf{w}.$$

The mean of the portfolio return is

$$E(R^p) = E(\mathbf{w}'\mathbf{R}) = \mathbf{w}'E(\mathbf{R}) = \mathbf{w}'\mathbf{E}.$$

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

$$\beta^p = \mathbf{w}'\boldsymbol{\beta}$$

The variance of the portfolio return is

$$\text{var}(R^p) = \mathbf{w}'\mathbf{V}\mathbf{w}$$

where \mathbf{V} is the variance-covariance matrix of returns. The problem is then

$$\min_{\mathbf{w}} \frac{1}{2} \mathbf{w}'\mathbf{V}\mathbf{w} \quad \text{s.t.} \quad \mathbf{w}'\mathbf{E} = \mu; \mathbf{w}'\mathbf{1} = 1; \mathbf{w}'\boldsymbol{\beta} = \beta^p$$

The Lagrangian is

$$\mathcal{L} = \frac{1}{2} \mathbf{w}'\mathbf{V}\mathbf{w} - \lambda_0 (\mathbf{w}'\mathbf{E} - \mu) - \lambda_1 (\mathbf{w}'\mathbf{1} - 1) - \lambda_2 (\mathbf{w}'\boldsymbol{\beta} - \beta^p).$$

The first order conditions with respect to \mathbf{w} give

$$\begin{aligned} \mathbf{w} &= \mathbf{V}^{-1} (\mathbf{E}\lambda_0 + \mathbf{1}\lambda_1 + \boldsymbol{\beta}\lambda_2) = \mathbf{V}^{-1} \mathbf{A}\boldsymbol{\lambda} \\ \mathbf{A}'\mathbf{w} &= \boldsymbol{\delta} \end{aligned}$$

where

$$\begin{aligned} A &= \begin{bmatrix} \mathbf{E} & \mathbf{1} & \boldsymbol{\beta} \end{bmatrix} \\ \boldsymbol{\lambda} &= \begin{bmatrix} \lambda_0 & \lambda_1 & \lambda_2 \end{bmatrix}' \\ \delta &= \begin{bmatrix} \mu & 1 & \beta^p \end{bmatrix}' \end{aligned}$$

Plugging this value of \mathbf{w} into the constraint equations, we get

$$\begin{aligned} A'\mathbf{w} &= \delta \\ A'\mathbf{V}^{-1}A\boldsymbol{\lambda} &= \delta \\ \boldsymbol{\lambda} &= (A'\mathbf{V}^{-1}A)^{-1}\delta \\ \mathbf{w} &= \mathbf{V}^{-1}A(A'\mathbf{V}^{-1}A)^{-1}\delta \end{aligned}$$

The portfolio variance is then

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

If A is a second moment matrix, to handle a risk free rate,

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

Or, writing out the some of the matrix notation,

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

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.

7.2 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,

$$s^{*2} = E[s_t^2]. \tag{5}$$

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

$$\frac{E(Z)}{\sigma(Z)} \leq \frac{\sigma(m)}{E(m)} \quad (6)$$

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 easily to solve this problem in quite general situations. They phrase their result as a "lower bound on discount factor volatility" but given (6), 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 (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

$$s^* = \frac{\sqrt{s_0^2 + R^2}}{\sqrt{1 - R^2}} \quad (7)$$

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

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

The last line demonstrates equation (7). 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.

High quality versions of figure 4 and 5. Please use these rather than the versions in the text (or the accompanying eps files). Rotate as indicated by the figures in the text






Figure III: $\psi = 0.25$ and $\gamma = 4.00, 20.00$