Portfolio Theory

John H. Cochrane[1]

First draft: February 2003. This draft: February 2007

[1] Universtity of Chicago Graduate School of Business, 5807 S. Woodlawn, Chicago IL 60636, john.cochrane@chicagogsb.edu. http://faculty.chicagogsb.edu/john.cochrane/research/Papers.

Contents

1 Introduction

Now we turn to one of the classic questions of finance–portfolio theory. Given a set of available assets, i.e. given their prices and the (subjective) distribution of their payoffs, what is the optimal portfolio? This is obviously an interesting problem, to Wall Street as well as to academics.

We can also view this problem as an alternative approach to the asset pricing question. So far, we have modeled the consumption process, and then found prices from marginal utility, following Lucas' (1978) "endowment economy" logic. If you correctly model the consumption process resulting from the actual general equilibrium economy, you get the right answer for the joint distribution of consumption and asset prices from this simplified method. We could instead model the *price process*, implicitly specifying linear technologies, and derive the optimal *quantities*, i.e. optimal portfolio holdings and the consumption stream they support. This is in fact the way asset pricing was originally developed. If you correctly model the price process, you again derive the correct joint distribution of prices and quantities in this way.

I start by developing portfolio theory by the choice of *final payoff*. This is often a very easy way to approach the problem, and it ties portfolio theory directly into the $p = E(mx)$ framework of the rest of the book. Dynamic portfolio choice is, unsurprisingly, the same thing as static portfolio choice of managed portfolios, or contingent claims. I then develop the "standard approach" to portfolio theory, in which we choose the *weights* in a given set of assets, and I compare the two approaches.

2 Choosing payoffs in one-period portfolio problems

2.1 Complete markets

The investor invests, and then values consumption at a later period. We summarize prices and payoffs by a discount factor m. We solve first order conditions $u'(c) = \lambda m$ for the optimal portfolio $c = u'^{-1}(\lambda m)$. If consumption is driven by an asset payoff $\hat{x}$ and outside income e, then $\hat{x} = u'^{-1}(\lambda m) - e$. The investor sells off outside income, then invests in a portfolio driven by contingent claims prices.

Complete markets are the simplest case. Given the absence of arbitrage opportunities there is a unique, positive stochastic discount factor or contingent claims price m such that $p = E(mx)$ for any payoff x. This is a key step: rather than face the investors with prices and payoffs, we summarize the information in prices and payoff by a discount factor, and we restate the problem as the choice of how many contingent claims to buy. That summary makes the portfolio problem much easier.

Now, consider an investor with utility function over terminal consumption $E\left[u(c)\right]$, initial wealth W to invest, and random labor or business income e. The last ingredient is not common in portfolio problems, but I think it's really important, and it's easy to put it in. The business or labor income e is not directly tradeable, though there may be traded assets with similar payoffs that can be used to hedge it. In a complete market, of course, there are assets that can perfectly replicate the payoff e.

The investor's problem is to choose a portfolio. Let's call the payoff of his portfolio $\hat{x}$, so its price or value is $p(\hat{x}) = E(m\hat{x})$. He will eat $c = \hat{x} + e$. Thus, his problem is

$$\max_{\{\hat{x}\}} E\left[u(\hat{x} + e)\right] \;\; s.t. \; E(m\hat{x}) = W \tag{1}$$

$\text{Max}_{\{\hat{x}\}}$ means "choose the payoff in every state of nature. In a discrete state space, this means

$$\max_{\{\hat{x}_i\}} \sum_i \pi_i u\left(\hat{x}_i + e\right) \;\; \text{s.t.} \;\; \sum_i \pi_i m_i \hat{x}_i = W$$

This is an easy problem to solve. The first order conditions are

$$\begin{aligned} u'(c) &= \lambda m & (2) \\ u'\left(\hat{x} + e\right) &= \lambda m. & (3) \end{aligned}$$

The optimal portfolio sets marginal utility proportional to the discount factor. The optimal portfolio itself is then

$$\hat{x} = u'^{-1}(\lambda m) - e. \tag{4}$$

We find the Lagrange multiplier λ by satisfying the initial wealth constraint. Actually doing this is not very interesting at this stage, as we are more interested in how the optimal portfolio distributes across states of nature than we are in the overall level of the optimal portfolio.

Condition (4) is an old friend. The discount factor represents contingent claims prices, so condition (2) says that marginal rates of substitution should be proportional to contingent claim price ratios. The investor will consume less in high price states of nature, and consume more in low price states of nature. Risk aversion, or curvature of the utility function, determines how much the investor is willing to substitute consumption across states. Equation (4) says that the optimal *asset* portfolio $\hat{x}$ first sells off, hedges or otherwise accommodates labor income e one for one and then makes up the difference.

Condition (2) is the same first order condition we have been exploiting all along. If the investor values first period consumption c_0 as well, then the marginal utility of first period consumption equals the shadow value of wealth, $\lambda = u'(c_0)$. Adding a discount factor β for future utility, so (2) becomes our old friend

$$\beta \frac{u'(c)}{u'(c_0)} = m.$$

We didn't really need a new derivation. We are merely taking the same first order condition, and rather than fix *consumption* and solve for *prices* (and returns, etc.), we are fixing *prices* and payoffs, and solving for *consumption* and the portfolio that supports that consumption.

2.1.1 Power utility and the demand for options

For power utility $u'(c) = c^{-\gamma}$and no outside income, the return on the optimal portfolio is $\hat{R} = m^{-\frac{1}{\gamma}}/E(m^{1-\frac{1}{\gamma}})$ Using a lognormal iid stock return, this result specializes to $\hat{R} = e^{(1-\alpha)\left(r+\frac{1}{2}\alpha\sigma^2\right)} R_T^{\alpha}$ where is the stock return and $\alpha \equiv \frac{1}{\gamma}\frac{\mu-r}{\sigma^2}$. The investor wants a payoff which is a nonlinear, power function of the stock return, giving rise to demands for options.

The same method quickly extends to a utility function with a "habit" or "subsistence level", $u'(c) = (c-h)^{-\gamma}$. This example gives a strong demand for put options.

Let's try this idea out on our workhorse example, power utility. Ignoring labor income, the first order condition, equation (2), is

$$\hat{x}^{-\gamma} = \lambda m$$

so the optimal portfolio (4) is

$$\hat{x} = \lambda^{-\frac{1}{\gamma}} m^{-\frac{1}{\gamma}}$$

Using the budget constraint $W = E(m\hat{x})$ to find the multiplier,

$$\begin{aligned} W &= E(m\lambda^{-\frac{1}{\gamma}} m^{-\frac{1}{\gamma}}) \\ \lambda^{-\frac{1}{\gamma}} &= \frac{W}{E\left(m^{1-\frac{1}{\gamma}}\right)}, \end{aligned}$$

the optimal portfolio is

$$\hat{x} = W \frac{m^{-\frac{1}{\gamma}}}{E(m^{1-\frac{1}{\gamma}})}. \tag{5}$$

The $m^{-\frac{1}{\gamma}}$ term is the important one – it tells us how the portfolio $\hat{x}$ varies across states of nature. The rest just makes sure the scale is right, given this investor's initial wealth W.

In this problem, payoffs scale with wealth. This is a special property of the power utility function – richer people just buy more of the same thing. Therefore, the *return* on the optimal portfolio

$$\hat{R} = \frac{\hat{x}}{W} = \frac{m^{-\frac{1}{\gamma}}}{E(m^{1-\frac{1}{\gamma}})} \tag{6}$$

is independent of initial wealth. We often summarize portfolio problems in this way by the *return* on the optimal portfolio.

To apply this formula, we have to specify an interesting set of payoffs and their prices, and hence an interesting discount factor. Let's consider the classic Black-Scholes environment: there is a risk free bond and a single lognormally distributed stock. By allowing dynamic trading or a complete set of options, the market is "complete," at least enough for this exercise. (The next section discusses just how "complete" the market has to be.)

The stock, bond, and discount factor follow

$$\frac{dS}{S} = \mu dt + \sigma dz \tag{7}$$

$$\frac{dB}{B} = r dt \tag{8}$$

$$\frac{d\Lambda}{\Lambda} = -r dt - \frac{\mu - r}{\sigma} dz \tag{9}$$

(These are also equations (17.2) from Chapter 17, which discusses the environment in more detail. You can check quickly that this is a valid discount factor, i.e. $E(d\Lambda/\Lambda) = -rdt$ and $E(dS/S) - rdt = -E(d\Lambda/\Lambda\ dS/S)$). The discrete-time discount factor for time T payoffs is $m_T = \Lambda_T/\Lambda_0$. Solving these equations forward and with a bit of algebra below, we can evaluate Equation (6),

$$\hat{R} = e^{(1-\alpha)\left(r + \frac{1}{2}\alpha\sigma^2\right)}\ R_T^{\alpha} \tag{10}$$

where $R_T = S_T/S_0$ denotes the stock return, and

$$\alpha \equiv \frac{1}{\gamma} \frac{\mu - r}{\sigma^2}.$$

(α will turn out to be the fraction of wealth invested in stocks, if the portfolio is implemented by dynamic stock and bond trading.)

The optimal payoff is power function of the stock return. Figure 1 plots this function using standard values $\mu - r = 8\%$ and $\sigma = 16\%$ for a few values of risk aversion γ. For $\gamma = \frac{0.09-0.01}{0.16^2} = 3.125$, the function is linear – the investor just puts all his wealth in the stock. At lower levels of risk aversion, the investor exploits the strong risk-return tradeoff, taking a position that is much more sensitive to the stock return at $R_T = 1$. He gains enormous wealth if stocks go up (vertical distance past $R_T = 1$), and the cost of somewhat less consumption if stocks go down. At higher levels of risk aversion, the investor accepts drastically lower payoffs in the good states (on the right) in order to get a somewhat better payoff in the more expensive (high m) bad states on the left.

The optimal payoffs in Figure 1 are nonlinear. The investor does not just passively hold a stock and bond portfolio. Instead, he buys a complex set of contingent claims, trades dynamically, or buys a set of options, in order to create the nonlinear payoffs shown in the Figure. Fundamentally, this behavior derives from the nonlinearity of marginal utility, combined with the nonlinearity of the state-prices implied by the discount factor.

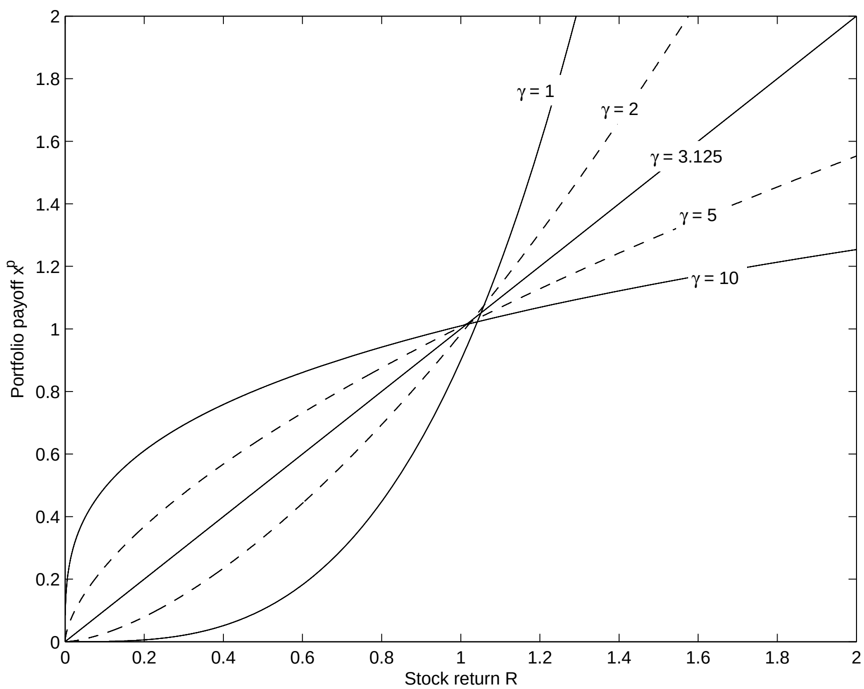

Figure 1: Return of an optimal portfolio. The investor has power utility $u\prime(c) = c^{-\gamma}$. He chooses an optimal portfolio in a complete market generated by a lognormal stock return with 9% mean and 16% standard deviation, and a 1% risk free rate.

Algebra. The solutions of the pair (7)-(9) are (see (17.5) for more detail),

$$\ln S_T = \ln S_0 + \left(\mu - \frac{\sigma^2}{2}\right) T + \sigma\sqrt{T}\varepsilon \tag{11}$$

$$\ln \Lambda_T = \ln \Lambda_0 - \left[r + \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 \right] T - \frac{\mu - r}{\sigma}\sqrt{T}\varepsilon \tag{12}$$

with $\varepsilon \tilde{} N(0,1)$. We thus have

$$m_T = \exp\left\{ - \left[r + \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 \right] T - \frac{\mu - r}{\sigma}\sqrt{T}\varepsilon \right\}$$

$$\begin{aligned}
E\left(m_T^{1-\frac{1}{\gamma}}\right) &= \exp\left[- \left(1 - \frac{1}{\gamma}\right) \left[r + \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 \right] T + \frac{1}{2}\left(1 - \frac{1}{\gamma}\right)^2 \left(\frac{\mu - r}{\sigma}\right)^2 T \right] \\
&= \exp\left\{ - \left(1 - \frac{1}{\gamma}\right) \left[r + \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 - \frac{1}{2}\left(1 - \frac{1}{\gamma}\right)\left(\frac{\mu - r}{\sigma}\right)^2 \right] T \right\} \\
&= \exp\left\{ - \left(1 - \frac{1}{\gamma}\right) \left[r + \frac{1}{2}\left(\frac{1}{\gamma}\right)\left(\frac{\mu - r}{\sigma}\right)^2 \right] T \right\}.
\end{aligned}$$

Using $R_T = S_T/S_0$ to substitute out ε in (12)

$$\begin{aligned}
m_T^{-\frac{1}{\gamma}} &= \exp\left\{ \frac{1}{\gamma}\left[r + \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 \right] T + \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}\left[\ln R_T - \left(\mu - \frac{\sigma^2}{2}\right) T \right] \right\} \\
&= \exp\left\{ \frac{1}{\gamma}\left[r + \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 - \frac{\mu - r}{\sigma^2}\left(\mu - \frac{\sigma^2}{2}\right) \right] T + \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}\ln R_T \right\} \\
&= \exp\left\{ \frac{1}{\gamma}\left[r - \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 + \left(\frac{\mu - r}{\sigma}\right)^2 - \frac{\mu - r}{\sigma^2}\left(\mu - \frac{\sigma^2}{2}\right) \right] T + \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}\ln R_T \right\} \\
&= \exp\left\{ \frac{1}{\gamma}\left[r - \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 + \frac{\mu - r}{\sigma^2}\left[\mu - r - \left(\mu - \frac{\sigma^2}{2}\right) \right] \right] T + \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}\ln R_T \right\} \\
&= \exp\left\{ \frac{1}{\gamma}\left[r - \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 - \frac{\mu - r}{\sigma^2}\left(r - \frac{\sigma^2}{2}\right) \right] T + \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}\ln R_T \right\}
\end{aligned}$$

Then

$$\begin{aligned}
\hat{R} &= \exp\left\{ \frac{1}{\gamma}\left[r - \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 - \frac{\mu - r}{\sigma^2}\left(r - \frac{\sigma^2}{2}\right) \right] T + \left(1 - \frac{1}{\gamma}\right)\left[r + \frac{1}{2}\left(\frac{1}{\gamma}\right)\left(\frac{\mu - r}{\sigma}\right)^2 \right] T \right\} \\
&\quad \times \exp\left\{ \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}\ln R_T \right\} \\
&= \exp\left\{ r - \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}\left(r - \frac{\sigma^2}{2}\right) - \frac{1}{2}\frac{1}{\gamma^2}\left(\frac{\mu - r}{\sigma}\right)^2 \right\} T \exp\left\{ \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}\ln R_T \right\} \\
&= \exp\left[r - \frac{1}{2}\sigma^2\alpha^2 - \alpha\left(r - \frac{\sigma^2}{2}\right) \right] T \times \exp\left\{ \alpha \ln R_T \right\} \\
&= \exp\left[(1 - \alpha)\left(r + \frac{1}{2}\alpha\sigma^2\right) T \right] \times R_T^{\alpha}
\end{aligned}$$

Implementation

This example will still feel empty to someone who knows standard portfolio theory, in which the maximization is stated over portfolio shares of specific assets rather than over the final payoff. Sure, we have characterized the optimal *payoffs*, but weren't we supposed to be finding optimal *portfolios*? What stocks, bonds or options does this investor actually hold?

Figure 1 *does* give portfolios. We are in a complete market. Figure 1 gives the number of contingent claims to each state, indexed by the stock return, that the investor should buy. In a sense, we have made the portfolio problem very easy by cleverly choosing a simple basis – contingent claims – for our complete market.

There remains a largely technical question: suppose you wanted to implement this pattern of contingent claims by explicitly buying standard put and call options, or by dynamic trading in a stock or bond, or any of the infinite number of equivalent repackaging of securities that span the complete market, rather than by explicitly buying contingent claims. How would you do it? I'll return to these questions below, and you'll see that they involve a lot more algebra. But really, they are technical questions. We've solved the important *economic* question, what the optimal payoff should *be*. Ideally, in fact, an intermediary (investment bank) would handle the financial engineering of generating most cheaply the payoff shown in Figure 1, and simply sell the optimal payoff directly as a retail product.

That said, there are two obvious ways to approximate payoffs like those Figure 1. First, we can approximate nonlinear functions by a series of linear functions. The low risk aversion ($\gamma = 1, \gamma = 2$) payoffs can be replicated by buying a series of call options, or by holding the stock and writing puts. The high risk aversion ($\gamma = 5, \gamma = 10$) payoffs can be replicated by writing call options, or by holding the stock and buying put options. The put options provide "portfolio insurance." Thus we see the demand and supply for *options* emerge from different attitudes towards risk. In fact many investors *do* explicitly buy put options to protect against "downside risk," while many hedge funds do, explicitly or implicitly, write put options.

Second, one can trade dynamically. In fact, as I will show below, the standard approach to this portfolio problem does not mention options at all, so one may wonder how I got options in here. But the standard approach leads to portfolios that are continually rebalanced. As it turns out, this payoff can be achieved by continually rebalancing a portfolio with α fraction of wealth held in stock. If you hold, say $\alpha = 60\%$ stocks and 40% bonds, then as the market goes up you will sell some stocks. This action leaves you less exposed to further stock market increases than you would otherwise be, and leads to the concave ($\gamma > 3.125$) discrete-period behavior shown in the graph.

Habits

A second example is useful to show some of the power of the method, and that it really can be applied past standard toy examples. Suppose the utility function is changed to have a subsistence or minimum level of consumption h,

$$u(c) = (c - h)^{1-\gamma}.$$

Now, the optimal payoff is

$$
\begin{aligned}
(\hat{x}-h)^{-\gamma} &= \lambda m \\
\hat{x} &= \lambda^{-\frac{1}{\gamma}} m^{-\frac{1}{\gamma}} + h
\end{aligned}
$$

Evaluating the wealth constraint,

$$
\begin{aligned}
W_0 &= E(m\hat{x}) = \lambda^{-\frac{1}{\gamma}} E\left(m^{1-\frac{1}{\gamma}}\right) + he^{-rT} \\
\lambda^{-\frac{1}{\gamma}} &= \frac{W_0 - he^{-rT}}{E\left(m^{1-\frac{1}{\gamma}}\right)} \\
\hat{x} &= \left(W_0 - he^{-rT}\right) \frac{m^{-\frac{1}{\gamma}}}{E\left(m^{1-\frac{1}{\gamma}}\right)} + h
\end{aligned}
$$

The *discount factor* has not changed, so we can use the discount factor terms from the last example unchanged. In the lognormal Black-Scholes example we have been carrying along, this result gives us, corresponding to (10),

$$
\hat{x} = \left(W_0 - he^{-rT}\right) e^{(1-\alpha)\left(r+\frac{1}{2}\alpha\sigma^2\right)T} R_T^{\alpha} + h
$$

This is a very sensible answer. First and foremost, the investor guarantees the payoff h. Then, wealth left over after buying a bond that guarantees h, $\left(W_0 - he^{-rT}\right)$ is invested in the usual power utility manner. Figure 2 plots the payoffs of the optimal portfolios. You can see the left end is higher and the right end is lower. The investor sells off some performance in good states of the world to make sure his portfolio never pays off less than h no matter how bad the state of the world.

2.2 Incomplete markets

> Most of the complete markets approach goes through in incomplete markets as well. The first order condition $\hat{x} = u'^{-1}(\lambda m) - e$ still gives the optimal portfolio, but in general there are many m and we don't know which one lands $\hat{x} \in \underline{X}$, the space available to the investor.

Well, what if markets are *not* complete? This is the case in the real world. Market incompleteness is also what makes portfolio theory challenging. So, let's generalize the ideas of the last section to incomplete markets.

When markets are incomplete, we have to be more careful about what actually is available to the investor. I start with a quick review of the setup and central results from Chapter

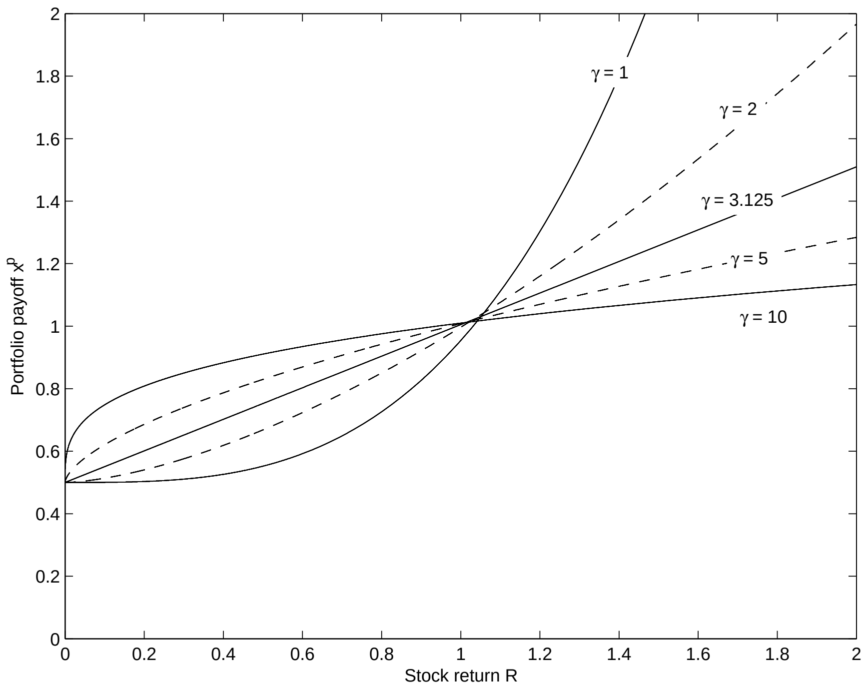

Figure 2: Portfolio problem with habit utility

4. The payoffs available to the investor are a space $\underline{X}$. For each payoff $x \in \underline{X}$ the investor knows the price $p(x)$. Returns have price 1, excess returns have price zero. The investor can form arbitrary portfolios without short-sale constraints or transactions costs (that's another interesting extension), so the space $\underline{X}$ of payoffs is closed under linear transformations:

$$x \in \underline{X}, y \in \underline{X} \Rightarrow ax + by \in \underline{X}$$

I assume that the law of one price holds, so the price of a portfolio is the same as the price of its constituent elements.

$$p(ax + by) = ap(x) + bp(y).$$

As before, let's follow the insight that summarizing prices and payoffs with a discount factor makes the portfolio theory problem easier. From Chapter 4, we know that the law of one price implies that there is a unique discount $x^* \in \underline{X}$ such that

$$p(x) = E(x^* x) \tag{13}$$

for all $x \in \underline{X}$. The discount factor x^* is often easy to construct. For example, if the payoff space is generated as all portfolios of a finite vector of basis payoffs $\mathbf{x}$ with price vector $\mathbf{p}$, $\underline{X} = \{\mathbf{c}'\mathbf{x}\}$, then $x^* = \mathbf{p}' E(\mathbf{x}\mathbf{x}')^{-1}\mathbf{x}$ satisfies $\mathbf{p} = E(x^*\mathbf{x})$ and $x^* \in \underline{X}$. Equation (9) is a continuous-time version of this equation.

If markets are complete, this is the unique discount factor. If markets are not complete, then there are many discount factors and any $m = x^* + \varepsilon$, with $E(\varepsilon x) = 0\ \forall x \in \underline{X}$ is a discount factor. Therefore, $x^* = proj(m|\underline{X})$ for any discount factor m. The return corresponding to the payoff $x^*$is $R^* = x^*/p(x^*) = x^*/E(x^{*2})$. R^* is the global minimum second moment return, and so it is on the lower portion of the mean-variance frontier. x^* and R^* need not be positive in every state of nature. Absence of arbitrage means there *exists* a positive discount factor $m = x^* + \varepsilon$, but the positive m may not lie in $\underline{X}$, and there are many non-positive discount factors as well.

The canonical one-period portfolio problem is now

$$\begin{aligned} & \max_{\{\hat{x}\in\underline{X}\}} E\left[u(c)\right] \text{ s.t.} \\ c &= \hat{x} + e; \;\; W = p(\hat{x}). \end{aligned} \tag{14}$$

This is different from our first problem (1) only by the restriction $\hat{x} \in \underline{X}$: markets are incomplete, and the investor can only choose a tradeable payoff.

The first order conditions are the same as before. We can see this most transparently in the common case of a finite set of basis payoffs $\underline{X} = \{\mathbf{c}'\mathbf{x}\}$. Then, the constrained portfolio choice is $\hat{x} = \boldsymbol{\alpha}'\mathbf{x}$ and we can choose the portfolio weights $\boldsymbol{\alpha}$, respecting in this way $\hat{x} \in \underline{X}$. The portfolio problem is

$$\max_{\{\alpha_i\}} E\left[u\left(\sum_i \alpha_i x_i + e\right)\right] \text{ s.t. } \; W = \sum_i \alpha_i p_i.$$

The first order conditions are

$$\frac{\partial}{\partial \alpha_i}: \quad p_i\lambda = E\left[u'(\hat{x}+e)x_i\right] \tag{15}$$

for each asset i, where λ is the Lagrange multiplier on the wealth constraint.

Equation (15) is our old friend $p = E(mx)$. It holds for each asset in $\underline{X}$ if and only if $u'(\hat{x}+e)/\lambda$ is a discount factor for all payoffs $\hat{x} \in \underline{X}$. We conclude that *marginal utility must be proportional to* **a** *discount factor*,

$$u'(\hat{x}+e) = \lambda m \tag{16}$$

where m satisfies $p = E(mx)$ for all $x \in \underline{X}$.

We can also apply the same derivation as before. It's prettier, but the logic is a little trickier. We know from the law of one price that there exists an m such that $p = E(mx) \forall x \in \underline{X}$, in fact there are lots of them. Thus, we can state the constraint as $W = E(m\hat{x})$ using any such m. Now the problem (14) is exactly the same as the original problem, so we can find the first order condition by choosing $\hat{x}$ in each state directly, with no mention of the original prices and payoffs.

The solution to the portfolio problem is thus once again

$$\hat{x} = u'^{-1}(\lambda m) - e.$$

If markets are complete, as above, the discount factor $m = x^*$ is unique and in $\underline{X}$. Every payoff is traded, so both λm and $u'^{-1}(\lambda m) - e$ are in $\underline{X}$. Hence, all we have to do is find the Lagrange multiplier to satisfy the initial wealth constraint.

If markets are not complete, we also have to pay attention to the constraint $\hat{x} \in \underline{X}$. We have derived *necessary* condition for an optimal portfolio, but not yet a *sufficient* condition. There are *many* discount factors that price assets, and for only *one* of them is the inverse marginal utility in the space of traded assets. While it's easy to construct $x^* \in \underline{X}$, for example, that may be the wrong discount factor.

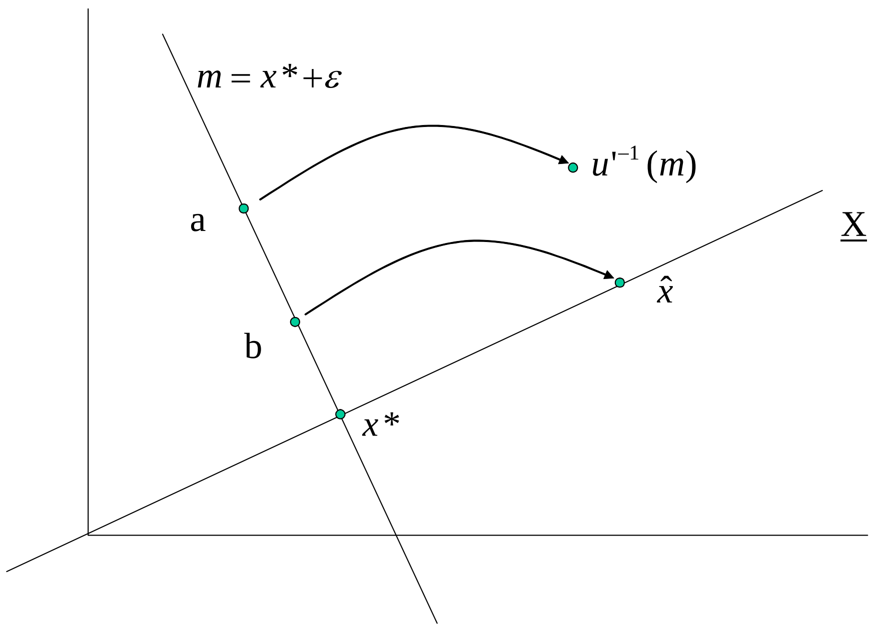

Figure 3: Portfolio problem in incomplete markets.

Figure 3 illustrates the problem for the case $e = 0$. $\underline{X}$ is the space of traded payoffs. x^* is the unique discount factor in $\underline{X}$. $m = x^* + \varepsilon$ gives the space of all discount factors. It is drawn at right angles to $\underline{X}$ since $E(\varepsilon x) = 0\ \forall x \in \underline{X}$. The optimal portfolio x^* satisfies $u'^{-1}(\lambda m) = x^*$ for some m. Case a shows what can go wrong if you pick the wrong m: $u'^{-1}(\lambda m)$ is not in the payoff space $\underline{X}$, so it can't be the optimal portfolio. Case b shows the optimal portfolio: we have chosen the right m so that $u'^{-1}(\lambda m)$ is in the payoff space $\underline{X}$. x^* is the wrong choice as well, since $u'^{-1}(\lambda x^*)$ takes you out of the payoff space.

As the figure suggests, markets don't *have to* be completely "complete" for $\hat{x} = u'^{-1}(\lambda x^*)$ to work. It is enough that the payoff space $\underline{X}$ is closed under (some) nonlinear transformations. If for every $x \in \underline{X}$, we also have $u'^{-1}(x) \in \underline{X}$, then $\hat{x} = u'^{-1}(\lambda x^*)$ will be tradeable, and we can again find the optimal portfolio by inverting marginal utility from the easy-to-compute unique discount factor $x^* \in \underline{X}$. A full set of options gives closure under *all* nonlinear transformations and this situation is often referred to as "complete markets," even though many shocks are not traded assets. Obviously, even less "completeness" than this can work in many applications.

What can we do? How can we pick the right m? In general, there are two ways to proceed. First, we can search over all possible m, i.e. slide up and down the m hyperplane and look for the one that sends $u'^{-1}(\lambda m) \in \underline{X}$. This isn't necessarily hard, since we can set up the search as a minimization, minimizing the distance between $u'^{-1}(m)$ and $\underline{X}$. Equivalently, we can invent prices for the missing securities, solve the (now unique) complete markets problem, and search over those prices until the optimal portfolio just happens to lie in the original space $\underline{X}$. Equivalently again, we can attach Lagrange multipliers to the constraint $\hat{x} \in \underline{X}$ and find "shadow prices" that satisfy the constraints.

Second, we can start all over again by explicitly choosing portfolio weights directly in the limited set of assets at hand. This approach also leads to a lot of complexity. In addition, in most applications there are a lot more assets at hand than can conveniently be put in a maximization. For example, we often do portfolio problems with a stock and a bond, ignoring the presence of options and futures. In almost all cases of practical importance, we have to result to numerical answers, which means some approximation scheme.

Third, we can simplify or approximate the *problem*, so that $u'^{-1}(\cdot)$ is an easy function.

2.3 Linear-quadratic approximation and mean-variance analysis

---

If marginal utility is *linear*, $u'(c) = c^b - c$, then we can easily solve for portfolios in incomplete markets. I derive $\hat{x} = \hat{c}^b - \hat{e} - \left[p(\hat{c}^b) - p(\hat{e}) - W\right] R^*$, where $\hat{c}^b$ and $\hat{e}$ are mimicking payoffs for a stochastic bliss point and outside income, W is initial wealth, and R^* is the minimum second moment return. The portfolio gets the investor as close as possible to bliss point consumption, after hedging outside income risk, and then accepting lower consumption in the high contingent claims price states.

We know how to find a *discount factor* in the payoff space $\underline{X}$, namely x^*. The problem is that marginal utility is nonlinear, while the payoff space $\underline{X}$ is only closed under linear combinations. This suggests a classic approximation: With quadratic utility, marginal utility is linear. Then we know that the inverse image of $x^* \in \underline{X}$ is also in the space $\underline{X}$ of payoffs, and this is the optimal portfolio.

Analytically, suppose utility is quadratic

$$u(c) = -\frac{1}{2}(c^b - c)^2$$

where c^b is a potentially stochastic bliss point. Then

$$u'(c) = c^b - c.$$

The first order condition (16) now reads

$$c^b - \hat{x} - e = \lambda m.$$

Now, we can project both sides onto the payoff space $\underline{X}$, and solve for the optimal portfolio. Since $proj(m|\underline{X}) = x^*$, this operation yields

$$\hat{x} = -\lambda x^* + proj\left(c^b - e|\underline{X}\right). \tag{17}$$

To make the result clearer, we again solve for the Lagrange multiplier λ in terms of initial wealth. Denote by $\hat{e}$ and $\hat{c}^b$ the mimicking portfolios for preference shocks and labor income risk,

$$\begin{aligned} \hat{e} &\equiv proj\left(e|\underline{X}\right) \\ \hat{c}^b &\equiv proj\left(c^b|\underline{X}\right) \end{aligned}$$

(Projection means linear regression. These are the portfolios of asset payoffs that are closest, in mean square sense, to the labor income and bliss points.) The wealth constraint then states

$$W = p(\hat{x}) = -\lambda p(x^*) + p(\hat{c}^b) - p(\hat{e})$$

$$\frac{p(\hat{c}^b) - p(\hat{e}) - W}{p(x^*)} = \lambda$$

Thus, the optimal portfolio is

$$\hat{x} = \hat{c}^b - \hat{e} - \left[p(\hat{c}^b) - p(\hat{e}) - W\right] R^*, \tag{18}$$

where again $R^* = x^*/p(x^*) = x^*/E(x^{*2})$ is the return corresponding to the discount-factor payoff x^*.

The investor starts by hedging as much of his preference shock and labor income risk as possible. If these risks are traded, he will buy a portfolio that gets rid of all labor income risk e and then buys bliss point consumption c^b. If they are not traded, he will buy a portfolio that is closest to this ideal – a mimicking portfolio for labor income and preference shock risk. Then, depending on initial wealth and hence risk aversion (risk aversion depends on wealth for quadratic utility), he invests in the minimum second moment return R^*. Typically (for all interesting cases) wealth is not sufficient to buy bliss point consumption, $W + p(\hat{e}) < p(\hat{c}^b)$. Therefore, the investment in R^* is negative. R^* is on the lower portion of the mean-variance frontier, so when you short R^*, you obtain a portfolio on the upper portion of the frontier. The investment in the risky portfolio is larger (in absolute value) for lower wealth. Quadratic utility has the perverse feature that risk aversion increases with wealth to infinity at the bliss point. Given that the investor cannot buy enough assets to consume $\hat{c}^b$, R^* tells him which states have the highest contingent claims prices. Obviously, sell what you have at the highest price.

In sum, *each investor's optimal portfolio is a combination of a mimicking portfolio to hedge labor income and preference shock risk, plus an investment in the (mean-variance efficient) minimum second moment return, whose size depends on risk aversion or initial wealth.*

2.3.1 The mean-variance frontier

With no outside income $e = 0$, we can express the quadratic utility portfolio problem in terms of local risk aversion,

$$\hat{R} = R^f + \frac{1}{\gamma}\left(R^f - R^*\right).$$

This expression makes it clear that the investor holds a mean-variance efficient portfolio, further away from the risk free rate as risk aversion declines.

Traditional mean-variance analysis focuses on a special case: the investor has no job, so labor income is zero, the bliss point is nonstochastic, and a riskfree rate is traded. This special case leads to a very simple characterization of the optimal portfolio. Equation (18) specializes to

$$\hat{x} = c^b - \left(\frac{c^b}{R^f} - W\right) R^* \tag{19}$$

$$\hat{R} = \frac{\hat{x}}{W} = R^* + \frac{c^b}{R^f W}\left(R^f - R^*\right) \tag{20}$$

In Chapter 5, we showed that the mean-variance frontier is composed of all portfolios of the form $R^* + \alpha(R^f - R^*)$. Therefore, *investors with quadratic utility and no labor income all hold mean-variance efficient portfolios.* As W rises or c^b declines, the investor becomes more risk averse. When W can finance bliss-point consumption for sure, $WR^f = c^b$, the investor becomes infinitely risk averse and holds only the riskfree rate R^f.

Obviously, these global implications – rising risk aversion with wealth – are perverse features of quadratic utility, which should be thought of as a local approximation. For this reason, it is interesting and useful to express the portfolio decision in terms of the local risk aversion coefficient.

Write (20) as

$$\hat{R} = R^f + \left(\frac{c^b}{R^f W} - 1\right)\left(R^f - R^*\right) \tag{21}$$

Local risk aversion for quadratic utility is

$$\gamma = -\frac{cu''(c)}{u'(c)} = \frac{c}{c^b - c} = \left(\frac{c^b}{c} - 1\right)^{-1}.$$

Now we can write the optimal portfolio

$$\hat{R} = R^f + \frac{1}{\gamma}\left(R^f - R^*\right). \tag{22}$$

where we evaluate local risk aversion γ at the point $c = R^f W$.

The investor invests in a mean-variance efficient portfolio, with larger investment in the risky asset the lower his risk aversion coefficient. Again, R^* is on the lower part of the mean-variance frontier, thus a short position in R^* generates portfolios on the upper portion of the frontier. $R^f W$ is the amount of consumption the investor would have in period 2 if he invested everything in the risk free asset. This is the sensible place at which to evaluate risk aversion. For example, if you had enough wealth to buy bliss point consumption $R^f W = c^b$, you would do it and be infinitely risk averse.

2.3.2 Formulas

I evaluate the mean-variance formula $\hat{R} = R^f + \frac{1}{\gamma}\left(R^f - R^*\right)$ for the common case of a riskfree rate R^f and vector of excess returns R^e with mean μ and covariance matrix Σ. The result is

$$R^f - R^* = \left(\frac{R^f}{1 + \mu'\Sigma^{-1}\mu}\right)\mu'\Sigma^{-1}R^e$$

The terms are familiar from simple mean-variance maximization: finding the mean-variance frontier directly we find that mean-variance efficient weights are all of the form $w = \lambda\mu'\Sigma^{-1}$ and the maximum Sharpe ratio is $\mu'\Sigma^{-1}\mu$.

Formula (22) is a little dry, so it's worth evaluating a common instance. Suppose the payoff space consists of a riskfree rate R^f and N assets with excess returns R^e, so that portfolio returns are all of the form $R^p = R^f + w'R^e$. Denote $\mu = E(R^e)$ and $\Sigma = cov(R^e)$. Let's find R^* and hence (22) in this environment.

Repeating briefly the analysis of Chapter 5, we can find

$$x^* = \frac{1}{R^f} - \frac{1}{R^f}\mu'\Sigma^{-1}(R^e - \mu).$$

(Check that $x^* \in \underline{X}$, $E(x^* R^f) = 1$ and $E(x^* R^e) = 0$, or derive it from $x^* = \alpha R^f + w'\left[R^e - \mu\right]$.) Then

$$p(x^*) = E(x^{*2}) = \frac{1}{R^{f2}} + \frac{1}{R^{f2}}\mu'\Sigma^{-1}\mu$$

so

$$R^* = \frac{x^*}{E(x^{*2})} = R^f \frac{1 - \mu'\Sigma^{-1}(R^e - \mu)}{1 + \mu'\Sigma^{-1}\mu}$$

and

$$\begin{aligned} R^f - R^* &= R^f - \frac{R^f}{1 + \mu'\Sigma^{-1}\mu} + \frac{R^f}{1 + \mu'\Sigma^{-1}\mu}\mu'\Sigma^{-1}(R^e - \mu) \\ R^f - R^* &= \frac{R^f}{1 + \mu'\Sigma^{-1}\mu}\mu'\Sigma^{-1}R^e \qquad (23) \end{aligned}$$

To give a reference for these formulas, consider the standard approach to finding the mean-variance frontier. Let R^{ep} be the excess return on a portfolio. Then we want to find

$$\begin{aligned}\min \sigma^2(R^{ep}) \text{ s.t. } E(R^{ep}) &= E\\ \min_{\{w\}} w'\Sigma w \text{ s.t. } w'\mu &= E\end{aligned}$$

The first order conditions give

$$w = \lambda \Sigma^{-1}\mu$$

where λ scales up and down the investment to give larger or smaller mean. Thus, the portfolios on the mean-variance frontier have excess returns of the form

$$R^{ep} = \lambda \mu'\Sigma^{-1}R^e$$

This is a great formula to remember: $\mu'\Sigma^{-1}$ gives the weights for a mean-variance efficient investment. You can see that (23) is of this form.

The Sharpe ratio or slope of the mean-variance frontier is

$$\frac{E(R^{ep})}{\sigma(R^{ep})} = \frac{\mu'\Sigma^{-1}\mu}{\sqrt{\mu'\Sigma^{-1}\mu}} = \sqrt{\mu'\Sigma^{-1}\mu}$$

Thus, you can see that the term scaling $R^f - R^*$ scales with the market Sharpe ratio.

We could of course generate the mean-variance frontier from the risk free rate and any efficient return. For example, just using $\mu'\Sigma^{-1}R^e$ might seem simpler than using (23), and it is simpler when making computations. The *particular* mean-variance efficient portfolio $R^f - R^*$ in (23) has the delicious property that it is the optimal portfolio for risk aversion equal to one, and the units of any investment have directly the interpretation as a risk aversion coefficient.

2.3.3 The market portfolio and two-fund theorem

In a market of quadratic utility, $e = 0$ investors, we can aggregate across people and express the optimal portfolio as

$$\hat{R}^i = R^f + \frac{\gamma^m}{\gamma^i}\left(R^m - R^f\right)$$

This is a "two-fund" theorem – the optimal portfolio for every investor is a combination of the risk free rate and the market portfolio. Investors hold more or less of the market portfolio according to whether they are less or more risk averse than the average investor.

I have used R^* so far as the risky portfolio. If you read Chapter 5, this will be natural. However, conventional mean-variance analysis uses the "market portfolio" on the top of the mean variance frontier as the reference risky return. It's worth developing this representation and the intuition that goes with it.

Write the portfolio choice of individual i from (21) as

$$\hat{R}^i = R^f + \frac{1}{\gamma^i}\left(R^f - R^*\right). \tag{24}$$

The market portfolio $\hat{R}^m$ is the wealth-weighted average of individual portfolios, or the return on the sum of individual payoffs,

$$\hat{R}^m \equiv \frac{\sum_{i=1}^N \hat{x}^i}{\sum_{j=1}^N W^j} = \frac{\sum_{i=1}^N W^i \hat{R}^i}{\sum_{j=1}^N W^j}$$

Summing (24) over individuals, then,

$$\hat{R}^m = R^f + \frac{\sum_{i=1}^N W^i \frac{1}{\gamma^i}}{\sum_{j=1}^N W^j}\left(R^f - R^*\right).$$

We can define an "average risk aversion coefficient" as the wealth-weighted average of (inverse) risk aversion coefficients[1],

$$\frac{1}{\gamma^m} \equiv \frac{\sum_{i=1}^N W^i \frac{1}{\gamma^i}}{\sum_{j=1}^N W^j}$$

so

$$\hat{R}^m = R^f + \frac{1}{\gamma^m}\left(R^f - R^*\right).$$

Using this relation to substitute $R^m - R^f$ in place of $R^f - R^*$ in (24), we obtain

$$\hat{R}^i = R^f + \frac{\gamma^m}{\gamma^i}\left(R^m - R^f\right) \tag{25}$$

The optimal portfolio is split between the risk free rate and the market portfolio. The weight on the market portfolio return depends on individual relative to average risk aversion.

The "market portfolio" here is the average of *all* assets held. If there are bonds in "net supply" then they are included in the market portfolio, and the remaining riskfree rate is in "zero net supply." Since $x^i = c^i$, the market portfolio is also the claim to total consumption.

[1] "Market risk aversion" is also the local risk aversion of an investor with the average blisspoint and average wealth,

$$\frac{1}{\gamma^m} = \frac{\frac{1}{N}\sum_{i=1}^N c^{bi}}{R^f \frac{1}{N}\sum_{i=1}^N W^i} - 1.$$

Since any two mean-variance efficient portfolios span the frontier, R^m and R^f for example, we see that optimal portfolios follow a *two-fund theorem*. This is very famous in the history of finance. It was once taken for granted that each individual needed a tailored portfolio, riskier stocks for less risk averse investors. Investment companies still advertise how well they listen. In this theory, the only way people differ is by their risk aversion, so all investors' portfolios can be provided by two funds, a "market portfolio" and a risk free security.

This is all illustrated in the classic mean-variance frontier diagram, Figure 4.

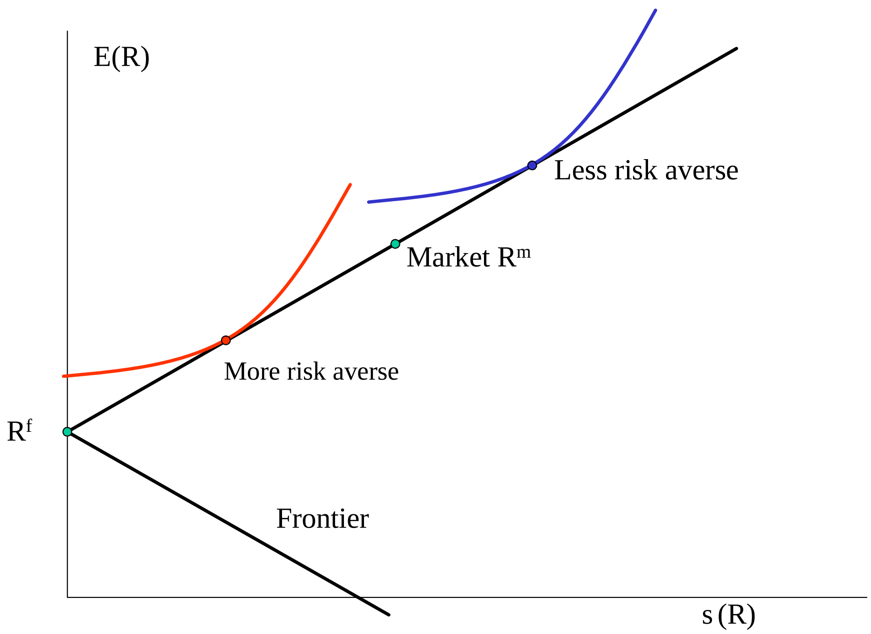

Figure 4: Mean-variance efficient portfolios.

2.4 Nontradeable income

I introduce two ways of expressing mean-variance portfolio theory with outside income. First, the *overall* portfolio, including the hedge portfolio for outside income, is still on the mean-variance frontier. Thus, we could use classic analysis to determine the right overall portfolio – keeping in mind that the overall market portfolio includes hedge portfolios for the average investors outside income too – and then subtract off the hedge portfolio for individual outside income in order to arrive at the individual's asset portfolio. Second, we can express the individual's portfolio as 1) the market *asset* portfolio, adjusted for risk aversion and the composition of wealth, 2) the average outside income hedge portfolio for all other investors, adjusted again for risk aversion and wealth and finally 3) the hedge portfolio for the individual's idiosyncratic outside income.

The mean-variance frontier is a beautiful and classic result, but most investors do in fact have jobs, business income or real estate. Here, I attempt some restatements of the more interesting case with labor income and preference shocks to bring them closer to mean-variance intuition.

One way to do this is to think of labor or business income as part of a "total portfolio". Then, the total portfolio *is* still mean-variance efficient, but we have to adjust the asset portfolio for the presence of outside income.

To keep it simple, keep a nonstochastic bliss point, c^b. Then, equation (18) becomes

$$\hat{x} = c^b - \hat{e} - \left[p(c^b) - p(\hat{e}) - W\right] R^*$$

We can rewrite this as

$$\hat{e} + \hat{x} = c^b - \left[p(c^b) - (W + p(\hat{e}))\right] R^*$$

The left hand side is the "total payoff", consisting of the asset payoff $\hat{x}$ and the labor income hedge portfolio $\hat{e}$ (Consumption is this payoff plus residual labor income, $c = \hat{x} + e = \hat{x} + (e - \hat{e}) + \hat{e}$.)

We define a rate of return on the "total portfolio" as the total payoff – asset portfolio plus human capital – divided by total value, and proceed as before,

$$\begin{aligned}
\hat{R}^{tp} &= \frac{\hat{e} + \hat{x}}{W + p(\hat{e})} = \frac{c^b}{W + p(\hat{e})} - \left[\frac{c^b}{R^f \left[W + p(\hat{e})\right]} - 1\right] R^* \\
\hat{R}^{tp} &= R^f + \frac{1}{\gamma}\left(R^f - R^*\right)
\end{aligned}$$

Now

$$\frac{1}{\gamma} = \frac{c^b}{R^f \left[W + p(\hat{e})\right]} - 1$$

is defined as the local risk aversion coefficient given c^b and using the value of initial wealth and the tradeable portfolio closest to labor income, invested at the risk free rate. Thus, we can say that the *total portfolio is mean-variance efficient.* We can also aggregate just as before, to express

$$\hat{R}^{tp,i} = R^f + \frac{\gamma^m}{\gamma^i}\left(R^{tp,m} - R^f\right) \tag{26}$$

where R^m is now the total wealth portfolio *including* the outside income portfolios,

This representation makes it seem like nothing much has changed. However the *asset* portfolio – the thing the investor actually buys – changes dramatically. $\hat{e}$ is a payoff the investor already owns. Thus, to figure out the *asset* market payoff, you have to *subtract* the labor income hedge portfolio from the appropriate mean-variance efficient portfolio.

$$\hat{R}^i = \frac{\hat{x}^i}{W^i} = \frac{p(\hat{e}^i) + W^i}{W^i}\left(\frac{\hat{e}^i + \hat{x}^i}{p(\hat{e}^i) + W^i} - \frac{\hat{e}^i}{p(\hat{e}^i) + W^i}\right) \tag{27}$$

$$= \left(1 + \frac{p(\hat{e}^i)}{W^i}\right)\hat{R}^{tp,i} - \left(\frac{p(\hat{e}^i)}{W^i}\right)\frac{\hat{e}^i}{p(\hat{e}^i)} \tag{28}$$

$$= \left(1 + \frac{p(\hat{e}^i)}{W^i}\right)\hat{R}^{tp,i} - \left(\frac{p(\hat{e}^i)}{W^i}\right)\hat{R}^{z,i} \tag{29}$$

where I use

$$\hat{R}^{z,i} = \frac{\hat{e}^i}{p(\hat{e}^i)}$$

to denote the return on the mimicking portfolio for outside income. (I can't use the natural notation $\hat{R}^{e,i}$ since R^e stands for excess return.) This can be a large correction. Also, in this representation the corresponding "market portfolio" $\hat{R}^{tp,m}$ includes everyone else's hedge portfolio. It is *not* the average of actual asset market portfolios.

For that reason, I prefer a slightly more complex representation. We can break up the "total" return to the two components, a "mimicking portfolio return" and the asset portfolio return,

$$\begin{aligned}
\hat{R}^{tp,i} &= \frac{\hat{e}^i + \hat{x}^i}{p(\hat{e}^i) + W^i} \\
&= \frac{p(\hat{e}^i)}{p(\hat{e}^i) + W^i}\frac{\hat{e}^i}{p(\hat{e}^i)} + \frac{W^i}{p(\hat{e}^i) + W^i}\frac{\hat{x}^i}{W^i} \\
&= (1 - w^i)\hat{R}^{z,i} + w^i\hat{R}^i
\end{aligned}$$

Here

$$\hat{R}^{z,i} = \frac{\hat{e}^i}{p(\hat{e}^i)};\ w^i = \frac{W^i}{p(\hat{e}^i) + W^i}; 1 - w^i = \frac{p(\hat{e}^i)}{p(\hat{e}^i) + W^i}.$$

The same decomposition works for $R^{tp,m}$. Then, substituting in (26),

$$(1 - w^i)\hat{R}^{z,i} + w^i\hat{R}^i = R^f + \frac{\gamma^m}{\gamma^i}\left((1 - w^m)\hat{R}^{z,m} + w^m\hat{R}^m - R^f\right)$$

and hence

$$\hat{R}^i - R^f = \frac{\gamma^m}{\gamma^i}\frac{w^m}{w^i}\left(\hat{R}^m - R^f\right) + \frac{\gamma^m}{\gamma^i}\frac{w^m}{w^i}\frac{(1 - w^m)}{w^m}\left(\hat{R}^{z,m} - R^f\right) - \frac{(1 - w^i)}{w^i}\left(\hat{R}^{z,i} - R^f\right) \tag{30}$$

This representation emphasizes a deep point, you only deviate from the market portfolio to the extent that you are *different* from everyone else. The first term says that an individual's actual portfolio scales up or down the market portfolio according to the individual's risk aversion and the relative weight of asset wealth in total wealth. If you have more outside wealth relative to total, w^i is lower, you hold a less risk averse position in your *asset* portfolio. The second term is the hedge portfolio for the average investor's labor income The next terms describe how you should change your portfolio if the character of your outside income is *different* from everyone else's.

A few examples will clarify the formula. First, suppose that outside income is nonstochastic, so $R^z = R^f$. The second terms vanish, and we are left with

$$\hat{R}^i - R^f = \frac{\gamma^m}{\gamma^i}\frac{w^m}{w^i}\left(\hat{R}^m - R^f\right)$$

This is the usual formula except that risk aversion is now multiplied by the share of asset wealth in total wealth,

$$\gamma^i w^i = \gamma^i \frac{W^i}{W^i + p(e^i)}.$$

An individual with a lot of outside income $p(e^i)$ is sitting on a bond. Therefore, his asset market portfolio should be shifted towards risky assets; his asset market portfolio is the same as that of an investor with no outside income but a lot less risk aversion. This explains why "effective risk aversion" for the asset market portfolio is in (30) multiplied by wealth.

Second, suppose that the investor has the same wealth and relative wealth as the market, $\gamma^i = \gamma^m$ and $w^i = w^m$, but outside income is stochastic. Then expression (30) simplifies to

$$\hat{R}^i - R^f = \left(\hat{R}^m - R^f\right) + \frac{(1-w)}{w}\left[\hat{R}^{z,m} - \hat{R}^{e,i}\right]$$

This investor holds the market portfolio (this time the actual, traded-asset market portfolio), plus a hedge portfolio derived from the *difference* between his income and the average investor's income. If the investor is just like the average investor in this respect as well, then he just holds the market portfolio of traded assets. But suppose this investor's outside income is a bond, $\hat{R}^{ej} = R^f$, while the average investor has a stochastic outside income. Then, the investor's *asset* portfolio will include the hedge portfolio for aggregate outside income. He will do better in a mean-variance sense by providing this "outside income insurance" to the average investor.

2.5 The CAPM, multifactor models, and four fund theorem

3 Choosing payoffs in intertemporal problems

One-period problems are fun and pedagogically attractive, but not realistic. People live a long time. One-period problems would still be a useful guide if the world were i.i.d., so that each day looked like the last. Alas, the overwhelming evidence from empirical work is that the world is *not* i.i.d. Expected returns, variances and covariances all change through time. Even if this were not the case, individual investors' outside incomes vary with time, age and the lifecycle. We need a portfolio theory that incorporates long-lived agents, and allows for time-varying moments of asset returns. Furthermore, many dynamic setups give rise to incomplete markets, since shocks to forecasting variables are not traded.

This seems like a lot of complexity, and it is. Fortunately, with a little reinterpretation of symbols, we can apply everything we have done for one-period markets to this intertemporal dynamic world.

I start with a few classic examples that should be in every financial economists' toolkit, and then draw the general point.

3.1 Portfolios and discount factors in intertemporal models

The same optimal portfolio formulas hold in an intertemporal model,

$$\beta^t u'(\hat{x}_t + e_t) = \lambda m_t \tag{31}$$

$$\hat{x}_t = u'^{-1}(\lambda m_t / \beta^t) - e_t. \tag{32}$$

where we now interpret $\hat{x}_t$ to be the flow of dividends (payouts) of the optimal portfolio, and e_t is the flow of outside income.

Start with an investor with no outside income; his utility function is

$$E \sum_{t=1}^{\infty} \beta^t u(c_t).$$

He has initial wealth W and he has a *stream* of outside income $\{e_t\}$. His problem is to pick a *stream* of payoffs or dividends $\{\hat{x}_t\}$, which he will eat, $c_t = \hat{x}_t + e_t$.

As before, we summarize the assets available to the investor by a discount factor m. Thus, the problem is

$$\max_{\{\hat{x}_t \in \underline{X}\}} E \sum_{t=1}^{\infty} \beta^t u(\hat{x}_t + e_t) \ s.t. \ W = E \sum_{t=1}^{T} m_t \hat{x}_t$$

Here m_t represents a discount factor *process*, i.e. for every payoff x_t, m_t generates prices p by

$$p = E\sum_{t=1}^{\infty} m_t x_t.$$

As before, absence of arbitrage and the law of one price guarantee that we can represent the prices and payoffs facing the investor by such a discount factor process.

The first order conditions to this problem are ($\partial/\partial \hat{x}_{it}$ in state i at time t)

$$\beta^t u'(\hat{x}_t + e_t) = \lambda m_t \tag{33}$$

Thus, once again the optimal payoff is characterized by

$$\hat{x}_t = u'^{-1}(\lambda m_t/\beta^t) - e_t. \tag{34}$$

The formula is only different because utility of consumption at time t is multiplied by β^t. If m is unique (complete markets), then we are done. If not, then again we have to choose the right $\{m_t\}$ so that $\{\hat{x}\} \in \underline{X}$. (We have to think in some more detail what this payoff space looks like when markets are not complete.)

As before, this condition characterizes the solution up to initial wealth. To match it to a specific initial wealth (or to find what wealth corresponds to a choice of λ), we impose the constraint,

$$E\sum_t m_t u'^{-1}(\lambda m_t) = W.$$

The corresponding continuous time formulation is

$$\max E\int_{t=0}^{\infty} e^{-\rho t}u(\hat{x}_t + e_t)dt \quad \text{s.t.} \quad W = E\int_{t=0}^{\infty} m_t \hat{x}_t dt$$

giving rise to the identical conditions

$$e^{-\rho t}u'(\hat{x}_t + e_t) = \lambda m_t \tag{35}$$

$$\hat{x}_t = u'^{-1}(\lambda m_t/e^{-\rho t}) - e_t. \tag{36}$$

3.2 The power-lognormal problem

We solve for the optimal infinite-horizon portfolio problem in the lognormal iid setup. The answer is that optimal consumption or dividend is a power function of the current stock value,

$$\hat{x}_t = (\text{const.}) \times \left(\frac{S_t}{S_0}\right)^{\alpha}; \; \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}$$

To see this analysis more concretely, and for its own interest, let's solve a classic problem. The investor has no outside income, lives forever and wants intermediate consumption, and has power utility

$$\max E\int_{t=0}^{\infty}\frac{\hat{x}_t^{1-\gamma}}{1-\gamma}dt.$$

He can dynamically trade, resulting in "complete" markets.

Once we have a discount factor m_t that represents asset markets, the answer is simple. From (36)

$$\hat{x}_t=\lambda^{-\frac{1}{\gamma}}\left(e^{\rho t}m_t\right)^{-\frac{1}{\gamma}}$$

As before, we can solve for λ,

$$\begin{aligned} W &= E\int_{t=0}^{\infty}m_t\lambda^{-\frac{1}{\gamma}}\left(e^{\rho t}m_t\right)^{-\frac{1}{\gamma}}dt \\ W &= \lambda^{-\frac{1}{\gamma}}E\int_{t=0}^{\infty}e^{-\frac{\rho}{\gamma}t}\ m_t^{1-\frac{1}{\gamma}}\ dt \end{aligned}$$

so the optimal payoff is

$$\frac{\hat{x}_t}{W}=\frac{e^{-\frac{\rho}{\gamma}t}\ m_t^{-\frac{1}{\gamma}}}{E\int_{t=0}^{\infty}e^{-\frac{\rho}{\gamma}t}\ m_t^{1-\frac{1}{\gamma}}\ dt} \tag{37}$$

The analogy to the one-period result (6) is strong. However, the "return" is now a dividend at time t divided by an initial value, an insight I follow up on below.

We might insist that the problem be *stated* in terms of a discount factor. But in practical problems, we will first face the technical job of find the discount factor that represents a given set of asset prices and payoffs, so to make the analysis concrete and to solve a classic problem, let's introduce some assets and find their discount factor. As before, a stock and bond follow

$$\frac{dS}{S} = \mu dt+\sigma dz \tag{38}$$

$$\frac{dB}{B} = rdt. \tag{39}$$

(Think of S and B as the cumulative value process with dividends reinvested, if you're worried about transversality conditions. What matters is a stock return $dR=\mu dt+\sigma dz$ and bond return rdt.) This is the same setup as the iid lognormal environment of section 2.1.1, but the investor lives forever and values intermediate consumption rather than living for one period and valuing terminal wealth.

Fortunately, we've already found the discount factor, both in chapter 17 and in equation (9) above, $m_t=\Lambda_t/\Lambda_0$ where

$$\frac{d\Lambda}{\Lambda}=-rdt-\frac{\mu-r}{\sigma}dz. \tag{40}$$

We can substitute $d\ln S$ for dz and solve (38)-(40), (algebra below) resulting in

$$\frac{\Lambda_t}{\Lambda_0} = e^{\frac{1}{2}\left(\frac{\mu-r}{\sigma^2}-1\right)(\mu+r)t} \times \left(\frac{S_t}{S_0}\right)^{-\frac{\mu-r}{\sigma^2}}.$$

And thus, for power utility, (37) becomes

$$\hat{x}_t = (\text{const.}) \times \left(\frac{S_t}{S_0}\right)^{\alpha}$$

where again

$$\alpha = \frac{1}{\gamma}\frac{\mu-r}{\sigma^2}$$

Optimal consumption at date t is a power function of the stock value at that date. As you can guess, and as I'll show below, one way to implement this rule is to invest a constantly rebalanced fraction of wealth α in stocks, and to consume a constant fraction of wealth as well. But this is a complete market, so there are lots of equivalent ways to implement this rule.

Evaluating the constant – the denominator of (37) takes a little more algebra and is not very revealing, but here is the final answer:

$$\frac{\hat{x}_t}{W} = \frac{1}{\gamma}\left[\rho + (\gamma-1)\left(r + \frac{1}{2}\gamma\alpha^2\sigma^2\right)\right] e^{-\frac{1}{\gamma}\left[\rho+\frac{1}{2}(\gamma\alpha-1)(\mu+r)\right]t}\left(\frac{S_t}{S_0}\right)^{\alpha} \tag{41}$$

Algebra:

$$d\ln\Lambda = \frac{d\Lambda}{\Lambda} - \frac{1}{2}\frac{d\Lambda^2}{\Lambda^2} = -\left[r + \frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2\right]dt - \frac{\mu-r}{\sigma}dz$$

$$d\ln S = \frac{dS}{S} - \frac{1}{2}\frac{dS^2}{S^2} = \left(\mu - \frac{1}{2}\sigma^2\right)dt + \sigma dz$$

For the numerator, we want to express the answer in terms of S_t. Substituting $d\ln S$ for dz,

$$d\ln\Lambda = -\left[r + \frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2\right]dt - \frac{\mu-r}{\sigma^2}\left[d\ln S - \left(\mu - \frac{1}{2}\sigma^2\right)dt\right]$$

$$d\ln\Lambda = \left[-r - \frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2 + \frac{\mu(\mu-r)}{\sigma^2} - \frac{1}{2}(\mu-r)\right]dt - \frac{\mu-r}{\sigma^2}d\ln S$$

$$d\ln\Lambda = \left[-r - \frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2 + \frac{\mu(\mu-r)}{\sigma^2} - \frac{1}{2}(\mu-r)\right]dt - \frac{\mu-r}{\sigma^2}d\ln S$$

$$d\ln\Lambda = \frac{1}{2}\left[\frac{\mu-r}{\sigma^2} - 1\right](\mu+r)dt - \frac{\mu-r}{\sigma^2}d\ln S$$

$$\ln\Lambda_t - \ln\Lambda_0 = \frac{1}{2}\left[\frac{\mu-r}{\sigma^2} - 1\right](\mu+r)t - \frac{\mu-r}{\sigma^2}\left(\ln S_t - \ln S_0\right)$$

$$m_t = \frac{\Lambda_t}{\Lambda_0} = e^{\frac{1}{2}\left(\frac{\mu-r}{\sigma^2}-1\right)(\mu+r)t}\left(\frac{S_t}{S_0}\right)^{-\frac{\mu-r}{\sigma^2}}.$$

$$\begin{aligned}
e^{-\frac{\rho}{\gamma}t}\, m_t^{-\frac{1}{\gamma}} &= e^{-\frac{\rho}{\gamma}t-\frac{1}{\gamma}\frac{1}{2}\left(\frac{\mu-r}{\sigma^2}-1\right)(\mu+r)t}\left(\frac{S_t}{S_0}\right)^{\frac{\mu-r}{\gamma\sigma^2}} \\
&= e^{-\frac{1}{\gamma}\left[\rho+\frac{1}{2}\left(\frac{\mu-r}{\sigma^2}-1\right)(\mu+r)\right]t}\left(\frac{S_t}{S_0}\right)^{\frac{\mu-r}{\gamma\sigma^2}}
\end{aligned}$$

For the denominator, it's easier to express Λ in terms of a normal random variable.

$$\begin{aligned}
\ln\Lambda_t - \ln\Lambda_0 &= -\left[r+\frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2\right]t - \frac{\mu-r}{\sigma}\sqrt{t}\varepsilon \\
m_t^{1-\frac{1}{\gamma}} &= e^{-\left(1-\frac{1}{\gamma}\right)\left[r+\frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2\right]t-\frac{\mu-r}{\sigma}\left(1-\frac{1}{\gamma}\right)\sqrt{t}\varepsilon} \\
E\left(m_t^{1-\frac{1}{\gamma}}\right) &= e^{\left\{-\left(1-\frac{1}{\gamma}\right)\left[r+\frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2\right]+\frac{1}{2}\left[\frac{\mu-r}{\sigma}\left(1-\frac{1}{\gamma}\right)\right]^2\right\}t} \\
&= e^{-\left(1-\frac{1}{\gamma}\right)\left\{r+\frac{1}{2}\frac{1}{\gamma}\left(\frac{\mu-r}{\sigma}\right)^2\right\}t}
\end{aligned}$$

$$\begin{aligned}
\int_0^\infty e^{-\frac{\rho}{\gamma}t}E\left(m_t^{1-\frac{1}{\gamma}}\right)dt &= \int_0^\infty e^{-\frac{\rho}{\gamma}t}e^{-\left(1-\frac{1}{\gamma}\right)\left\{r+\frac{1}{2}\frac{1}{\gamma}\left(\frac{\mu-r}{\sigma}\right)^2\right\}t}dt \\
&= \int_0^\infty e^{-\frac{1}{\gamma}\left\{\rho+(\gamma-1)\left[r+\frac{1}{2}\frac{1}{\gamma}\left(\frac{\mu-r}{\sigma}\right)^2\right]\right\}t}dt \\
&= \frac{\gamma}{\rho+(\gamma-1)\left[r+\frac{1}{2}\frac{1}{\gamma}\left(\frac{\mu-r}{\sigma}\right)^2\right]}
\end{aligned}$$

Thus,

$$\begin{aligned}
\frac{\hat{x}_t}{W} &= \frac{\rho+(\gamma-1)\left[r+\frac{1}{2}\frac{1}{\gamma}\left(\frac{\mu-r}{\sigma}\right)^2\right]}{\gamma}e^{-\frac{1}{\gamma}\left[\rho+\frac{1}{2}\left(\frac{\mu-r}{\sigma^2}-1\right)(\mu+r)\right]t}\left(\frac{S_t}{S_0}\right)^{\frac{\mu-r}{\gamma\sigma^2}} \\
\frac{\hat{x}_t}{W} &= \frac{1}{\gamma}\left(\rho+(\gamma-1)\left[r+\frac{1}{2}\gamma\alpha^2\sigma^2\right]\right)e^{-\frac{1}{\gamma}\left[\rho+\frac{1}{2}(\gamma\alpha-1)(\mu+r)\right]t}\left(\frac{S_t}{S_0}\right)^{\alpha}.
\end{aligned}$$

3.3 A mapping to one-period problems

I introduce a little notation that makes even clearer the analogy between multipleriod and one-period problems. $\mathcal{E}(x) \equiv E\sum_{t=1}^{\infty}\beta^t x_t$ treats time and states symmetrically. Then, we write $p = E\left(\sum_{t=1}^{\infty}\beta^t m_t x_t\right) = \mathcal{E}(mx)$. In this notation, the infinite-period dynamic problem looks exactly the same as the one period problem.

The analogy in the above examples to the one-period analysis is striking. Obviously, one-period and multiperiod models are the same in a deep sense.

To make the analogy closest, let us define an expectation operator that adds over *time* using β^t or $e^{-\rho t}$ as it adds over *states* using probabilities. Thus, define

$$\begin{aligned} \text{one period:} \quad & \mathcal{E}(x) \equiv E(x_1) = \sum_s \pi(s) x_1(s) \\ \text{infinite period, discrete:} \quad & \mathcal{E}(x) \equiv E\sum_{t=1}^{\infty} \beta^t x_t = \sum_{t=1}^{\infty}\sum_{s_t} \beta^t \pi(s_t) x_t(s_t) \\ \text{infinite period, continuous} \ : \quad & \mathcal{E}(x) \equiv E\int_0^{\infty} e^{-\rho t} x_t dt \end{aligned}$$

It is convenient to take β as the investor's discount factor, but not necessary.

With this definition, infinite horizon portfolio theory looks *exactly* like one period theory. We write asset pricing as

$$p(x) = E\sum_{t=1}^{\infty} \beta^t m_t x_t = \mathcal{E}(mx).$$

Here it is convenient to start with a discount factor that is scaled by β^t in order to then multiply by β^t. In the canonical example which was expressed $m_t = \beta^t u\prime(c_t)/u'(c_0)$ we now have $m_t = u\prime(c_t)/u'(c_0)$.

(One problem with this definition is that the weights over time do not add up to one, $\mathcal{E}(1) = \beta/(1-\beta)$. One can define $\mathcal{E}(x) = \frac{(1-\beta)}{\beta}\sum \beta^t E(x_t)$ to restore this property, but then we must write pricing as $p(x) = \beta/(1-\beta)\mathcal{E}(mx)$. I choose the simpler pricing equation, at the cost that you have to be careful when taking long run means $\mathcal{E}$ of constants.)

The investor's objective is

$$\begin{aligned} \max E\sum_t \beta^t u(c_t) &= \max \mathcal{E}\left[u(c)\right] \\ c_t &= \hat{x}_t + e_t \end{aligned}$$

The constraint is

$$W = E\sum_t \beta^t m_t \hat{x}_t = \mathcal{E}(m\hat{x})$$

In sum, we are *exactly* back to

$$\max \mathcal{E}\left[u(\hat{x}_t + e_t)\right] \text{ s.t. } W = \mathcal{E}(m\hat{x})$$

The first order conditions are

$$\begin{aligned} u'(\hat{x} + e) &= \lambda m \\ \hat{x} &= u'^{-1}(\lambda m) - e \end{aligned}$$

exactly as before. (We rescaled m, which is why it's not m/β^t as in (34).)

With power utility and no outside income, we can evaluate the constraint as

$$W = \mathcal{E}(m\lambda^{-\frac{1}{\gamma}}m^{-\frac{1}{\gamma}})$$

so again the complete problem is

$$\frac{\hat{x}_t}{W} = \frac{m_t^{-\frac{1}{\gamma}}}{\mathcal{E}(m^{1-\frac{1}{\gamma}})}$$

All the previous analysis goes through unchanged!

3.3.1 Units

> Reinterpreting the symbols from one period problems, "Returns" $x/p(x)$ are now dividends at time t divided by initial price, and the "long-run mean-variance frontier" values stability over time as well as states of nature.

We do, however, have to reinterpret the symbols. $\hat{x}/W$ is now a *dividend stream* divided by its time-0 price. This, apparently is the right generalization of "return" to an infinite-horizon model. More generally, for any payoff stream I think it is better to call the "return" a "yield,"

$$y_t = \frac{x_t}{p(\{x_t\})} = \frac{x_t}{\mathcal{E}(mx)}.$$

Its typical size will be something like 0.04 not 1.04. Similarly, we can define"excess yields", which are the zero-price objects as

$$y_t^e = y_t^1 - y_t^2.$$

The *risk free* payoff is thus one in all states and dates, a perpetuity

$$x_t^f = 1.$$

risk free yield is therefore

$$y_t^f = \frac{1}{p(\{1\})}$$

This is, in fact, the coupon yield of the perpetuity.

I think this observation alone makes a good case for looking at prices and payoff streams rather than one period returns. In the standard Merton-style or period to period analysis, a long term bond is a security that is attractive because its price happens to go up a lot when interest rates decline. Thus, it provides a good hedge for a long-term highly risk averse investor. The fact that a 10 year bond is the riskless asset for an investor with a 10 year horizon, or an indexed perpetuity is the riskless asset for an investor with an infinite horizon, is a feature hidden deep in value functions. But once you look at prices and payoffs, it's just obvious that the indexed perpetuity is the riskless asset for a long-term investor.

In place of our usual portfolios and payoff spaces, we have spaces of yields,

$$\underline{Y} \equiv \{y \in \underline{X} : p(y) = 1\},$$

$$\underline{Y}^e \equiv \{y^e \in \underline{X} : p(y^e) = 0\}.$$

It's natural to define a *long-run mean / long-run variance frontier* which solves

$$\min_{\{y \in \underline{Y}\}} \mathcal{E}(y^2) \text{ s.t. } \mathcal{E}(y) = \mu.$$

"Long run variance" prizes stability over *time* as well as stability across states of nature. If we redo exactly the same algebra as before, we find that the long-run frontier is generated as

$$y^{mv} = y^* + wy^{e*}. \tag{42}$$

Here, y^* is the discount-factor mimicking portfolio return,

$$y^* = \frac{x^*}{p(x^*)} = \frac{x^*}{\mathcal{E}(x^{*2})}. \tag{43}$$

If a riskfree rate is traded, y^{e*} is simply

$$y^{e*} = \frac{y^f - y^*}{y^f}. \tag{44}$$

The long-run mean-variance frontier of excess returns is

$$\min_{\{y^e \in \underline{Y}^e\}} \mathcal{E}(y^{e2}) \text{ s.t. } \mathcal{E}(y^e) = \mu.$$

This frontier is generated simply by

$$y^{emv} = wy^{e*} \;\; w \in \Re$$

Our payoff spaces need to include dynamic trading. For example, if you see a variable z_t (e.g. dividend yields, cay) that forecasts returns R_{t+1}, you want to include in your portfolio dynamic trading strategies that depend on z_t. One easy way to do this is simply to include managed portfolios, as we did in Chapter 8. If a variable z_t is useful for describing time-variation in future returns or payoffs, and if x_{t+1} is a price-zero payoff, then we just include payoffs of the form $f(z_t)\, x_{t+1}$ in the payoff space $\underline{X}$, and inspired by a Taylor approximation, $z_t x_{t+1}, z_t^2 x_{t+1}$, etc. This is equivalent to including payoffs (dividend streams) from real managed portfolios, for example mutual funds or hedge funds that implement dynamic trading. In the real world as well as in this formal sense, dynamic trading means that funds or managed portfolios can synthesize payoffs not available from an original set of assets. Therefore, a time-invariant choice of *managed portfolios* is exactly equivalent to a *dynamic strategy.*

A particularly important kind of market incompleteness occurs with dynamic trading, however. If a variable z_t forecasts returns, the discount factor and hence optimal portfolios will typically depend on shocks to z_t, which may not be traded. For example, there is no tradeable security that pays the innovation to dividend yields in a VAR. This observation in particular motivates us to understand incomplete markets in a dynamic setting.

3.3.2 Incomplete markets and the long-run mean long-run variance frontier

---

With quadratic utility and no outside income, a long-run investor in a dynamic, interetemporal, incomplete market wants a portfolio on the long-run mean-variance frontier. All investors split their payoffs between an indexed perpetuity and the market payoff, which is a claim to the aggregate consumption stream. The representations derived before for portfolio theory with outside income now apply as well, using outside income flows.

As before, with *incomplete* markets we face the same issue of finding the one of many possible discount factors m which leads to a tradeable payoff. Again, however, we can use the quadratic utility approximation

$$\begin{aligned} u(c) &= -\frac{1}{2}\left(c^b - c\right)^2 \\ U &= \mathcal{E}\left[-\frac{1}{2}\left(c^b - c\right)^2\right] = E\sum_t \beta^t\left(-\frac{1}{2}\right)(c_t^b - c_t)^2 \end{aligned}$$

and the above analysis goes through *exactly.* Again, all we have to do is to reinterpret the symbols.

The optimal portfolio with a nonstochastic bliss point and no labor income is

$$\hat{y} = y^f + \frac{1}{\gamma}\left(y^f - y^*\right).$$

We recognize a long-run mean/long-run variance efficient portfolio on the right hand side. Aggregating across identical individuals we have

$$\hat{y}^i = y^f + \frac{\gamma^a}{\gamma^i}\left(\hat{y}^m - y^f\right).$$

Thus, the classic propositions have straightforward reinterpretations:

1. *Each investor holds a portfolio on the long-run mean/ long-run variance frontier.*
2. *The market portfolio is also on the long-run mean / long-run variance frontier.*
3. *Each investor's portfolio can be spanned by a real perpetuity y^f and a claim to aggregate consumption $\hat{y}^m$*

In the absence of outside income, a "long-run" version of the CAPM holds in this economy, since the market is "long-run" efficient.

Keep in mind that *all* of this applies with arbitrary return dynamics – we are *not* assuming iid returns – and it holds with incomplete markets, in particular that innovations to state variables are not traded.

As conventional mean-variance theory gave a useful approximate *characterization* of optimal portfolios without actually *calculating* them – finding the mean-variance frontier is hard – so here we give an approximate *characterization* of optimal portfolios in a fully dynamic, intertemporal, incomplete markets context. Calculating them – finding x^*, y^*, the long run mean-long run variance frontier, or supporting a payoff $\hat{x}$ with dynamic trading in specific assets – will also be hard.

4 Portfolio theory by choosing portfolio weights

The standard approach to portfolio problems is quite different. Rather than summarize assets by a discount factor and choose the final *payoff*, you specify the assets explicitly and choose the portfolio weights.

For example, we can solve a one-period problem in which the investor chooses among returns R^f and R

$$\max_{\{\alpha\}} Eu\left[W_0\left(R^f+\alpha' R^e\right)\right]$$

The first order condition is our old friend,

$$E\left[u'(W_T)R^e\right]=0$$

$$E\left[u'\left(W_0\left(R^f+\alpha' R^e\right)\right)\right]=0$$

The obvious easy case to solve will be quadratic utility,

$$E\left[(c^b-W_0\left(R^f+\alpha' R^e\right))R^e\right]=0$$

$$\begin{aligned}
0 &= \left(c^b-W_0R^f\right)E\left(R^e\right)-W_0\alpha' E\left[R^eR^{e\prime}\right]=0\\
\alpha &= \left(\frac{c^b}{W_0}-R^f\right)E\left[R^eR^{e\prime}\right]^{-1}E\left(R^e\right)\\
\alpha &= \frac{1}{\gamma}R^fE\left[R^eR^{e\prime}\right]^{-1}E\left(R^e\right)
\end{aligned}$$

This is the same mean-variance efficient portfolio we've seen before.

Solving interetemporal problems this way is a little harder because the set of portfolios explodes. For example, suppose there is a signal, z_t that predicts returns $R^e_{t+1}=a+bz_t+\varepsilon_{t+1}$, and suppose we want to maximize $E(U(W_T))$ Now, the weights change each period, as functions of z_t and time.

One still can attack a problem like this using a static approach, using the equivalence I stressed in Chapter 8 between managed portfolios and conditioning information. We can form a few sensible but reasonable ad-hoc trading rules, for example $\alpha=a+bz_t$ (portfolios will give us arbitrary linear functions of these rules, thus changing the intercept and slope as needed), and then expand the return space to include returns of these trading rules. The *static* choice between managed portfolios is, in principle, equivalent to the fully *dynamic* portfolio theory, just as in Chapter 8 *unconditional* moments with managed portfolios could, in principle, deliver all the testable implications of a fully dynamic model.

As in Chapter 8, the limitation of this approach is that we don't know for sure that a finite set of ad-hoc trading rules will encompass the truly optimal portfolio. As in that context, we can appeal to the universal practice in static problems: We do not include all individual stocks, bonds, currencies, etc., but first reduce the problem to a small number of portfolios that we feel capture the interesting cross section of returns. If a few cleverly chosen *portfolios* of assets are enough to capture the cross-section, then surely a few cleverly chosen *trading strategies* are enough to capture the dynamic portfolio problem.

However, if one really wants the exact optimum, there is no substitute for searching over the infinite-dimensional space of potential trading strategies. This is the traditional approach to the problem, which I present here along with the classic results of that investigation.

4.1 One period, power-lognormal

I re-solve the one period, power utility, lognormal example by explicitly choosing portfolio weights. The answer is the same, but we learn how to implement the answer by dynamically trading the stock and bond. The portfolio holds a constantly-rebalanced share $\alpha_t = \frac{1}{\gamma}\frac{\mu-r}{\sigma^2}$ in the risky asset.

This is a classic theorem: the fraction invested in the risky asset is independent of investment horizon. It challenges conventional wisdom that young people should hold more stocks since they can afford to wait out any market declines.

This approach is easiest to illustrate in a canonical example, the power-lognormal case we have already studies. At each point in time, the investor puts a fraction α_t of his wealth in the risky asset. Thus the problem is

$$\max_{\{\alpha_t\}} Eu(c_T). \text{ s.t.}$$

$$\begin{aligned} dW_t &= W_t\left[\alpha_t \frac{dS_t}{S_t} + (1-\alpha_t) r dt\right] \\ c_T &= W_T;\ W_0 \text{ given} \end{aligned}$$

I start with the canonical lognormal iid environment,

$$\begin{aligned} \frac{dS_t}{S_t} &= \mu dt + \sigma dz_t \\ \frac{dB}{B} &= r dt. \end{aligned}$$

Substituting, wealth evolves as

$$\frac{dW_t}{W_t} = \left[r + \alpha_t(\mu - r)\right] dt + \alpha_t \sigma dz. \tag{45}$$

We find the optimal weights α_t by dynamic programming. The value function satisfies

$$V(W,t) = \max_{\{\alpha_t\}} E_t V(W_{t+dt}, t+dt)$$

and hence, using Ito's lemma,

$$
\begin{aligned}
0 &= \max_{\{\alpha_t\}} E_t \left\{ V_W dW + \frac{1}{2} V_{WW} dW^2 + V_t dt \right\} \\
0 &= \max_{\{\alpha_t\}} W V_W \left[r + \alpha_t(\mu - r) \right] + \frac{1}{2} W^2 V_{WW} \alpha_t^2 \sigma^2 + V_t \qquad (46)
\end{aligned}
$$

The first order condition for portfolio choice α_t leads directly to

$$
\alpha_t = -\frac{V_W}{W V_{WW}} \frac{\mu - r}{\sigma^2} \qquad (47)
$$

We will end up proving

$$
V(W, t) = k(t) W_t^{1-\gamma}
$$

and thus

$$
\alpha_t = \frac{1}{\gamma} \frac{\mu - r}{\sigma^2}. \qquad (48)
$$

The proportion invested in the risky asset is a constant, independent of wealth and investment horizon. It is larger, the higher the stock excess return, lower variance, and lower risk aversion. Conventional wisdom says you should invest more in stocks if you have a longer horizon; the young should invest in stocks, while the old should invest in bonds. The data paint an interesting converse puzzle: many young people invest in bonds until they build up a safe "nest egg," and the bulk of stock investment is done by people in their mid 50s and later. In this model, the conventional wisdom is wrong.

Of course, models are built on assumptions. A lot of modern portfolio theory is devoted to changing the assumptions so that the conventional wisdom is right, or so that the "safety-first" stylized fact is optimal. For example, time-varying expected returns can raise the Sharpe ratio of long-horizon investments, and so can make it optimal to hold more in stocks for longer investment horizons.

(Actually, the quantity $-\frac{V_W}{WV_{WW}}$ *is* the risk aversion coefficient. Risk aversion is often measured by people's resistance to taking bets. Bets affect your wealth, not your consumption, so aversion to wealth bets measures this quantity. The special results are that in this model, risk aversion is also equal to the local curvature of the utility function γ, and therefore risk aversion is independent of time and wealth, even though $V(W, t)$.)

With the optimal portfolio weights in hand, invested wealth W follows

$$
W_T = W_0 e^{(1-\alpha)\left(r + \frac{1}{2}\sigma^2 \alpha\right)T} \left(\frac{S}{S_0} \right)^{\alpha} \qquad (49)
$$

This is exactly the result we derived above. If $\alpha = 1$, we obtain $W = W_0 \left(S_T / S_0 \right)$, and if $a = 0$ we obtain $W_T = W_0 e^{rT}$, sensibly enough.

Algebra The algebra for (49) is straightforward if uninspiring.

$$\begin{aligned}
\frac{dW_t}{W_t} &= (1-\alpha)\, r dt + \alpha \frac{dS}{S} \\
d\ln W_t &= \frac{dW_t}{W_t} - \frac{1}{2}\frac{dW^2}{W^2} = (1-\alpha)\, r dt + \alpha \frac{dS_t}{S_t} - \frac{1}{2}\alpha^2\sigma^2 dt \\
d\ln S_t &= \frac{dS_t}{S_t} - \frac{1}{2}\frac{dS^2}{S^2} = \frac{dS_t}{S_t} - \frac{1}{2}\sigma^2 dt \\
d\ln W_t &= (1-\alpha)\, r dt + \alpha \left(d\ln S_t + \frac{1}{2}\sigma^2 dt \right) - \frac{1}{2}\alpha^2\sigma^2 dt \\
d\ln W_t &= \left[(1-\alpha)\, r + \frac{1}{2}\sigma^2 \alpha (1-\alpha) \right] dt + \alpha d\ln S_t \\
d\ln W_t &= (1-\alpha) \left(r + \frac{1}{2}\sigma^2 \alpha \right) dt + \alpha d\ln S_t \\
\ln W_T - \ln W_0 &= (1-\alpha) \left(r + \frac{1}{2}\sigma^2 \alpha \right) T + \alpha \left(\ln S_T - \ln S_0 \right) \\
W_T &= W_0 e^{(1-\alpha)\left(r+\frac{1}{2}\sigma^2\alpha\right)T} \left(\frac{S_T}{S_0} \right)^{\alpha}
\end{aligned}$$

The value function

It remains to prove that the value function V really does have the form$V(W,t) = k(t)W_t^{1-\gamma}/(1-\gamma)$, and to find $k(t)$.

Substituting the optimal portfolio α_t into (88), the value function solves the differential equation

$$\begin{aligned}
0 &= WV_W \left[r + \alpha_t(\mu - r) \right] + \frac{1}{2}W^2 V_{WW}\alpha^2\sigma^2 + V_t \\
0 &= \left[r - \frac{V_W}{WV_{WW}}\frac{(\mu-r)}{\sigma^2}(\mu - r) \right] + \frac{1}{2}\frac{W^2 V_{WW}}{WV_W}\left(\frac{V_W}{WV_{WW}}\frac{(\mu-r)}{\sigma^2} \right)^2 \sigma^2 + \frac{V_t}{WV_W} \\
0 &= r - \frac{V_W}{WV_{WW}}\frac{(\mu-r)}{\sigma^2}(\mu - r) + \frac{1}{2}\frac{V_W}{WV_{WW}}\frac{(\mu-r)}{\sigma^2}^2 + \frac{V_t}{WV_W}
\end{aligned}$$

$$0 = r - \frac{1}{2}\frac{V_W}{WV_{WW}}\frac{(\mu-r)^2}{\sigma^2} + \frac{V_t}{WV_W}, \tag{50}$$

subject to the terminal condition

$$u(W_T) = V(W_T).$$

The usual method for solving such equations is to guess the solution up to undertermined parameters or simple functions, and then figure out what those parameters have to be in order for the guess to work. In this case, guess a solution of the form

$$V(W,t) = e^{\eta(T-t)}W^{1-\gamma}$$

Hence,

$$
\begin{aligned}
V_t &= -\eta e^{\eta(T-t)} W^{1-\gamma} \\
V_W &= (1-\gamma) e^{\eta(T-t)} W^{-\gamma} \\
V_{WW} &= -\gamma(1-\gamma) e^{\eta(T-t)} W^{-\gamma-1} \\
-\frac{V_W}{W V_{WW}} &= \frac{1}{\gamma} \\
\frac{V_t}{W V_W} &= -\frac{\eta}{1-\gamma}
\end{aligned}
$$

Plugging in to the PDE (90), that equation holds if the undetermined coefficient η solves

$$0 = r + \frac{1}{2}\frac{1}{\gamma}\frac{(\mu - r)^2}{\sigma^2} - \frac{\eta}{1-\gamma}$$

Hence,

$$\eta = (1-\gamma)\left[r + \frac{1}{2}\frac{1}{\gamma}\frac{(\mu - r)^2}{\sigma^2}\right]$$

and

$$V(W,t) = e^{(1-\gamma)\left[r + \frac{1}{2}\frac{1}{\gamma}\frac{(\mu-r)^2}{\sigma^2}\right](T-t)} W^{1-\gamma}$$

Since our guess works, the portfolio weights are in fact as given by equation (48). You might have guessed just $W^{1-\gamma}$, but having more time to trade and asset to grow makes success more likely.

In a complete-market problem, we don't have to guess the value function. We've already solved the problem and found the final payoffs, so we can compute the value function from the previous contingent-claim solution.

4.2 Comparison with the payoff approach

Having both the discount factor approach and the portfolio weight approach in hand, you can see the appeal of the discount factor-complete markets approach. It took us two lines to get to $\hat{x} = (\text{const}) \times R_T^{\alpha}$, and only a few more lines to evaluate the constant in terms of initial wealth. The portfolio weight approach, by contrast took a lot of algebra. One reason it did so, is that we solved for a lot of stuff we didn't really need. We found not only the optimal *payoff*, but we found a specific dynamic trading strategy to support that payoff. That might be useful. On the other hand, you might want to implement the optimal payoff with a portfolio of call and put options at time zero and not have to spend the entire time dynamically trading. Or you might want to use 2 or 3 call options and then limit your amount of dynamic trading. The advantage of the portfolio choice approach is that you really know the answer is in the payoff space. The disadvantage is that if you make a slight change in the payoff space, you have to start the problem all over again.

Sometimes problems cannot be easily solved by choosing portfolio weights, yet we can easily characterize the payoffs. The habit example with $u'(c) = (c-h)^{-\gamma}$ above is one such example. We solved very quickly for final payoffs. You can try to solve this problem by choosing portfolio weights, but you will fail, in a revealing manner. Equation (89) will still describe portfolio weights. We had not used the form of the objective function in getting to this point. Now, however, the risk aversion coefficient will depend on wealth and time. If you are near $W = h$, you become much more risk averse! We need to solve the value function to see how much so. The differential equation for the value function (90) is also unchanged. The only thing that changes is the terminal condition. Now, we have a terminal condition

$$V(W,T) = (W-h)^{1-\gamma}.$$

Of course, our original guess $V(W,t) = e^{\eta(T-t)}W^{1-\gamma}$ won't match this terminal condition. A natural guess $V(W,t) = e^{\eta(T-t)}(W - f(t)h)^{1-\gamma}$, alas, does not solve the differential equation. The only way I know to proceed analytically is to use the general solution of the differential equation

$$V(W,t) = \int a(\xi) e^{(1-\xi)\left[r+\frac{1}{2}\frac{1}{\gamma}\frac{(\mu-r)^2}{\sigma^2}\right](T-t)} W_t^{1-\xi} d\xi$$

and then find $a(\xi)$ to match the terminal condition. Not fun.

You can see the trouble. We have complicated the problem by asking not just for the answer – the time T payoff or the number of contingent claims to buy –but also by asking for a trading strategy to synthesize those contingent claims from stock and bond trading. We achieved success by being able to stop and declare victory before the hard part. Certainly in this complete market model, it is simpler *first* to characterize the optimal payoff $\hat{x}$, and *then* to choose how to implement that payoff by a specific choice of assets, i.e. put and call options, dynamic trading, pure contingent claims, digital options, etc.

On the other hand, in general *incomplete* markets problems, choosing portfolio weights means you know you always stay in the asset space $\hat{x} \in \underline{X}$.

5 Dynamic intertemporal problems

Now we remove the iid assumption and allow mean returns, variance of returns and outside income to vary over time. I also introduce intermediate consumption in the objective.

5.1 A single-variable Merton problem

We allow mean returns, return volatility and labor income to vary over time. This section simplifies by treating a single risky return and a single state variable. The optimal portfolio weight on the risky asset becomes

$$\alpha_t = \frac{1}{\gamma_t}\frac{\mu_t - r_t}{\sigma_t^2} + \eta_t \beta_{dy,dR}$$

where γ_t and η_t are risk aversion and aversion to the risk that the state variable changes, defined by corresponding derivatives of the value function, and $\beta_{dy,dR}$ is the regression coefficient of state-variable innovations on the risky return.

We see two new effects: 1) "Market timing." The allocation to the risky asset may rise and fall over time, for example if the mean excess return $\mu_t - r_t$ varies and γ_t and σ_t do not. 2) "Hedging" demand. If the return is good for "bad" realizations of the state variable, this raises the desirability and thus overall allocation to the risky asset.

These results simply characterize the optimal portfolio problem without solving for the actual value function. That step is much harder in general.

Here's the kind of portfolio problem we want to solve. We want utility over consumption, not terminal wealth; and we want to allow for time-varying expected returns and volatilities.

$$\max E \int_0^\infty e^{-\rho t} u(c_t) dt \text{ s.t.} \tag{51}$$

$$dR_t = \mu(y_t)dt + \sigma(y_t)dz_t \tag{52}$$

$$dy_t = \mu_y(y_t)dt + \sigma_y(y_t)dz_t \tag{53}$$

The objective can also be or include terminal wealth,

$$\max E \int_0^T e^{-\rho t} u(c_t) dt + EU(W_T).$$

In the traditional Merton setup, the y variables are considered only as state variables for investment opportunities. However, we can easily extend the model to think of them as state variables for labor or proprietary income e_t and include $c_t = x_t + e_t$ as well. I start in this section by specializing to a single state variable y, which simplifies the algebra and gives one set of classic results. The next section uses a vector of state variables and generates a different set of classic results.

If the investor puts weights α in the risky asset, wealth evolves as

$$\begin{aligned} dW &= W\alpha dR + W(1-\alpha)rdt + (e-c)\,dt \\ dW &= [Wr + W\alpha\,(\mu - r) + (e-c)]\,dt + W\alpha\sigma dz \end{aligned}$$

e (really $e(y_t)$) is outside income.

The value function must include the state variable y, so the Bellman equation is

$$V(W,y,t) = \max_{\{c,\alpha\}} u(c)dt + E_t\left[e^{-\rho dt}V(W_{t+dt}, y_{t+dt}, t+dt)\right],$$

using Ito's lemma as usual,

$$\begin{aligned} 0 &= \max_{\{c,\alpha\}} u(c)dt - \rho V dt + V_t dt + V_W E_t\,(dW) + V_y E_t(dy) \\ &\quad + \frac{1}{2}V_{WW}dW^2 + \frac{1}{2}V_{yy}dy^2 + V_{Wy}dWdy. \end{aligned}$$

Next we substitute for dW, dy. The result is

$$\begin{aligned} 0 &= \max_{\{c,\alpha\}} u(c) - \rho V(W,y,t) + V_t + V_W\,[Wr + W\alpha\,(\mu - r) + e - c] + V_y\mu_y \\ &\quad + \frac{1}{2}V_{WW}W^2\alpha^2\sigma^2 + \frac{1}{2}V_{yy}\sigma_y^2 + WV_{Wy}\alpha\sigma\sigma_y. \end{aligned} \tag{54}$$

Now, the first order conditions. Differentiating (54),

$$\frac{\partial}{\partial c} : u'(c) = V_W$$

Marginal utility of consumption equals marginal value of wealth. A penny saved has the same value as a penny consumed.

Next, we find the first order condition for portfolio choice:

$$\begin{aligned} \frac{\partial}{\partial \alpha} &: \quad WV_W(\mu - r) + W^2V_{WW}\sigma^2\alpha + W\sigma\sigma_y V_{Wy} = 0 \\ \alpha &= -\frac{V_W}{WV_{WW}}\frac{(\mu - r)}{\sigma^2} - \frac{\sigma_y}{\sigma}\frac{V_{Wy}}{WV_{WW}} \end{aligned}$$

This is the all-important answer we are looking for: the weights of the optimal portfolio. $\sigma\sigma_y = cov(dR, dy)$ is the covariance of return innovations with state variable innovations, so $\sigma\sigma_y/\sigma^2 = \beta_{dy,dR}$ is the regression coefficient of state variable innovations on return innovations. Thus, we can write the optimal portfolio weight in the risky asset as

$$\alpha = -\frac{V_W}{WV_{WW}}\frac{\mu_t - r_t}{\sigma_t^2} - \frac{V_{Wy}}{WV_{WW}}\beta_{dy,dR} \tag{55}$$

$$= \frac{1}{\gamma}\frac{\mu_t - r_t}{\sigma_t^2} + \frac{\eta}{\gamma}\beta_{dy,dR} \tag{56}$$

In the second line, I have introduced the notation γ for risk aversion and η which measures "aversion" to state variable risk.

$$\gamma \equiv -\frac{WV_{WW}}{V_W}; \eta \equiv \frac{V_{Wy}}{V_W}.$$

The γ here is the local curvature of the value function at time t. It is not necessarily the power of a utility function, nor is it necessarily constant over time.

The first term is the same as we had before. However, the mean and variance change over time – that's the point of the Merton model. Thus, *Investors will "time the market," investing more in times of high mean or low variance.* The second term is new: *Investors will increase their holding of the risky asset if it covaries negatively with state variables of concern to the investor.* "Of concern" is measured by V_{Wy}. This is the "hedging" motive. A long term bond is a classic example. Bond prices go up when subsequent yields go down. Thus a long-term bond is an excellent hedge for the risk that interest rates decline, meaning your investment opportunities decline. Investors thus hold more long term bonds than they otherwise would, which may account for low long-term bond returns. Since stocks now mean-revert too, we should expect important quantitative results from the Merton model: *mean-reversion in stock prices will make stocks even more attractive.*

(This last conclusion depends on risk aversion, i.e. whether substitution or wealth effects dominate. Imagine that news comes along that expected returns are much higher. This has two effects. First there is a "wealth effect." The investor will be able to afford a lot more consumption in the future. But there is also a "substitution effect." At higher expected returns, it pays the investor to consume *less* now, and then consume even more in the future, having profited by high returns. If risk aversion, equal to intertemporal substitution, is high, the investor will not pay attention to the latter incentive. Raising consumption in the future means consumption rises now, so $V_W = u'(c)$ declines now, i.e. $V_{Wy} < 0$. However, if risk aversion is very low, the substitution effect will dominate. The investor consumers less now, so as to invest more. This means $V_W = u'(c)$ rises, and $V_{Wy} > 0$. Log utility is the knife edge case in which substitution and wealth effects offset, so $V_{Wy} = 0$. We usually think risk aversion is greater than log, so that case applies.)

Of course, risk aversion and state variable aversion are not constants, nor are they determined by preferences alone. This discussion presumes that risk aversion and state variable aversion do not change. They may. Only by fully solving the Merton model can we really see the portfolio implications.

5.1.1 Completing the Merton model

Actually finding the value function in the Merton problem is not easy, and has only been accomplished in a few special cases. Most applied work uses approximations or numerical methods.

Conceptually this step is simple, as before: we just need to find the value function. We plug optimal portfolio and consumption decisions into 54 and solve the resulting partial differential equation. However, even a brief look at the problem will show you why so little has been done on this crucial step, and thus why *quantitative* use of Merton portfolio theory languished for 20 years until the recent revival of interest in approximate solutions. The partial differential equation is, from (54) (algebra below)

$$\begin{aligned} 0 &= u\left[u'^{-1}(V_W)\right] - \rho V + V_t + WV_W r - V_W u'^{-1}(V_W) + V_W e + V_y \mu_y + \frac{1}{2} V_{yy} \sigma_y^2 \\ &\quad -\frac{1}{2} \frac{1}{\sigma^2 V_{WW}} \left[V_W(\mu - r) + \sigma \sigma_y V_{Wy}\right]^2. \end{aligned}$$

This is not a pleasant partial differential equation to solve, and analytic solutions are usually not available. The nonlinear terms $u(u'^{-1}(V_W))$ and $u'^{-1}(V_W)$ are especially troublesome, which accounts for the popularity of formulations involving the utility of terminal wealth, for which these terms are absent.

There are analytical solutions for the following special cases:

1. Power utility, infinite horizon, no state variables. As you might imagine, $V(W) = W^{1-\gamma}$ works again. This is a historically important result as it establishes that the CAPM holds even with infinitely lived, power utility investors, so long as returns are i.i.d. over time and there is no labor income. I solve it in the next subsection

2. Log utility, no labor income. In this case, $V_{Wy} = 0$, the investor does no intertemporal hedging. Now we recover the log utility CAPM, even when there are state variables.

3. Power utility of terminal wealth (no consumption), $AR(1)$ state variable, no labor income, power (or more generally HARA) utility. (Kim and Omberg 1996). Here the natural guess that $V(W, y, t) = W^{1-\gamma} \exp\left[a(T-t) + b(T-t)y + c(T-T)y^2\right]$ works, though solving the resulting differential equation is no piece of cake.

4. Approximations or numerical evaluation. This is the approach taken by most of the huge literature that studies Merton problems in practice.

Algebra: Plugging optimal consumption c and portfolio α decisions into (54),

$$\begin{aligned} 0 &= u(c) - \rho V + V_t + V_W\left[Wr + W\alpha(\mu - r) + e - c\right] + V_y \mu_y \\ &\quad + \frac{1}{2} V_{WW} W^2 \alpha^2 \sigma^2 + \frac{1}{2} V_{yy} \sigma_y^2 + WV_{Wy} \alpha \sigma \sigma_y \end{aligned}$$

$$\begin{aligned} 0 &= u(u'^{-1}(V_W)) - \rho V + V_t + WV_W r - V_W u'^{-1}(V_W) + V_W e + V_y \mu_y + \frac{1}{2} V_{yy} \sigma_y^2 \\ &\quad + \frac{1}{2} V_{WW} W^2 \alpha^2 \sigma^2 + W\left(V_W(\mu - r) + V_{Wy} \sigma \sigma_y\right) \alpha \end{aligned}$$

$$
\begin{aligned}
0 &= u(u'^{-1}(V_W)) - \rho V + V_t + WV_W r - V_W u'^{-1}(V_W) + V_W e + V_y \mu_y + \frac{1}{2} V_{yy} \sigma_y^2 \\
&+ \frac{1}{2} V_{WW} W^2 \sigma^2 \left[\frac{V_W}{WV_{WW}} \frac{(\mu - r)}{\sigma^2} + \frac{\sigma \sigma_y}{\sigma^2} \frac{V_{Wy}}{WV_{WW}} \right]^2 \\
&- W \left[V_W(\mu - r) + V_{Wy} \sigma \sigma_y \right] \left[\frac{V_W}{WV_{WW}} \frac{(\mu - r)}{\sigma^2} + \frac{\sigma \sigma_y}{\sigma^2} \frac{V_{Wy}}{WV_{WW}} \right]
\end{aligned}
$$

$$
\begin{aligned}
0 &= u(u'^{-1}(V_W)) - \rho V + V_t + WV_W r - V_W u'^{-1}(V_W) + V_W e + V_y \mu_y + \frac{1}{2} V_{yy} \sigma_y^2 \\
&+ \frac{1}{2} \frac{1}{\sigma^2 V_{WW}} \left[V_W(\mu - r) + \sigma \sigma_y V_{Wy} \right]^2 \\
&- \frac{1}{\sigma^2 V_{WW}} \left[V_W(\mu - r) + V_{Wy} \sigma \sigma_y \right] \left[V_W(\mu - r) + \sigma \sigma_y V_{Wy} \right]
\end{aligned}
$$

$$
\begin{aligned}
0 &= u(u'^{-1}(V_W)) - \rho V + V_t + WV_W r - V_W u'^{-1}(V_W) + V_W e + V_y \mu_y + \frac{1}{2} V_{yy} \sigma_y^2 \\
&- \frac{1}{2} \frac{1}{\sigma^2 V_{WW}} \left[V_W(\mu - r) + \sigma \sigma_y V_{Wy} \right]^2 .
\end{aligned}
$$

5.2 The power-lognormal iid model with consumption

I solve the power utility infinite-horizon model with iid returns and no outside income. The investor consumes a constant proportion of wealth, and invests a constant share in the risky asset.

In the special case of power utility, no outside income and iid returns, the differential equation (54) specializes to

$$
0 = \frac{V_W^{-\frac{1}{\gamma}(1-\gamma)}}{1-\gamma} - \rho V + V_t + WV_W r - V_W V_W^{-\frac{1}{\gamma}} - \frac{1}{2} \frac{1}{\sigma^2 V_{WW}} \left[V_W(\mu - r) \right]^2
$$

To solve it, we guess a functional form

$$
V = k \frac{W^{1-\gamma}}{1-\gamma}.
$$

Plugging in, we find that the differential equation holds if

$$
k^{-\frac{1}{\gamma}} = \frac{\rho}{\gamma} - \frac{1-\gamma}{\gamma} \left[r + \frac{1}{2} \frac{(\mu - r)^2}{\gamma \sigma^2} \right].
$$

Hence, we can fully evaluate the policy: Optimal consumption follows

$$
c = V_W^{-\frac{1}{\gamma}} = \frac{1}{\gamma} \left[\rho - (1-\gamma) \left(r + \frac{1}{2} \frac{(\mu - r)^2}{\gamma \sigma^2} \right) \right] W \tag{57}
$$

The investor consumes a constant share of wealth W. For log utility ($\gamma = 1$) we have $c = \rho W$. The second term only holds for utility different than log. If $\gamma > 1$, higher returns (either a higher risk free rate or the higher squared Sharpe ratio in the second term) lead you to raise consumption. Income effects are greater than substitution effects (high γ resists substitution), so the higher "wealth effect" means more consumption now. If $\gamma < 1$, the opposite is true; the investor takes advantage of higher returns by consuming less now, building wealth up faster and then consuming more later. The risky asset share is, from (55),

$$\alpha = \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}. \tag{58}$$

We already had the optimal consumption stream in (41). What we learn here is that we can support that stream by the consumption rule (57) and portfolio rule (58).

The Algebra

$$\begin{aligned} V &= k\frac{W^{1-\gamma}}{1-\gamma} \\ V_W &= kW^{-\gamma} \\ V_{WW} &= -\gamma k W^{-\gamma-1} \end{aligned}$$

$$\begin{aligned} 0 &= \frac{k^{-\frac{1}{\gamma}(1-\gamma)} W^{(1-\gamma)}}{1-\gamma} - \rho k \frac{W^{1-\gamma}}{1-\gamma} + WkW^{-\gamma} r - \left(kW^{-\gamma}\right)^{1-\frac{1}{\gamma}} + \frac{1}{2}\frac{\left(kW^{-\gamma}\right)^2}{\gamma k W^{-\gamma-1}}\frac{(\mu-r)^2}{\sigma^2} \\ 0 &= \frac{k^{1-\frac{1}{\gamma}}}{1-\gamma} W^{1-\gamma} - \frac{\rho k}{1-\gamma} W^{1-\gamma} + rkW^{1-\gamma} - k^{1-\frac{1}{\gamma}} W^{1-\gamma} + \frac{1}{2}\frac{(\mu-r)^2}{\sigma^2}\frac{k}{\gamma} W^{1-\gamma} \\ 0 &= \frac{k^{-\frac{1}{\gamma}}}{1-\gamma} - \frac{\rho}{1-\gamma} + r - k^{-\frac{1}{\gamma}} + \frac{1}{2}\frac{(\mu-r)^2}{\gamma\sigma^2} \\ 0 &= \left(\frac{\gamma}{1-\gamma}\right) k^{-\frac{1}{\gamma}} - \frac{\rho}{1-\gamma} + r + \frac{1}{2}\frac{(\mu-r)^2}{\gamma\sigma^2} \\ k^{-\frac{1}{\gamma}} &= \frac{\rho}{\gamma} - \frac{1-\gamma}{\gamma}\left[r + \frac{1}{2}\frac{(\mu-r)^2}{\gamma\sigma^2}\right] \end{aligned}$$

5.3 Multivariate Merton problems and the ICAPM

I characterize the infinite-period portfolio problem with multiple assets and multiple state variables. The optimal portfolio weights

$$\alpha = \frac{1}{\gamma}\Sigma^{-1}(\mu - r) + \beta'_{dy,dR}\frac{\eta}{\gamma}$$

include a mean-variance efficient portfolio, but also include mimicking portfolios for state-variable risks

Now, let's solve the same problem with a vector of asset returns and a vector of state variables. This generalization allows us to think about how the investor's choice *among* assets may be affected by time-varying investment opportunities and by labor income and state variables for labor income. We start as before,

$$\max E \int_0^\infty e^{-\rho t} u(c_t) \text{ s.t.} \tag{59}$$

$$dR_t = \mu(y_t)dt + \sigma(y_t)dz_t \tag{60}$$

$$dy_t = \mu_y(y_t)dt + \sigma_y(y_t)dz_t \tag{61}$$

$$de_t = \mu_e(y_t)dt + \sigma_e(y_t)dz_t \tag{62}$$

Now I use dR to denote the vector of N returns dS_i/S_i, so μ is an N dimensional vector. y is a K dimensional vector of state variables. dz is an (at least) $N+K$ dimensional vector of independent shocks, $E_t(dzdz') = I$. Thus, σ is an $N \times (N+K)$ dimensional matrix and σ_y is a $K \times (N+K)$ dimensional matrix. I'll examine the case in which one asset is a risk free rate, r_t Since it varies over time, it is one of the elements of y_t. The conventional statement of the problem ignores outside income and only thinks of state variables that drive the investment opportunity set, but since labor income is important and all the results we will get to accommodate it easily, why not include it.

Now, if the investor puts weights α on the risky assets, wealth evolves as

$$\begin{aligned} dW &= W\left(\alpha' dR\right) + W(1 - 1'\alpha)rdt + (e-c)\,dt \\ dW &= \left[Wr + W\alpha'(\mu - r) + (e-c)\right]dt + W\alpha'\sigma dz. \end{aligned}$$

The Bellman equation is

$$V(W,y,t) = \max_{\{c,\alpha\}} u(c)dt + E_t\left[e^{-\rho dt}V(W_{t+dt}, y_{t+dt}, t+dt)\right],$$

and using Ito's lemma as usual,

$$\begin{aligned} 0 &= \max_{\{c,\alpha\}} u(c)dt - \rho V dt + V_t dt + V_W E_t(dW) + V_{y'}E_t(dy) \\ &\quad + \frac{1}{2}V_{WW}dW^2 + \frac{1}{2}dy'V_{yy'}dy + dWV_{Wy'}dy. \end{aligned}$$

I use the notation $V_{y'}$ to denote the row vector of derivatives of V with respect to y. V_y would be a corresponding column vector. $V_{yy'}$ is a matrix of second partial derivatives.

Next we substitute for dW, dy. The result is

$$\begin{aligned} 0 &= \max_{\{c,\alpha\}} u(c) - \rho V(W,y,t) + V_t + V_W\left[Wr + W\alpha'(\mu - r) - c\right] + V_{y'}\mu_y \\ &\quad + \frac{1}{2}V_{WW}W^2\alpha'\sigma\sigma'\alpha + \frac{1}{2}Tr(\sigma_y' V_{yy'}\sigma_y) + W\alpha'\sigma\sigma_y' V_{Wy} \end{aligned} \tag{63}$$

This is easy except for the second derivative terms. To derive them

$$E(dz'Adz) = \sum_{i,j} dz_i A_{ij} dz_j = \sum_i A_{ii} = Tr(A).$$

Then,

$$dy'V_{yy'}dy = (\sigma_y dz)'V_{yy'}(\sigma_y dz) = dz'\sigma_y' V_{yy'}\sigma_y dz = Tr(\sigma_y' V_{yy'}\sigma_y).$$

We can do the other terms similarly,

$$\begin{aligned}
dWV_{Wy'}dy &= (W\alpha'\sigma dz)'V_{Wy'}(\sigma_y dz) = Wdz'\sigma'\alpha V_{Wy'}\sigma_y dz \\
&= WTr(\sigma'\alpha V_{Wy'}\sigma_y) = WTr(\alpha'\sigma\sigma_y' V_{Wy}) = W\alpha'\sigma\sigma_y' V_{Wy} \\
V_{WW}dW^2 &= V_{WW}(W\alpha'\sigma dz)'(W\alpha'\sigma dz) = W^2 V_{WW}dz'\sigma'\alpha\alpha'\sigma dz \\
&= W^2 V_{WW}Tr\left(\sigma'\alpha\alpha'\sigma\right) = W^2 V_{WW}Tr\left(\alpha'\sigma\sigma'\alpha\right) = W^2 V_{WW}\alpha'\sigma\sigma'\alpha
\end{aligned}$$

(I used $Tr(AA') = Tr(A'A)$ and $Tr(AB) = Tr(A'B')$. These facts about traces let me condense a $(N+K) \times (N+K)$ matrix to a 1×1 quadratic form in the last line, and let me transform from an expression for which it would be hard to take α derivatives, $Tr(\sigma'\alpha\alpha'\sigma)$, to one that is easy, $\alpha'\sigma\sigma'\alpha$).

Now, the first order conditions. Differentiating (63), we obtain again

$$\frac{\partial}{\partial c} : u'(c) = V_W$$

Differentiating with respect to α,

$$\begin{aligned}
\frac{\partial}{\partial \alpha} &: \quad WV_W(\mu - r) + W^2 V_{WW}\sigma\sigma'\alpha + W\sigma\sigma_y' V_{Wy} = 0 \\
\alpha &= -\frac{V_W}{WV_{WW}}\left(\sigma\sigma'\right)^{-1}(\mu - r) - \left(\sigma\sigma'\right)^{-1}\sigma\sigma_y' \frac{V_{Wy}}{WV_{WW}}
\end{aligned}$$

This is the all-important answer we are looking for: the weights of the optimal portfolio. It remains to make it more intuitive. $\sigma\sigma' = cov(dR, dR') = \Sigma$ is the return innovation covariance matrix. $\sigma\sigma_y' = cov(dR, dy') = \sigma_{dR,y'}$ is the covariance of return innovations with state variable innovations, and $\left(\sigma\sigma'\right)^{-1}\sigma\sigma_y' = \Sigma^{-1}\sigma_{dR,y'} = \beta'_{dy,dR}$ is a matrix of multiple regression coefficients of state variable innovations on return innovations. Thus, we can write the optimal portfolio weights as

$$\alpha = -\frac{V_W}{WV_{WW}}\Sigma^{-1}(\mu - r) - \beta'_{dy,dR}\frac{V_{Wy}}{WV_{WW}} \tag{64}$$

The first term is exactly the same as we had before, generalized to multiple assets. We recognize in $\Sigma^{-1}(\mu - r)$ the weights of a mean-variance efficient portfolio. Thus we obtain an important result: *In an iid world, investors will hold an instantaneously mean-variance efficient portfolio.* Since we're using diffusion processes which are locally normal, this is the proof behind the statement that normal distributions result in mean-variance portfolios. Mean variance portfolios do not *require* quadratic utility, which I used above to start thinking about mean-variance efficiency. However, note that even if α is constant over time, this means dynamically trading and rebalancing, so that portfolios will not be mean-variance efficient at discrete horizons. In addition, the risky asset share α will generally change over time, giving even more interesting and mean-variance inefficient discrete-horizon returns.

The second term is new: *Investors will shift their portfolio weights towards assets that covary with, and hence can hedge, outside income or changes in the investment opportunity set.* Investors will differ in their degree of risk aversion and "aversion to state variable risk" so we can write the optimal portfolio as

$$\alpha = \frac{1}{\gamma}\Sigma^{-1}(\mu - r) + \beta'_{dy,dR}\frac{\eta}{\gamma} \tag{65}$$

where again

$$\gamma \equiv -\frac{WV_{WW}}{V_W}; \eta \equiv \frac{V_{Wy}}{V_W}.$$

If a positive return on an asset is associated with an increase in the state variable y, and if this increase is associated with an increase in the marginal value of wealth, i.e. $V_{Wy} < 0$, then this tendency leads to a greater average return. Increasing the marginal value of wealth indirectly, by changing a state variable, is as important as increasing it directly, by lowering consumption.

5.3.1 Multifactor efficiency and K fund theorems.

The optimal portfolios solve a generalized mean-variance problem: minimize the variance of portfolio returns for given values of portfolio mean return *and* for given values of covariance with state variable innovations. This is the *multifactor efficient* frontier.

We can nicely interpret this result as a generalization of mean-variance portfolio theory, following Fama (1996).*The Merton investor minimizes the variance of return subject to mean return, and subject to the constraint that returns have specified covariance with innovations to state variables.* Let's form portfolios

$$dR^p = \alpha' dR + (1 - \alpha' 1) r dt$$

The suggested mean, variance, covariance problem is

$$\begin{aligned} \min var_t(dR^p) \text{ s.t. } E_t dR^p &= E; \;\; cov_t(dR^p, dy) = \xi \\ \min_{\{\alpha\}} \alpha' \sigma\sigma' \alpha \text{ s.t.} r + \alpha' (\mu - r) &= E \; ; \; \alpha' \sigma\sigma'_y = \xi \end{aligned}$$

Introducing Lagrange multipliers λ_1, λ_2, the first order conditions are

$$\begin{aligned} \sigma\sigma'\alpha &= \lambda_1 (\mu - r) + \sigma\sigma'_y \lambda_2 \\ \alpha &= \lambda_1 (\sigma\sigma')^{-1} (\mu - r) + (\sigma\sigma')^{-1} \sigma\sigma'_y \lambda_2 && (66) \\ \alpha &= \lambda_1 \Sigma^{-1} (\mu - r) + \beta'_{dy,dR} \lambda_2 && (67) \end{aligned}$$

This is exactly the same answer as (64)!

Figure 5 illustrates. As the mean-variance frontier is a hyperbola, the mean-variance-covariance frontier is a revolution of a hyperbola. Fama calls this frontier the set of *multifactor efficient portfolios.* (Covariance with a state variable is a linear constraint on returns, as is the mean. Thus, the frontier is the revolution of a parabola in mean-variance-covariance space, and the revolution of hyperbola in mean-standard deviation-covariance space as shown. I draw the prettier case with no risk free rate. With a risk free rate, the frontier is a cone.) As shown in the picture, we can think of the investor as maximizing preferences defined over mean, variance and covariance of the portfolio, just as previously we could think of the investor as maximizing preferences defined over mean and variance of the portfolio.

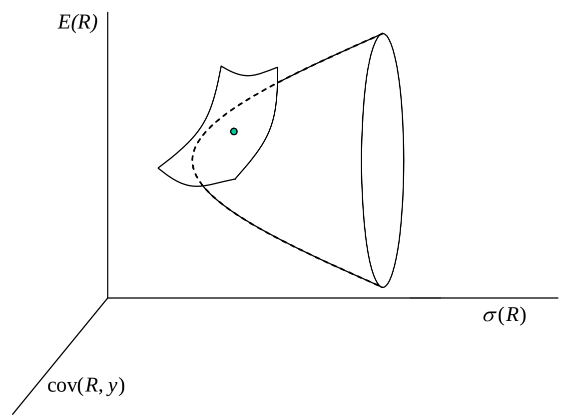

Figure 5: Multifactor effcient portfolio and "indifference curve."

The first term in (64) and (67) is the mean-variance frontier, or a tangency portfolio. (Set $\lambda_2 = 0$ and equation (67) derives this result.) Thus, we see that *typical investors do not hold mean-variance efficient portfolios.* They are willing to give up some mean-variance efficiency in return for a portfolio that hedges the state variable innovations dy.

What do they hold? Mean-variance portfolio theory led to the famous "two fund" theorem. This generalization leads naturally to a $K+2$ fund theorem. *Investors splits their wealth between the tangency portfolio and K mimicking portfolios for state variable innovation risk.* To see this, let's write the investor's optimal portfolio return, not just its weights.

$$\begin{aligned} dR_i &= \alpha^{i\prime} dR + (1 - \alpha^{i\prime} 1) r dt; \\ &= r dt + \alpha^{i\prime} (dR - r dt) \end{aligned}$$

In the latter expression, I split up the investor's portfolio into a risk free investment and an investment in a zero cost portfolio. Following (65), we can split up this portfolio return

$$\begin{aligned} dR_i &= r dt + \frac{1}{\gamma^i} dR^T + \frac{\eta^{i\prime}}{\gamma^i} dR^z \\ dR^T &= (\mu - r)' \Sigma^{-1} (dR - r dt) \qquad (68) \\ dR^z &= \beta_{dy,dR} (dR - r dt) \qquad (69) \end{aligned}$$

We recognize dR^T as a zero-cost investment in a mean-variance efficient or tangency portfolio. The dR^z portfolios are zero-cost portfolios formed from the fitted values of regressions of state variable innovations on the set of asset returns. They are *mimicking portfolios* for the state variable innovations, projection of the state variable innovations on the payoff space. They are also "maximum correlation" portfolios, as regression coefficients minimize residual variance, $\min_{\{\beta_{dy,dR}\}} var\left(dy - \beta_{dy,dR} dR\right)$. Of course, any two K+2 independent multifactor-efficient portfolios will span the multifactor efficient frontier, so you may see other expressions. In particular, I will in a moment express portfolios relative to the *market* portfolio rather than the *tangency* portfolio. The key is to find an *interesting* set of portfolios that span the frontier.

5.3.2 The ICAPM

The ICAPM states that expected returns are proportional to covariances with the market return, but also to covariances with state variables, or mimicking portfolio returns

$$\mu - r = cov(dR, dR^m)\gamma^m - cov(dR, dy')\eta^m$$
$$\mu - r = cov(dR, dR^m)\gamma^m - cov(dR, dR^{z\prime})\eta^m$$

It's always interesting to express portfolio theory with reference to the market portfolio. since in the end we can only hold portfolios other than the market if we are different from everybody else. The market portfolio is the average of individual portfolios, weighted by wealth $\alpha^m = \sum_i W^i \alpha^i / \sum_i W^i$. Thus, summing (65) over investors,

$$\alpha^m = \frac{1}{\gamma^m}\Sigma^{-1}(\mu - r) + \beta'_{dR,dy}\frac{\eta^m}{\gamma^m} \tag{70}$$

Here,

$$\frac{1}{\gamma^m} = \frac{\sum_i W^i \frac{1}{\gamma^i}}{\sum_j W^j}; \;\; \frac{\eta^m}{\gamma^m} = \frac{\sum_i W^i \frac{\eta^i}{\gamma^i}}{\sum_j W^j}$$

The ICAPM solves this expression for the mean excess return

$$\mu - r = \gamma^m \Sigma \alpha^m - \sigma_{dR,dy'}\eta^m \tag{71}$$

(I used $\beta'_{dR,dy} = \Sigma^{-1}\sigma_{dR,dy'}$.) The market portfolio return is

$$dR^m = rdt + \alpha^{m\prime}\left(dR - r\right)dt$$

Thus, we recognize

$$\Sigma\alpha^m = cov(dR, dR')\alpha^m = cov(dR, dR^m)$$

and we have *The ICAPM:*

$$\mu - r = cov(dR, dR^m)\gamma^m - cov(dR, dy')\eta^m.$$

Mean excess returns are driven by covariance with the market portfolio and covariance with each of the state variables. The risk aversion and "state-variable aversion" coefficients give the slopes of average return on covariances.

Since the mimicking portfolio returns dR^z in (69) are the projections of state variables on the space of excess returns, $cov(dR, dy') = cov(dR, dR^{z\prime})$. Directly,

$$\begin{aligned} dR^z &= \beta_{dy,dR}\left(dR - rdt\right) \\ &= cov(dy, dR')cov(dR, dR')^{-1}\left(dR - rdt\right) \\ cov(dR, dR^{z\prime}) &= cov(dR, dR')cov(dR, dR')^{-1}cov(dR, dy') \end{aligned}$$

Thus, we can express the ICAPM in terms of covariances with mimicking portfolios,

$$\mu - r = cov(dR, dR^m)\gamma^m - cov(dR, dR^{z\prime})\eta^m.$$

Since the state variables are often nebulous or hard to measure, this form is used widely in practice.

(Historically, "ICAPM" only refers to models in which the other variables are state variables for investment opportunities, not state variables for outside income, since Merton's original paper did not include outside income. "Multifactor models" encompasses the latter. However, since it's clearly so trivial to include state variables for outside income at least this far, I'll use "ICAPM" anyway.)

This expression with covariance on the right hand side is nice, since the slopes are related to preference (well, value function) parameters. However, it's traditional to express the right hand side in terms of regression betas, and to forget about the economic interpretation of the λ slope coefficients (especially because they are often embarrassingly large). This is easy to do:

$$\begin{aligned} \mu - r &= \beta_{dR,dR^m}\lambda_m - \beta_{dR,dy'}\lambda_{dy} \\ \beta_{dR,dR^m} &= \frac{cov(dR, dR^m)}{\sigma^2_{dR^m}}; \\ \beta_{dR,dy'} &= (cov(dy, dy')^{-1}cov(dy, dR'))' \\ \lambda_m &= \frac{\sigma^2_{dR^m}}{\gamma^m}; \\ \lambda_{dy} &= cov(dy, dy')\eta^m \end{aligned} \tag{72}$$

Now we have expected returns as a linear function of market betas, *and* betas on state variable innovations (or their mimicking portfolios).

Don't forget that all the moments are conditional! The whole point of the ICAPM is that at least one of the conditional mean or conditional variance must vary through time.

This derivation may seem strange. Isn't the ICAPM about "market equilibrium?" How do we jump from a "demand curve" to an equilibrium without saying anything about supply? The answer is that the implicit general equilibrium behind the ICAPM has linear technologies: investors can change the aggregate amount in each security costlessly, without affecting its rate of return. If a security rises in value, investors can and do collectively rebalance away from that security. It is *not* a demand curve which one intersects with a fixed supply of shares to find market prices. In this sense, in a dynamic model, the CAPM and ICAPM are models of the *composition of the market portfolio.* They are not models of price determination.

This insight is important to understand the conundrum, if everybody is like this, how do time-varying returns etc. survive? Suppose all investors have standard power preferences. One would think that return dynamics would be driven out. The answer is that *quantities* adjust. The average investor *does*, for example, buy more when $\mu - r$ is high; as a result he becomes more exposed to risks.

5.3.3 Portfolios relative to the market portfolio

Relative to the market portfolio, rather than the tangency portfolio,

$$dR^i = rdt + \frac{\gamma^m}{\gamma^i}\left(dR^m - rdt\right) + \frac{1}{\gamma^i}\left(\eta^{i\prime} - \eta^{m\prime}\right)dR^z$$

An investor holds more or less of the market portfolio according to his risk aversion, and then more or less of the mimicking portfolios for state variable risk, as his "aversion" to these is greater or less than those of the market.

A mean-variance investor will thus shade his portfolio toward the mimicking portfolio of average state variable risk, therefore providing insurance to other investors for a fee. But for this to work, there must be other investors who hedge even more than average, accepting portfolios with even worse mean-variance properties than the market.

State variables for outside income risk will likely have an idiosyncratic component, for which $\eta^m = 0$. Optimal portfolios thus will contain this individual-specific component.

It's always interesting to express portfolio theory with reference to the market portfolio. since in the end we can only hold portfolios other than the market if we are different from everybody else. The market portfolio is also easier to identify than the tangency portfolio. To this end, use (71) to eliminate $(\mu - r)$ on the right hand side of the individual portfolio weights (65), to obtain

$$\begin{aligned}
\alpha^i &= \frac{1}{\gamma^i}\Sigma^{-1}\left[\gamma^m\Sigma\alpha^m - \sigma_{dR,dy'}\eta^m\right] + \beta'_{dy,dR}\frac{\eta^i}{\gamma^i} \\
\alpha^i &= \frac{\gamma^m}{\gamma^i}\alpha^m + \beta'_{dy,dR}\frac{\left(\eta^i - \eta^m\right)}{\gamma^i}
\end{aligned}$$

If you like looking at the actual portfolio return rather than just the weights,

$$\begin{aligned}
dR^i &= rdt + \alpha^{i\prime}(dR - rdt) \\
dR^i &= rdt + \frac{\gamma^m}{\gamma^i}\left(dR^m - rdt\right) + \frac{1}{\gamma^i}\left(\eta^{i\prime} - \eta^{m\prime}\right)dR^z
\end{aligned}$$

where again $dR^z = \beta_{dy,dR}\left(dR - rdt\right)$ are the returns on the mimicking portfolios for state-variable risk.

The investor first holds more or less of the *market* portfolio according to risk aversion. Then, he holds more or less of the mimicking portfolios for state variable risk according to whether the investor "feels" *differently* about these risks than does the average investor.

Some special cases of this portfolio advice are particularly interesting. First, return to the market portfolio in (70). The first term – and only the first term – gives the mean-variance efficient portfolio. Thus, *the market portfolio is no longer mean-variance efficient.* Referring to Figure 5, you can see that the optimal portfolio has slid down from the vertical axis of the nose-cone shaped multifactor efficient frontier. The average investor, and hence the market portfolio, gives up some mean-variance efficiency in order to gain a portfolio that better hedges the state variables.

This prediction is the source of much portfolio advice from multifactor models, for example, why the ICAPM interpretation of the Fama-French 3 factor model, is used as a sales tool for value-stock portfolios. If you find a mean-variance investor, an investor who does not fear the state variable changes and so has $\eta^i = 0$; this investors can now profit by deviating from market weights. He should slide up the nose-cone shaped multifactor efficient frontier in Figure 5, in effect selling state-variable insurance to other investors, and charging a fee to do so. His optimal portfolio is now

$$dR^i = rdt + \frac{\gamma^m}{\gamma^i}\left(dR^m - rdt\right) - \frac{\eta^{m\prime}}{\gamma^i}dR^z$$

he should sell the mimicking portfolio for *aggregate* state-variable risk. This expression tells us, *quantitatively*, how the mean-variance investor should deviate from the market portfolio in order to profit from the ICAPM. An estimate of the ICAPM will tell us the slope coefficients (of average returns on covariances) γ^m, η^m.

On the other hand, for everyone who is long someone else must be short. For every investor who wants to profit from, say, the value premium in this way, there must be an investor whose η^i is for example twice η^m, thus his optimal portfolio return is

$$dR^i = rdt + \frac{\gamma^m}{\gamma^i}\left(dR^m - rdt\right) + \frac{\eta^{m\prime}}{\gamma^i}dR^z$$

He sees the great news of the Sharpe ratio of value portfolios, but wants to sell not buy, since he is already too exposed to that state variable risk.

Third, this formulation includes idiosyncratic state variable risk, perhaps the most important (and overlooked) risk of all. $\eta = V_{Wy}/W_W$ is a *vector* and different for different y. We should include in investor i's problem state variables for his *individual* outside income risk, even though $\eta_m = 0$ for such risks. If, for example, this investor is no different from everybody else about his feelings towards *aggregate* state variables, then his optimal portfolio will be

$$dR^i = rdt + \frac{\gamma^m}{\gamma^i}\left(dR^m - rdt\right) + \frac{1}{\gamma^i}\eta^{i\prime}dR^z$$

This investor holds the market portfolio, plus a portfolio of assets that best offsets his individual outside income risks. (State variables for investment opportunities are by definition common to all investors.)

The hedge portfolios for individual risks, with $\eta^m = 0$ obtain no extra premium; η^i does not enter the ICAPM (71). Thus, *unpriced* mimicking portfolios are likely to be the most interesting and important for the average investor. The K funds that span the multifactor frontier are *not* likely to be the same for every individual when we incorporate idiosyncratic outside income risk into the analysis.

5.4 Completing the Merton model

The parameters γ and η governing the optimal portfolio come from the value function. The value function does not have a closed-form solution for cases other than the simple ones studied above, so, alas, this is where the analysis ends.

Again, we still need to compute the levels of risk aversion and "state variable aversion" from the primitives of the model, the utility function and formulas for the evolution of stock prices. Conceptually this step is simple, as before: we just need to find the value function. Alas, the resulting partial differential equation is so ugly that work on this multivariate model has pretty much stopped at the above qualitative analysis. From (54), the equation is

$$\begin{aligned} 0 &= u\left[u'^{-1}(V_W)\right] - \rho V + V_t + V_W W r - V_W u'^{-1}(V_W) + V_{y'}\mu_y + \frac{1}{2}Tr(\sigma_y' V_{yy'}\sigma_y) \\ &\quad + W\alpha^{*\prime}\left[(\mu - r)V_W + \sigma\sigma_y' V_{Wy}\right] + \frac{1}{2}V_{WW}W^2\alpha^{*\prime}\sigma\sigma'\alpha^* \end{aligned}$$

where

$$\alpha^* = -\frac{V_W}{WV_{WW}}\left(\sigma\sigma'\right)^{-1}(\mu - r) - \left(\sigma\sigma'\right)^{-1}\sigma\sigma_y'\frac{V_{Wy}}{WV_{WW}}$$

6 Portfolios with time-varying expected returns

Given that market returns are forecastable, e.g. from the dividend yield, let's try to find the optimal portfolio. This is a classic and interesting single-state variable Merton problem. How much market timing should one do? How strong is the "hedging demand" that makes stocks more attractive than they would be if returns were iid? This classic problem has been attacked by Kim and Ohmberg (1996), who find an exact (but difficult) solution with no intermediate consumption, Brennan, Schwartz and Lagnado (1997) who solve it numerically, Brandt (1999), Campbell and Vicera (1999), Wachter (1999a), Barberis (1999), who shows how and many others. (Cochrane 1999 attempts a summary.)

6.1 Payoff approach

Let μ_t denote the conditional mean return. Then the portfolio problem has a return with a time-varying mean, in which the state variable follows a continuous-time $AR(1)$.

$$\max E \int_0^\infty e^{-\rho t} u(c_t) dt \text{ or } \max E e^{-\rho T} u(c_T) \text{ s.t.}$$

$$\begin{aligned} dR - rdt &= \mu_t dt + \sigma dz_t \\ d\mu_t &= \phi(\bar{\mu} - \mu) dt + \sigma_{\mu z} dz_t + \sigma_{\mu w} dw_t \end{aligned}$$

and the usual budget constraint. The return and state-variable shocks are not in general perfectly correlated. I express the state-variable shock as a sum of two orthogonal elements to express this fact. This is a "complete-markets problem" when $\sigma_{\mu w} = 0$. In that case, shocks to μ_t can be perfectly hedged using asset returns dR. If $\sigma_{\mu w} \neq 0$, then, as we will see, the investor might want to hold a portfolio that loads on dw, and we have to impose the constraint that he cannot do that.

The optimal payoff is

$$\begin{aligned} e^{-\rho t} c_t^{-\gamma} &= \lambda \frac{\Lambda_t}{\Lambda_0} \\ c_t &= \left[e^{\rho t} \lambda \frac{\Lambda_t}{\Lambda_0} \right]^{-\frac{1}{\gamma}} \end{aligned}$$

Again, portfolio theory is done, but we have to solve the financial-engineering issue of finding the discount factor for this asset structure. Since there is an extra shock dw, the discount factor is

$$\frac{d\Lambda}{\Lambda} = -rdt - \frac{(\mu - r)}{\sigma} dz_t - \eta dw$$

The constant η is undertermined, since the dw shock does not correspond to any traded asset. This is the equivalent of adding an ε shock in Figure 3.

It's easier to define the Sharpe ratio as the state variable

$$x_t = \frac{\mu_t - r}{\sigma}$$

$$
\begin{aligned}
d\left(\frac{\mu - r}{\sigma}\right) &= \phi\left[\frac{\bar{\mu} - r}{\sigma} - \frac{\mu - r}{\sigma}\right] dt + \frac{\sigma_{\mu z}}{\sigma} dz + \frac{\sigma_{\mu w}}{\sigma} dw \\
dx_t &= -\phi\left(x_t - \bar{x}\right) dt + \sigma_{xz} dz_t + \sigma_{xw} dw_t \\
\frac{d\Lambda_t}{\Lambda_t} &= -r dt - x_t dz - \eta dw
\end{aligned}
$$

Now we can express the current Sharpe ratio as an AR(1)

$$
x_t - \bar{x} = \sigma_{xz} \int_0^T e^{-\phi s} dz_{t-s} + \sigma_{xw} \int_0^T e^{-\phi s} dw_{t-s} + e^{-\phi T}\left(x_0 - \bar{x}\right)
$$

and then the discount factor is

$$
\begin{aligned}
d \ln \Lambda_t &= \frac{d\Lambda}{\Lambda} - \frac{1}{2}\frac{d\Lambda^2}{\Lambda^2} = -\left(r + \frac{1}{2}x_t^2\right) dt - x_t dz \\
\ln \Lambda_t - \ln \Lambda_0 &= -\int_0^t \left(r + \frac{1}{2}x_s^2\right) ds - \int_0^t x_s dz_s
\end{aligned}
$$

$$
\begin{aligned}
d \ln \Lambda_t &= \frac{d\Lambda}{\Lambda} - \frac{1}{2}\frac{d\Lambda^2}{\Lambda^2} = -\left(r + \frac{1}{2}x_t^2 + \frac{1}{2}\eta^2\right) dt - x_t dz - \eta dw \\
\ln \Lambda_t - \ln \Lambda_0 &= -\int_0^t \left(r + \frac{1}{2}x_s^2 + \frac{1}{2}\eta^2\right) ds - \int_0^t x_s dz_s - \eta(w_T - w_0) \qquad (73)
\end{aligned}
$$

Now, the "complete" markets case requires that there is a single shock dz, so shocks to μ are perfectly correlated with shocks to dR, and thus the investor can perfectly hedge them. Markets are "incomplete" when this is not the case, because the investor would like to hedge risks to the investment opportunity set. Complete markets may not be that bad an approximation, since especially at high frequency discount rate shocks dominate changes in market prices. In the "complete" markets case, then, we have

$$
\begin{aligned}
x_t - \bar{x} &= \sigma_{xz} \int_0^T e^{-\phi s} dz_{t-s} + e^{-\phi T}\left(x_0 - \bar{x}\right) \\
\ln \Lambda_t - \ln \Lambda_0 &= -\int_0^t \left(r + \frac{1}{2}x_s^2\right) ds - \int_0^t x_s dz_s \qquad (74) \\
c_t &= \left[e^{\rho t} \lambda \frac{\Lambda_t}{\Lambda_0}\right]^{-\frac{1}{\gamma}} \qquad (75)
\end{aligned}
$$

This system is not quite so pleasant, though it is a well-studied class. (It's a "stochastic volatility" model often used to describe stock prices.) Even in the absence of a closed-form solution, however, you can straightforwardly simulate it and watch the discount factor and then optimal payoffs respond to the state variables

In the "incomplete" markets case, we're back to (73) together with (75). Our task is to pick the choice of η so that the final c_t is not driven by shocks dw, which is not so easy. It looks like $\eta = 0$ will do the trick, since then the explicit dependence of Λ on w in the last term of (73) vanishes. This would be enough for a quadratic utility investor, but not for power utility.

6.2 Portfolio approach

6.3 Quantitative evaluation

7 Wacky weights and Bayesian portfolios

Portfolio maximization depends on parameters, and uncertainty about the parameters can change the *estimate* of the optimal portfolio as well as provide some *standard errors* of that portfolio. The basic idea is, rather than maximizing conditional on parameters θ,

$$\max_{\{\alpha\}} \int u(\alpha' R_{t+1} W_0) p(R_{t+1}|\theta) dR_{t+1}$$

We integrate over the uncertainty about parameters as well,

$$\max_{\{\alpha\}} \int u(\alpha' R_{t+1} W_0) \left[\int p(R_{t+1}|\theta) p(\theta) d\theta \right] dR_{t+1} \tag{76}$$

This approach can usefully tame the wild advice of most portfolio calculations, and it advises you to place less weight on less well measured aspects of the data.

Warning: this section is even more preliminary than the rest of the chapter.

The parameters of any model of asset returns are uncertain. We don't know exactly what the equity premium is, we don't know exactly what the regression of returns on dividend yields and other predictor variables looks like, and we don't know exactly what the cross section of mean returns or the "alphas" in a factor model are. Now, in a statistical framework, we would suppose that uncertainty in the inputs just translates into standard errors on the outputs. It's interesting to track down how much uncertainty we have about optimal portfolio weights, but that consideration doesn't change the optimal portfolio itself.

However, it makes intuitive sense that one will take less advantage of a poorly estimated model, rather than just invest according to its point estimates. This intuition is correct, and Bayesian portfolio theory is a way to formalize it and think about *how much* to take advantage of a poorly estimated model. The central idea is that parameter uncertainty is also a risk faced by the investor. When some assets or strategies returns are less well estimated than others, this source of uncertainty can skew the portfolio weights as well as tilt the overall allocation towards less risky assets.

So far, when we have made a portfolio calculation, we maximized expected utility, treating the parameters as known. In a one-period problem, we solved

$$\max_{\{\alpha\}} \int u(\alpha' R_{t+1} W_0) p(R_{t+1}|\theta) dR_{t+1}$$

where θ denotes the parameters of the return distribution. But it is a mistake to treat the parameters as fixed. Uncertainty about the parameters is real uncertainty to the investor. Instead, we should integrate over all the parameter values as well, i.e.

$$\max_{\{\alpha\}} \int u(\alpha' R_{t+1} W_0) \left[\int p(R_{t+1}|\theta) p(\theta) d\theta \right] dR_{t+1} \tag{77}$$

The "predictive density"

$$p(R_{t+1}) = \int p(R_{t+1}|\theta) p(\theta) d\theta$$

expresses the true probability density of returns faced by the investor. This is a very intuitive formula. It says to generate the probability density of returns by integrating over all possible values of the parameters, using the information one has about the chance that those parameters are correct to weight the possibilities.

With these ideas in mind, Bayesian portfolio theory can capture three effects. First, it can capture the effect of parameter uncertainty on asset allocation. Here, we take a "diffuse prior", so that $p(\theta)$ reflects sampling error in the parameters. Second, it can let us mix "prior" information with "sample" information. Often, we have a new model or idea suggesting a portfolio strategy, but we also have a wealth of experience that alphas are pretty small. $p(\theta)$ now weights our "prior" information, really summing up all the other data we've looked at, relative to the statistical information in a new model, to interpolate between the two views. Third, an investor keeps learning as time goes by; each new data point is not only good and bad luck, but it also is extra information that changes the investor's probability assessments. Thus, the investor's shifting probability views become a "state variable." I only consider the first two effects here, though the third is also interesting.

7.1 Wacky weights

To motivate the Bayesian approach, let's just construct some optimal portfolios. We will see that implementing the portfolio formulas using sample statistics leads to dramatic overfitting of portfolio advice. The examples suggest that we find a way to scale back the overfitting to produce more sensible advice.

7.1.1 Optimal market timing

I calculate an approximate optimal portfolio that takes advantage of the predictability of returns from dividend yields. We see that it recommends a very strong market timing strategy, too strong really to be believed.

Given that variables such as the dividend yield forecast returns, how *much* should an optimal portfolio market-time, i.e. invest more in stocks when expected returns are higher? The complete solution to the Merton problem in this case, (55)

$$\alpha = -\frac{V_W}{WV_{WW}}\frac{\mu_t - r_t}{\sigma_t^2} - \frac{V_{Wy}}{WV_{WW}}\beta_{dy,dR}$$

is hard to calculate, especially if one wants the value function. However, we can evaluate the first, "market timing" component for fixed values of risk aversion, to get an idea of the strength of market-timing that a full solution will recommend.

Figures 7 and 6 present market timing rules based on a regression of returns on dividend yields, using the CRSP value-weighted index 1926-2007. The optimal amount of market timing is suspiciously large, especially for the $\gamma = 2 - 5$ range that seems reasonable for generating the overall allocation to stocks.

$$\begin{array}{ccccccc} R^e_{t+1} = & -7.20 & + & 3.75 & D/P_t & + & \varepsilon_{t+1} \\ (t) & (-1.20) & & (2.66) & & \sigma^2_\varepsilon = & 19.81\% \end{array} \quad (78)$$

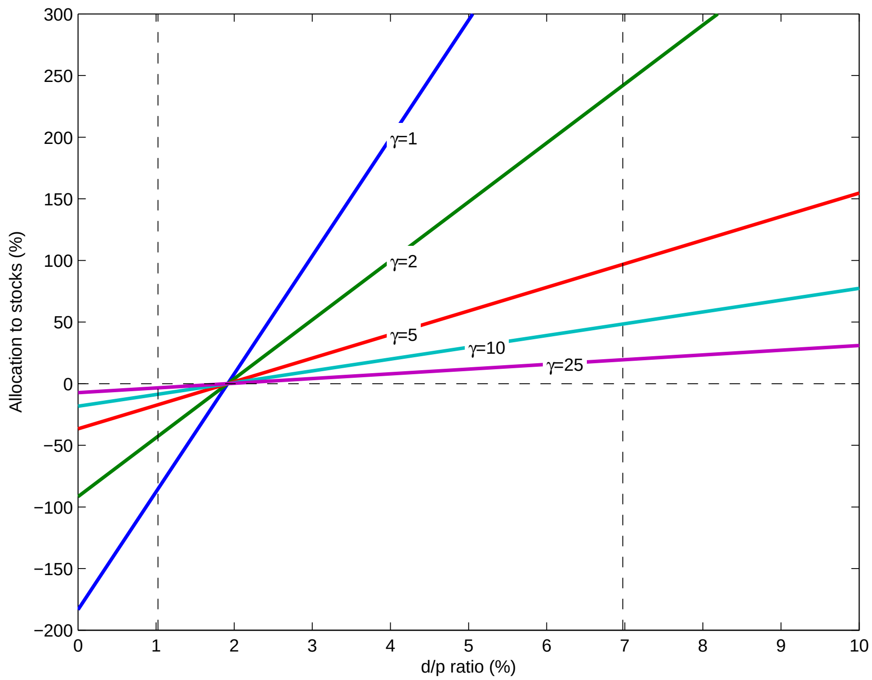

Figure 6: Market timing portfolio allocation. The allocation to risky stocks is $\alpha_t = \frac{1}{\gamma}\frac{E_t(R^e)}{\sigma^2}$ .Expected excess returns come from the fitted value of a regression of returns on dividend yields, $R^e_{t+1} = a + b(D_t/P_t) + \varepsilon_{t+1}$. The vertical lines mark $E(D/P) \pm 2\sigma(D/P)$.

In this case especially though, the *statistical* reliability of the regression gives one pause. The t statistic of only 2.55 in 80 years of data is not that large. Since the original dividend yield regressions which found coefficients of 5 or so in the mid 1980s, the boom of the 1990s despite low dividend yields cut the coefficient and t statistics below 2 by the late 1990s. Anyone who took the advice lost out on the 1990s boom. The market decline of the early 2000 made the strategy look somewhat better, and the coefficient has risen to 3.5 and regained its significance. But the literature on dividend yield predictability still worries about a "structural shift", that the still negative advice from this regression is too pessimistic. All of this just reflects natural hesitance to adopt the portfolio advice, and motivates a Bayesian approach to formally shading back the strong market timing of the figure.

This example is in fact a very *mild* one. I estimated parameters using data from 1926, which is much longer than the 5-10 year samples that are considered "long" when hedge funds estimate market timing strategies. I only considered one signal. You can imagine how quickly the overfitting problem explodes if one includes multiple variables on the right hand side, generating much stronger portfolio advice.

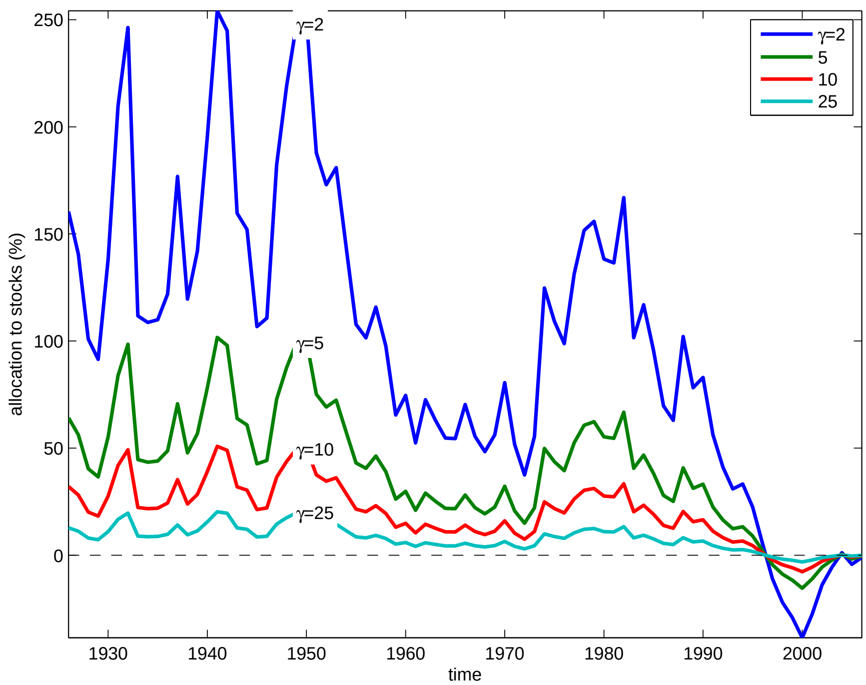

Figure 7: Market timing portfolio allocation over time. The allocation to risky stocks is $\alpha_t = \frac{1}{\gamma}\frac{E_t(R^e)}{\sigma^2}$ .Expected excess returns come from the fitted value of a regression of returns on dividend yields, $R^e_{t+1} = a + b(D_t/P_t) + \varepsilon_{t+1}$.

7.1.2 A simple cross-section

I calculate a very simple mean-variance frontier, and we see that the maximizer recommends "wacky weights" with very strong long and short positions.

To evaluate how a typical mean-variance problem might work in practice, I estimate the mean and covariance matrix of the 25 Fama-French portfolios and 3 factors in 20 years of data. This too is a small and well-behaved problem by real-world standards. 5 years is a long estimation period, and hedge funds consider hundreds of assets, along with time-varying means and covariance matrices, which all make the problems much worse. Figure 8 presents the results.

I transform all the returns to excess returns, so we are only considering the composition of the risky portfolio. The weights w in the portfolio $w'R^e$ are

$$w = \frac{1}{\gamma}\Sigma^{-1}\mu$$

where μ and Σ are the mean and covariance matrix of excess returns. The composition of the optimal portfolio is the same for all γ, I choose the scale to report so that the variance of the resulting portfolio is the same as the variance of the market index,

$$w = \frac{\sigma^2(rmrf)}{\mu'\Sigma^{-1}\mu}\Sigma^{-1}\mu$$

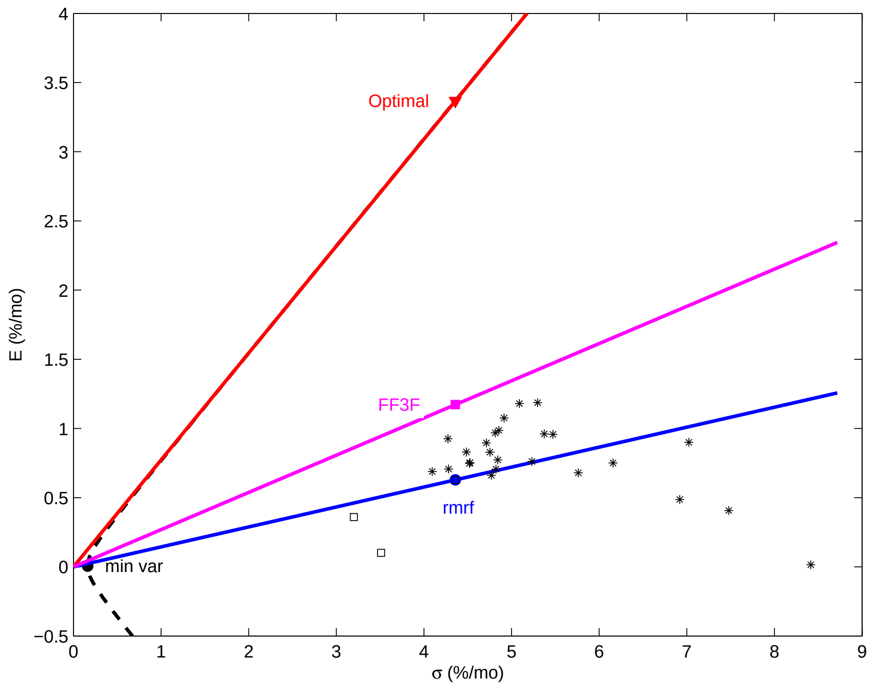

Figure 8: mean-variance optimization with excess returns of the Fama French 25 size and book/market portfolios, together with the 3 Fama French factors. "Optimal" is the mean-variance optimal portfolio, at the same variance as the market return. "FF3F" is the mean-variance optimal combination of the 3 Fama-French factors. The stars are the means and standard deviations of the individual portfolio returns. The squares are the means and variances of the factors. The thin black line gives the mean-variance frontier when weights sum to one. Min var is the minimum variance excess return with weights that sum to one. Mean and covariance estimates based on 20 years of data.

With this choice, $\sigma^2(w'R^e) = w'\Sigma w = \sigma^2(rmrf)$. This portfolio is graphed at the "optimal" point in Figure 8. The "FF3F" point gives the mean-variance optimal combination of the 3 Fama-French factors, and the "rmrf" point gives the market return, whose Sharpe ratio gives the equity premium so much trouble.

The figure suggests that portfolio optimization can deliver huge gains in mean returns and Sharpe ratios. Why earn the market's measly 0.5%/month when 3.5%/month is available with no increase in volatility? However, let's take a look at the actual portfolio the optimizer is recommending, in Table 1 below:

	low	2	3	4	high
small	-149	51	69	96	52
2	-19	-57	190	-13	-60
3	29	-34	-31	-93	41
4	116	-39	-42	35	-2
large	87	-19	8	-22	2

rmrf	hml	smb
-94	77	-69

Table 1. Optimal percent portfolio weights $w = \frac{\sigma^2(rmrf)}{\mu'\Sigma^{-1}\mu}\Sigma^{-1}\mu$ in the Fama-French 25 size and book/market portfolios plus the 3 Fama French Factors.

These are "wacky weights!" 149% short position in small growth, then +190% in the (2,3) portfolio, and -93% in the (3,4) portfolio, seem like awfully large numbers. A -99% investment in the market portfolio seems awfully strong as well. One is very hard-pressed to take these numbers seriously.

Figure 8 also shows the minimum variance portfolio. This is the portfolio with minimum variance, subject to the constraint that all weights sum to one. It has weights

$$w = \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}'\Sigma^{-1}\mathbf{1}}$$

This portfolio is useful to see troubles with the covariance matrix by itself, since it does not involve means. If we believe this portfolio, by clever long and short positions among the Fama French 25 returns and the 3 factors, one can synthesize a nearly riskfree portfolio!

A moment's reflection shows the problem. The assets are highly correlated with each other. Thus, small, and perhaps even insignificant, differences in average returns can show up as huge differences in portfolio weights. Figure 9 illustrates the situation. Suppose even that securities A and B are identical, in that they have the same true mean, standard deviation, and betas. They are also highly correlated. In a sample, A may be lucky and B unlucky. Given the strong correlation between A and B, the mean-variance frontier connecting them is very flat, resulting in an optimal portfolio that is very long A and very short B.

The minimum-variance portfolio is an instance of the same problem, in which A and B don't have the same mean. Now the line connecting them passes very close to the riskfree rate.

7.1.3 Common "solutions"

Again, all these problems are much worse in real-world applications. One of the most frequent uses of portfolio theory is to incorporate subjective views ("trader skill") about individual securities into an overall portfolio. You can imagine how the portfolio starts to jump around when a trader can express views about individual alphas.

These problems have been with portfolio optimization for 50 years. Industry practice has developed a lot of common-sense procedures to guard against them. For example, most portfolio

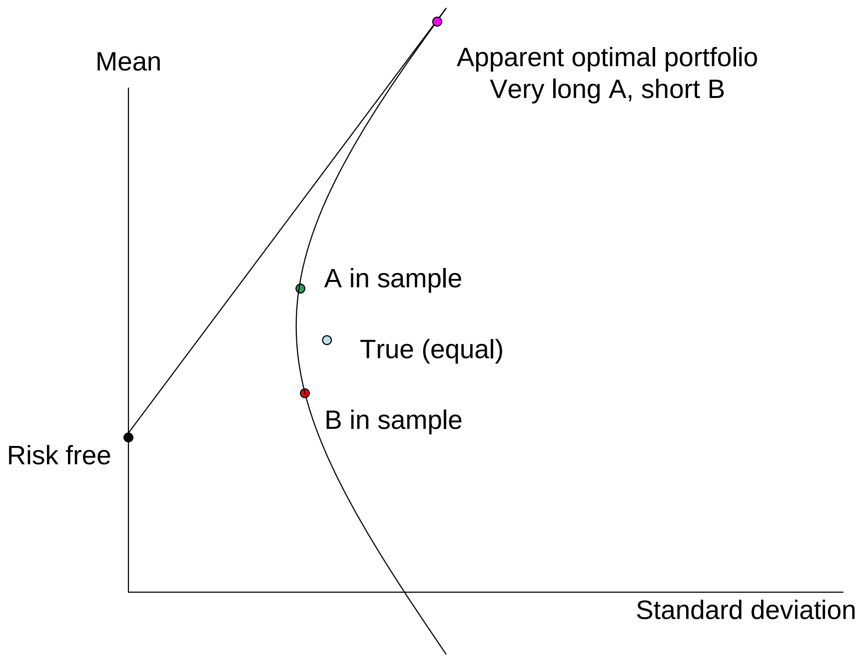

Figure 9: Example that generates huge portfolio weights

optimization adopts short-sale constraints or at least limitations on the extent of short-selling. Most professional portfolio managers also adopt (or are forced to adopt) a long list of additional, seemingly arbitrary rules: don't invest more than x percent in a single security, make sure the quantities are diversified across industry categories, etc.

Figure 9 suggests that most of these ad-hoc solutions have their limits. In that example merely going 100% long A and 0 B is better than the extreme long-short position, but it does not recover the true answer 50% A and 50% B. Similarly, most portfolio rules must reflect doubt that the portfolio maximization has done its job correctly. If so, it seems better to figure out why the portfolio maximizer isn't working than to use it but slap ad-hoc rules on the results.

Both examples derive ultimately from the fact that we don't know the parameters generating asset returns precisely. Small errors in those parameters give rise to nonsense portfolios. The right answer, it seems, is to understand parameter uncertainty and to incorporate it into the portfolio model.

7.2 Parameter uncertainty

With that motivation, let's adapt the ideas surrounding equation (77) to some simpler portfolio calculations.

7.2.1 Lognormal iid power allocation

Parameter uncertainty acts like an additional risk to investors. With normal distributions, we simply add the mean return standard error to the variance of returns; the investor acts as if returns are $R_{t+1} \sim \mathcal{N}\left[\bar{\mu}, \sigma^2 + \sigma_\mu^2\right]$ where $\bar{\mu}$ is the mean of his estimate of the mean returns, and σ_μ^2 is the variance of that estimate.

I solve a simple lognormal iid power allocation problem, and find that the allocation to the risky asset is $\alpha = (\bar{\mu} - r)/\gamma\sigma^2\left[1 + (\gamma - 1)h/T\right]$ where h is the investor's horizon and T is the data sample used to estimate the mean return. Parameter uncertainty is more important for long run investors. I present some evaluations of this formula, which suggest that the phenomenon is quantitatively important.

In the classic lognormal, power, iid world, the allocation to stocks is

$$\alpha = \frac{1}{\gamma}\frac{\mu_t - r_t}{\sigma_t^2}$$

For typical numbers $\mu - r = 8\%$ and $\sigma = 16\%$, with a market Sharpe ratio of 0.5,

$$\frac{0.08}{0.16^2} = \frac{0.08}{0.0256} = 3.123$$

so an investor with risk aversion 3.125 should put all his portfolio in stocks. The conventional 60/40 allocation happens at a risk aversion of about 5.2. These seem like reasonable numbers, at least so long as one ignores the implication that consumption growth should have the same volatility as stock returns. This fraction is invariant to horizon if the investor desires terminal wealth.

Now, the equity premium is notoriously hard to estimate. $\sigma = 16\%$ means that in 50 years of data, our uncertainty about the sample mean is $\sigma/\sqrt{T} = 2.26$ percentage points. At a standard error level, this means, roughly, that the one-standard error band for our optimal allocation is between

$$\frac{1}{\gamma}\frac{0.08 \pm 0.026}{0.16^2} = \frac{1}{\gamma}\left(2.11 - 4.14\right).$$

For $\gamma = 3.125$, that is the range 66%-132%. This is a large band of uncertainty, which should leave one uncomfortable with the calculation. (I have always marveled at the common practice of carefully rebalancing when allocations are only a few percent off their targets.) Still, this calculation does not capture the idea that uncertainty about the market premium might lead us to take *less* market risk overall.

To do that, let's implement a simple version of the Bayesian calculation. The parameter that really matters here is μ. Let us then specify that R is normally distributed with mean μ and variance σ^2, and that our posterior about μ is also normally distributed with mean $\bar{\mu}$ and variance σ_μ^2. For example, we could take $p(\mu)$ as a normal distribution centered around the point estimate $\bar{\mu}$, using the standard error to calibrate the variance σ_μ^2. By collecting terms and expressing the integral in normal form, one can show that

$$\begin{aligned} p(R_{t+1}) &= \int p(R_{t+1}|\mu)p(\mu)d\mu \\ &= \frac{1}{\sqrt{2\pi\sigma^2}\sqrt{2\pi\sigma_\mu^2}}\int e^{-\frac{1}{2}\frac{(r-\mu)^2}{\sigma^2}}e^{-\frac{1}{2}\frac{(\mu-\bar{\mu})^2}{\sigma_\mu^2}}d\mu = \frac{1}{\sqrt{2\pi\left(\sigma^2+\sigma_\mu^2\right)}}e^{-\frac{1}{2}\frac{(r-\bar{\mu})^2}{(\sigma^2+\sigma_\mu^2)}} \qquad (79) \end{aligned}$$

Thus, we have

$$R_{t+1} \sim \mathcal{N}\left[\bar{\mu}, \sigma^2 + \sigma_\mu^2\right] = \mathcal{N}\left[\bar{\mu}, \sigma^2(1+1/T)\right] \quad (80)$$

This is an important and very intuitive result. *Parameter uncertainty about the mean means we add the parameter variance to our estimate of the return variance.* Parameter uncertainty about the *mean* return does not lower one's estimate of the mean. Instead, that parameter uncertainty raises the perceived *riskiness* of the return. Plugged into a mean-variance calculation, you can see that the investor will act more risk averse than ignoring parameter uncertainty.

Equation (80) suggests that parameter uncertainty is more important for investors with longer horizons. This makes sense. First of all, uncertainty about the mean is much more important than uncertainty about variances, which in theory at least collapses to zero as you observe data at finer intervals. Then, at short horizons, risk comes almost entirely from the variance of returns, while at 10 year or longer horizons, uncertainty about the mean has more of a chance to compound. For example, if the annualized mean market return has a 5 percentage point standard error, this adds only $5/365 = 0.0137$ percentage points to daily volatility, a tiny contribution compared to a $20/\sqrt{365} = 1.0$ percentage point daily standard deviation of returns. However, at a 10 year horizon, $10 \times 5 = 50$ percentage points is a much larger amount of risk compared to $20 \times 10 = 200\%$ 10 year volatility. Fundamentally, the fact that means scale with horizon but standard deviations scale with the square root of horizon makes uncertainty about the mean more important for longer horizons.

We can see this effect to some extent in (80). As we apply the formula to longer horizons, the number of nonoverlapping intervals T that goes in the standard error formula shrinks. Thus, to a good approximation, we can use

$$R_{t+h} \sim \mathcal{N}\left[\bar{\mu}, \sigma^2(1+h/T)\right]. \quad (81)$$

It's more satisfying to be a bit more explicit both about the horizon and about the portfolio optimization, so here is a concrete example. Let's find the optimal portfolio for a power utility investor in an iid world with horizon h. I'll presume a constantly rebalanced fraction α in the risky asset, which we know from above is the answer when we allow α_t to be freely chosen. (I am ignoring the learning effect that σ_μ^2 declines over the investment horizon.) Given the portfolio choice α, wealth evolves as

$$\begin{aligned}
\frac{dW}{W} &= rdt + \alpha\left(\mu - r\right)dt + \alpha\sigma dz \\
d\ln W &= \left[r + \alpha(\mu - r) - \frac{1}{2}\alpha^2\sigma^2\right]dt + \alpha\sigma dz \\
W_h &= W_0 e^{\left[r+\alpha(\mu-r)-\frac{1}{2}\alpha^2\sigma^2\right]h+\sigma\alpha\sqrt{h}\varepsilon};\ \varepsilon\tilde{\ }N(0,1)
\end{aligned}$$

The investor's problem is

$$\max_\alpha E\left(\frac{W_T^{1-\gamma}}{1-\gamma}\right)$$

$$\max_\alpha \frac{W_0^{1-\gamma}}{1-\gamma} E\left(e^{(1-\gamma)\left[r+\alpha(\mu-r)-\frac{1}{2}\alpha^2\sigma^2\right]h+(1-\gamma)\alpha\sigma\sqrt{h}\varepsilon}\right)$$

Now, the investor is unsure about the return shock ε (of course), but *also about the mean* μ. You can see directly that μ and ε are sources of risk that act similarly in the objective. The two sources of risk are independent – uncertainty about the mean comes from the past sample, and uncertainty about returns comes from the future. (Again, I am ignoring the fact that the investor learns a bit more about the mean during the investment period.) Thus, the problem is

$$\max_{\alpha} \frac{W_0^{1-\gamma}}{1-\gamma} \int e^{(1-\gamma)\left[r+\alpha(\mu-r)-\frac{1}{2}\alpha^2\sigma^2\right]h+(1-\gamma)\alpha\sigma\sqrt{h}\varepsilon} f(\varepsilon)f(\mu)d\varepsilon d\mu$$

Doing the conventional ε integration first,

$$\max_{\alpha} \frac{W_0^{1-\gamma}}{1-\gamma} \int e^{(1-\gamma)\left[r+\alpha(\mu-r)-\frac{1}{2}\alpha^2\sigma^2\right]h+\frac{1}{2}(1-\gamma)^2\alpha^2\sigma^2 h} f(\mu)d\mu$$

$$\max_{\alpha} \frac{W_0^{1-\gamma}}{1-\gamma} \int e^{(1-\gamma)\left[r+\alpha(\mu-r)-\frac{1}{2}\gamma\alpha^2\sigma^2\right]h} f(\mu)d\mu$$

Let's model the investor's uncertainty about μ also as normally distributed, with mean $\bar{\mu}$ (sample mean) and standard deviation σ_μ (standard error). Then we have

$$\max_{\alpha} \frac{W_0^{1-\gamma}}{1-\gamma} e^{(1-\gamma)\left[r+\alpha(\bar{\mu}-r)-\frac{1}{2}\gamma\alpha^2\sigma^2\right]h+\frac{1}{2}(1-\gamma)^2\alpha^2\sigma_\mu^2 h^2}$$

The actual maximization is a little anticlimactic. Taking the derivative with respect to α, and canceling $W_0^{1-\gamma}e^{(..)}$ we have

$$(1-\gamma)\left[(\bar{\mu}-r)-\gamma\alpha\sigma^2\right]h+(1-\gamma)^2\alpha\sigma_\mu^2 h^2=0$$

$$\left[\gamma\alpha\sigma^2\right]-(1-\gamma)\alpha\sigma_\mu^2 h=(\bar{\mu}-r)$$

and thus finally

$$\alpha=\frac{\bar{\mu}-r}{\gamma\sigma^2+(\gamma-1)\sigma_\mu^2 h}.$$

If the variance of μ comes from a standard error $\sigma_\mu=\sigma/\sqrt{T}$in a sample T, then

$$\alpha=\frac{\bar{\mu}-r}{\gamma\sigma^2\left[1+\frac{\gamma-1}{\gamma}\frac{h}{T}\right]} \tag{82}$$

You can see in these formulas the special case $\alpha=(\mu-r)/\left(\gamma\sigma^2\right)$ we recovered above when there is no parameter uncertainty. Again, parameter uncertainty adds to the risk that the agent faces. The extra risk is worse as horizon increases, almost exactly[2] as in the simple calculation (81). You also can see that the effect disappears for log utility, $\gamma=1$, one of the many special properties of that utility function.

To give a quantitative evaluation, I reproduce the setup in Barberis's (2000) Figure 1. Barberis studies investors with horizons from 1 to 10 years, in a 43 year and a 11 year long data set, and risk aversion of 5 and 10. Figure 10 presents the allocation to stocks for this case as a function of horizon, calculated using (82) and using the mean and standard deviation of returns[3] given in Barberis' Table 1. The figure is almost identical to Barberis' figure 1. As in Barberis' calculations,

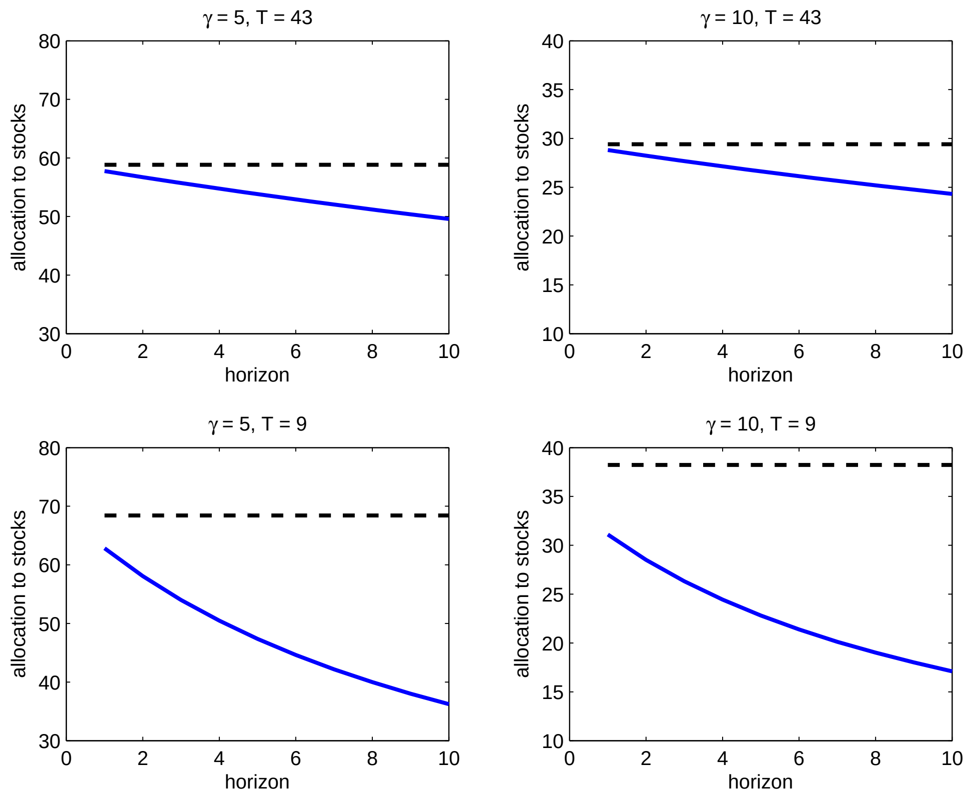

Figure 10: Portfolio allocation to stocks with parameter uncertainty. The solid lines present the case with parameter uncertainty, and the dashed lines ignore parameter uncertainty. The allocation to stocks is $\alpha = \frac{\mu - r}{\sigma^2[\gamma + (\gamma-1)h/T)]}$.$\mu = 0.06, \sigma = 0.1428$ for $T = 43$ and $\mu = 0.078, \sigma = 0.151$ for $T = 9$ with $r = 0$.

this calculation challenges the standard result that the allocation to stocks is the same for all horizons.

This calculation is tremendously simplified of course. Real Bayesian portfolio theorists will derive the parameter density $f(\theta)$ from a prior and a likelihood function; they include estimation uncertainty about variances as well as means, and they calculate.

7.2.2 Optimal market timing

Applying the rule $R_{t+1} \sim \mathcal{N}\left[\bar{\mu}, \sigma^2 + \sigma_\mu^2\right]$ in a very back-of-the envelope manner, I show that parameter uncertainty substantially tames the strong market timing prediction from return forecastabil-

[2] I'm still not exactly sure why (81) and (82) differ at all, however, one reason for the "preliminary" disclaimer.

[3] I don't exactly match the results without learning, which is a bit of a puzzle since $\alpha = (\mu - r)/\left(\gamma\sigma^2\right)$ is uncontroversial. Also, I had to set $r = 0$ to get even a vaguely similar calculation when using the μ and σ from Barberis' Table 1.

ity using dividend yields. Conditional mean returns that are much different from the unconditional mean are less well measured, so an optimal portfolio takes less notice of them.

I implement a very simple and back-of-the-envelope calculation by applying the same formula as above. At each date the conditional mean return is

$$E_t\left(R^e_{t+1}\right) = \hat{\alpha} + \hat{b}(D_t/P_t)$$

Therefore, the sampling uncertainty about the conditional mean return is

$$\sigma^2\left[E_t\left(R^e_{t+1}\right)\right] = \sigma^2(\hat{a}) + \sigma^2(\hat{b})\left(D_t/P_t\right)^2 + 2cov(\hat{a},\hat{b})D_t/P_t \tag{83}$$

Figure 11 shows the conditional mean return with one and two standard error bounds. The most important point is that the standard error bounds widen as the dividend yield departs from its mean. Since we are more *uncertain* about mean returns further away from the center, the optimal allocation that includes parameter uncertainty will likely place less weight on those means, resulting in a flatter optimal-allocation line.

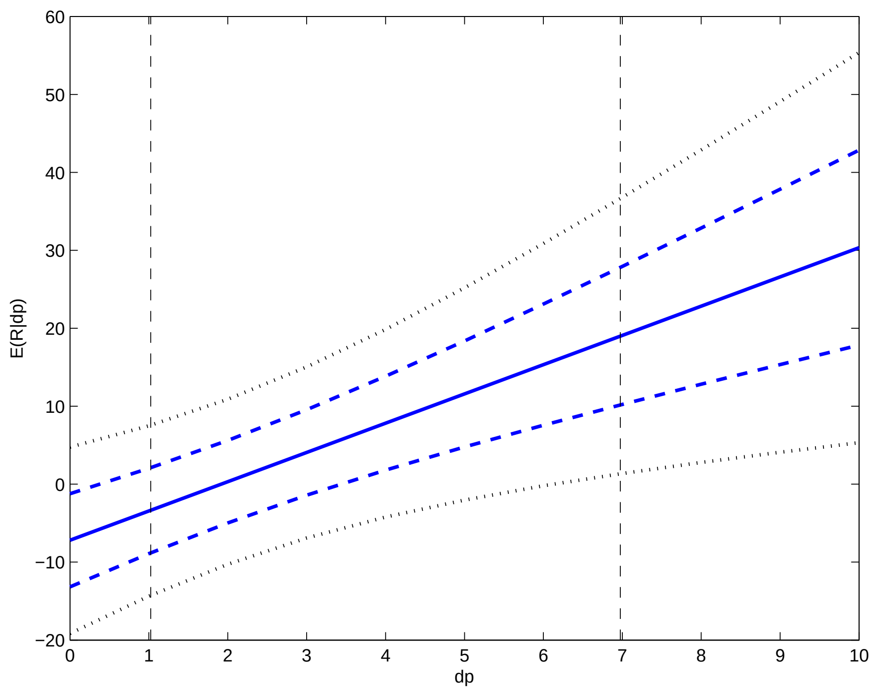

Figure 11: Expected excess returns as a function of dividend yield, with one and two standard error bands. Vertical lines are the mean dividend yield plus and minus two standard deviations.

Figure 12 presents the optimal allocation. The line without parameter uncertainty, marked "none" presents the standard allocation result for $\gamma = 5$, based on the regression (78)

$$\alpha = \frac{1}{\gamma}\frac{E(R^e_{t+1}|D_t/P_t)}{\sigma^2(\varepsilon)} = \frac{1}{\gamma}\frac{\hat{a} + \hat{b}(D_t/P_t)}{\sigma^2(\varepsilon)}.$$

This is the same as the $\gamma = 5$ line of Figure 7 and shows the same strong market-timing. The remaining lines include parameter uncertainty. Each line graphs

$$\alpha = \frac{1}{\gamma} \frac{\hat{a} + \hat{b}(D_t/P_t)}{\sigma^2(\varepsilon) + h \times \sigma^2 \left[E_t \left(R^e_{t+1}\right)\right]}$$

using (83) to calculate the bottom variance for each investment horizon h, as indicated in the graph. The results are visually quite similar to the much more complex calculation presented by Barberis (2000). (In particular, compare it to the presentation of Barberis' results in Figure 3 of Cochrane (2000) "New facts in finance")

In sum, Figure 12 says that investors with decently long horizons should substantially discount the market timing advice of dividend yield regressions, because of parameter uncertainty, and in particular because expected returns are *more* uncertain the further the dividend yield is from its mean.

Of course, this calculation is much too simplified. I assumed a "buy and hold investor" who does not change his allocation to stocks as the dividend yield varies through his investment life, I ignored the "learning" effect, and my parameter uncertainty does not come from a properly specified prior and likelihood.

Kandel and Stambaugh (1996) study a monthly horizon and come to the opposite conclusion: the optimal allocation based on the dividend yield is almost completely independent of parameter uncertainty. They conclude that the "economic" significance of dividend yield forecastability is much larger than its statistical significance, because optimal portfolios market-time quite strongly despite the poor statistical significance. As the figure shows, the apparent difference between Kandel and Stambaugh and Barberis is just a matter of horizons.

7.3 Incorporating prior information

The second, and perhaps more important, use of Bayesian portfolio theory, is to incorporate prior information into the analysis. When you see a new model, it often indicates aggressive portfolio advice. But the new model does not capture the wealth of information one has that random walks and the CAPM are decent approximations. Obviously, one wants to merge these two pieces of information, creating a sensible compromise between the new model and ancient wisdom. Furthermore, one should weight the new model more or less according to how well measured its recommendations are. Bayesian portfolio theory allows one to do this.

7.3.1 A simple model

We can tame "wacky weights" by shading the inputs back to a "prior" that alphas (relative to any chosen asset pricing model" are not as large as an estimate may suggest. I develop the formula $E(\alpha|\hat{\alpha}, \alpha_p) = \left(\frac{\hat{\alpha}}{\sigma^2(\hat{\alpha})} + \frac{\alpha_p}{\sigma^2(\alpha_p)}\right) / \left(\frac{1}{\sigma^2(\hat{\alpha})} + \frac{1}{\sigma^2(\alpha_p)}\right)$ where $\hat{\alpha}$ is an estimate, and α_p is a prior (typically zero) with confidence level expressed by $\sigma^2(\alpha_p)$. Sensibly, one weights evidence $\hat{\alpha}$ more strongly the better measured it is.

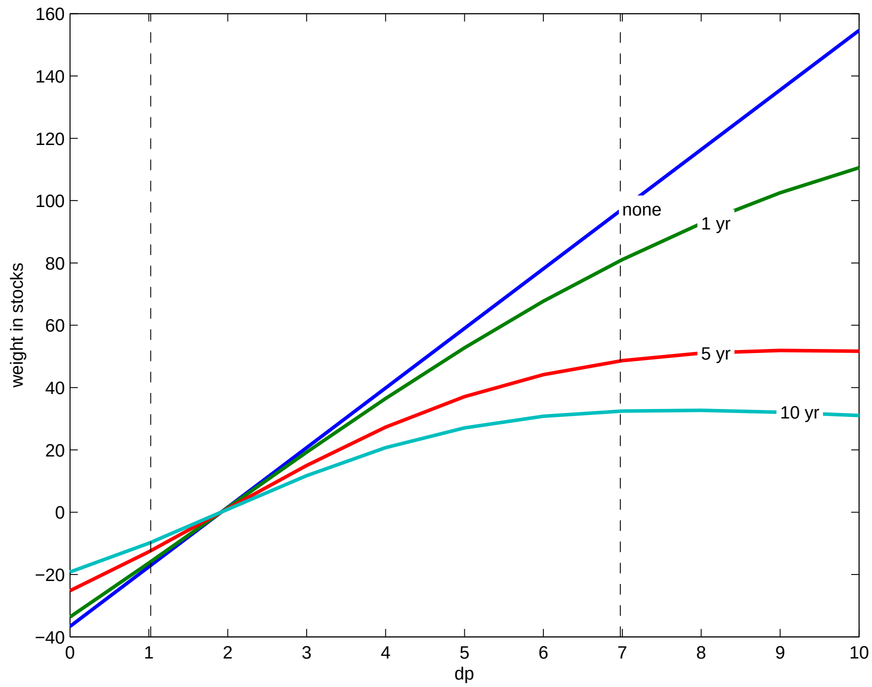

Figure 12: Optimal allocation to stocks given that returns are predictable from dividend yields, and including parameter uncertainty, for $\gamma = 5$.

Here is a very simplified version of the calculation. Suppose you have two uncorrelated normally distributed signals about the same quantity, α_1, with standard error σ_1 and α_2 with standard error σ_2. You can call α_1 the "prior" and σ_{α_1} your confidence in the prior, and α_2 the "estimate, " with σ_2 the standard error of the estimate. How do we combine these two signals to get the best estimate of α? The answer is (algebra below)

$$E(\alpha|\alpha_1, \alpha_2) = \frac{\frac{\alpha_1}{\sigma_1^2} + \frac{\alpha_2}{\sigma_2^2}}{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}} \tag{84}$$

Sensibly, you create a linear combination of the two signals, weighting each in inverse proportion to its variance. If α is a vector, the general case reads

$$E(\alpha|\alpha_1, \alpha_2) = \left(\Sigma_1^{-1} + \Sigma_2^{-1}\right)^{-1} \left(\Sigma_1^{-1}\alpha_1 + \Sigma_2^{-1}\alpha_2\right).$$

A natural application of course is that the "prior" is $\alpha_1 = 0$, with a common confidence level, $\Sigma_1 = \sigma_1 I$. This specification further simplifies the formulas.

Derivation

An easy "frequentist" way to derive (84)is to think of the signals α_1 and α_2 as generated from the true α,

$$\begin{aligned} \alpha_1 &= \alpha + \varepsilon_1 \\ \alpha_2 &= \alpha + \varepsilon_2 \end{aligned}$$

with ε_1 and ε_2 independent but having different variances σ_1 and σ_2. This is a simple instance of GLS with two data points – estimate alpha with two data points α_1 and α_2. The GLS formula is

$$\left(X'\Omega^{-1}X\right)^{-1}X'\Omega^{-1}Y$$
$$\left(\left[\begin{array}{cc} 1 & 1 \end{array}\right]\left[\begin{array}{cc} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{array}\right]^{-1}\left[\begin{array}{c} 1 \\ 1 \end{array}\right]\right)^{-1}\left[\begin{array}{cc} 1 & 1 \end{array}\right]\left[\begin{array}{cc} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{array}\right]^{-1}\left[\begin{array}{c} \alpha_1 \\ \alpha_2 \end{array}\right]$$
$$\left(\frac{1}{\sigma_1^2}+\frac{1}{\sigma_2^2}\right)^{-1}\left(\frac{\alpha_1}{\sigma_1^2}+\frac{\alpha_2}{\sigma_2^2}\right)$$

If $\alpha, \alpha_1, \alpha_2$ are $N \times 1$ vectors, the same GLS formula with $2N$ data points is

$$\left(\left[\begin{array}{cc} 1 & 1 \end{array}\right]\left[\begin{array}{cc} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{array}\right]^{-1}\left[\begin{array}{c} 1 \\ 1 \end{array}\right]\right)^{-1}\left[\begin{array}{cc} 1 & 1 \end{array}\right]\left[\begin{array}{cc} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{array}\right]^{-1}\left[\begin{array}{c} \alpha_1 \\ \alpha_2 \end{array}\right]$$
$$\left(\Sigma_1^{-1}+\Sigma_2^{-1}\right)^{-1}\left(\Sigma_1^{-1}\alpha_1+\Sigma_2^{-1}\alpha_2\right)$$

If you're not familiar with GLS, it's easy to find the scalar case directly by minimizing variance. Choose w to

$$\min_w var(w\alpha_1+(1-w)\alpha_2-\alpha)$$
$$\min_\lambda w^2\sigma_1^2+(1-w)^2\sigma_2^2$$

$$\begin{aligned} w\sigma_1^2 &= (1-w)\sigma_2^2 \\ w &= \frac{\sigma_2^2}{\sigma_1^2+\sigma_2^2} = \frac{\frac{1}{\sigma_1^2}}{\frac{1}{\sigma_1^2}+\frac{1}{\sigma_2^2}} \end{aligned}$$

Naturally, these formulas have more precise Bayesian derivations. Pastor (2000, p. 191, Eq 25-28) considers a standard factor model

$$R_t^e = \alpha + \beta F_t + u_t$$

where R^e is an excess return, and F is a factor (e.g. the market portfolio) which is also an excess return. He shows that the posterior mean alpha is

$$\tilde{\alpha} = w\alpha_{\text{sample}} + (1-w)\alpha_{\text{prior}}$$

$$w = \frac{\frac{1}{(\sigma_u^2/T)}}{\frac{1}{\sigma_\alpha^2}\left(1+\frac{E(F)^2}{var(F)}\right)+\frac{1}{(\sigma_u^2/T)}}$$

Here, $E(F), var(F)$ are sample moments of the factor, σ_u^2 is the (prior expected value of the) residual variance, and σ_α^2 is the variance of the prior about α. You recognize the same formula as in (84) but with a squared sharpe ratio of the factor adjustment, familiar from the Shanken correction to standard errors. As with the Shanken correction, this will be small in typical monthly data, but worth including. Pastor shows that betas also have a small Bayesian adjustment, basically resulting from the question whether one should include an intercept in the regression.

7.3.2 Covariances

Sample covariance matrices show spurious cross-correlation, especially for many assets and few time periods. I review simple techniques for solving this problem, including enhancing the diagonal elements, and using factor models of residual covariance.

Though in theory we learn covariances with arbitrary precision by looking at finely sampled data, in practice covariance matrices are also hard to estimate. This is especially true with large numbers of assets and small time spans. A covariance matrix can't have more than T degrees of freedom so a 200×200 covariance matrix estimated in anything less than 201 data points is singular by construction, showing spurious risk-free arbitrages. Even when this limit is not reached, sample covariance matrices tend to show "too much" correlation, weird-looking linear combinations of assets that appear nearly riskfree.

A common "Bayesian" solution to this problem is to emphasize the diagonals, i.e. to use

$$\Sigma = \hat{\Sigma} + \lambda D$$

where D is a diagonal matrix.

More generally, we usually estimate large covariance matrices by imposing a factor structure,

$$\begin{aligned} R_t &= \beta f_t + \varepsilon \\ (N \times 1) &= (K \times N)\,(K \times 1) + (N \times 1)\,;\ K << N \\ cov(RR') &= \beta \Sigma_f \beta' + \Sigma \end{aligned}$$

Then, more sensible return covariance matrices emerge by downweighting Σ, imposing $\Sigma = D$, etc.

7.4 Additional Literature and a paradox

Formula 80 is in Klein and Bawa 1976, Theorem 1. They use a more standard "noninformative prior", they consider variance estimation as well, with the result that the return distribution has a t distribution with $T-$(number of assets) degrees of freedom rather than normal.

Parameter uncertainty still can affect portfolio choice because the investor will *learn* more about the process as time goes by. This effect makes the parameter estimates a "state variable" in a Mertonian context, if utility is not logarithmic. Breannan (1998) studies this effect in a static context and Xia (2001)studies it in the classic single variable return forecasting context. In my example, I turned this effect off by specifying that the investor "learns" only at the terminal date, not as the investment proceeds.

The covariance matrix shrinkage literature is huge. Ledoit and Wolf (2003, 2004) is a good recent contribution.

7.5 Algebra

Algebra for (79) – though there must be an easier way!

$$\frac{(r-\mu)^2}{\sigma^2} + \frac{(\mu-\bar{\mu})^2}{\sigma_\mu^2}$$

$$\frac{r^2 - 2r\mu + \mu^2}{\sigma^2} + \frac{\mu^2 - 2\mu\bar{\mu} + \bar{\mu}^2}{\sigma_\mu^2}$$

$$\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)\mu^2 - 2\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)\mu + \left(\frac{r^2}{\sigma^2}+\frac{\bar{\mu}^2}{\sigma_\mu^2}\right)$$

$$\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)\left[\mu^2 - 2\frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\mu + \left(\frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\right)^2\right] - \frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)^2}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)} + \left(\frac{r^2}{\sigma^2}+\frac{\bar{\mu}^2}{\sigma_\mu^2}\right)$$

$$\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)\left[\mu - \frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\right]^2 - \frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)^2}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)} + \left(\frac{r^2}{\sigma^2}+\frac{\bar{\mu}^2}{\sigma_\mu^2}\right)$$

$$\frac{1}{\sqrt{2\pi\sigma^2}\sqrt{2\pi\sigma_\mu^2}}\int e^{-\frac{1}{2}\left[-\frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)^2}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}+\left(\frac{r^2}{\sigma^2}+\frac{\bar{\mu}^2}{\sigma_\mu^2}\right)\right]} e^{-\frac{1}{2}\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)\left[\mu-\frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\right]^2} d\mu$$

$$\frac{\sqrt{2\pi\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)^{-1}}}{\sqrt{2\pi\sigma^2}\sqrt{2\pi\sigma_\mu^2}} e^{-\frac{1}{2}\left[-\frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)^2}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}+\left(\frac{r^2}{\sigma^2}+\frac{\bar{\mu}^2}{\sigma_\mu^2}\right)\right]} \frac{1}{\sqrt{2\pi\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)^{-1}}}\int e^{-\frac{1}{2}\frac{\left[\mu-\frac{\left(\frac{r}{\sigma^2}+\frac{\bar{\mu}}{\sigma_\mu^2}\right)}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\right]^2}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)^{-1}}} d\mu$$

$$\frac{1}{\sqrt{2\pi\sigma^2\sigma_\mu^2\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}} e^{-\frac{1}{2}\left[-\frac{\frac{r^2}{\sigma^4}+\frac{\bar{\mu}^2}{\sigma_\mu^4}+2\frac{r}{\sigma^2}\frac{\bar{\mu}}{\sigma_\mu^2}}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}+\frac{r^2}{\sigma^2}+\frac{\bar{\mu}^2}{\sigma_\mu^2}\right]}$$

$$\frac{1}{\sqrt{2\pi\left(\sigma^2+\sigma_\mu^2\right)}} e^{-\frac{1}{2}\left[\left(1-\frac{\frac{1}{\sigma^2}}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\right)\frac{r^2}{\sigma^2}-2\frac{\frac{r}{\sigma^2}\frac{\bar{\mu}}{\sigma_\mu^2}}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}+\left(1-\frac{\frac{1}{\sigma_\mu^2}}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\right)\frac{\bar{\mu}^2}{\sigma_\mu^2}\right]}$$

$$\frac{1}{\sqrt{2\pi\left(\sigma^2+\sigma_\mu^2\right)}} e^{-\frac{1}{2}\left[\left(\frac{\frac{1}{\sigma_\mu^2}}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\right)\frac{r^2}{\sigma^2}-2\frac{\frac{r}{\sigma^2}\frac{\bar{\mu}}{\sigma_\mu^2}}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}+\left(\frac{\frac{1}{\sigma^2}}{\left(\frac{1}{\sigma^2}+\frac{1}{\sigma_\mu^2}\right)}\right)\frac{\bar{\mu}^2}{\sigma_\mu^2}\right]}$$

$$\frac{1}{\sqrt{2\pi\left(\sigma^2+\sigma_\mu^2\right)}} e^{-\frac{1}{2}\frac{(r-\bar{\mu})^2}{\left(\sigma^2+\sigma_\mu^2\right)}}$$

8 Bibliographical note

9 Comments on portfolio theory

Portfolio theory looks like a lot of fun, and of great practical importance. Fit your model to return dynamics, compute the optimal portfolio, start your hedge fund. However, there are a number of conceptual and practical limitations, in addition to the advice of Bayesian portfolio theory, that dampen one's enthusiasm.

The average investor holds the market

We tend quickly to forget that *the average investor must hold the market portfolio.* If everyone should follow our portfolio advice, then the return dynamics on which it is built do not represent an equilibrium. For every investor who should buy value or momentum stocks or invest more when the dividend yield is high, there must be someone else out there whose constellation of outside income, risk aversion, or state-variable sensitivity means he should *sell* value or momentum stocks or invest *less* when the dividend yield is high. Portfolio advice *cannot* apply to everyone. Who are the investors on the other side, generating the return anomaly? This is a great question to ask of any portfolio exercise.

For that reason, I have emphasized expressions of optimal portfolios in terms of each investor's characteristics *relative* to those of the average investor. When deciding on an allocation for stocks relative to bonds, the investor can ask "am I more or less risk averse than the average investor?" and not just compare his risk aversion to the market premium. The two calculations are equivalent, of course, but in practice investors may have a better sense of how risk averse they are relative to the average than the absolute calculation. And, on reflection, many investors may realize "that looked good, but I have no reason to really think I'm different than average," which brings us right back to the market portfolio.

Almost all portfolio theory is devoted to telling the (locally) mean-variance investor (i.e. power or recursive utility, no outside income) how to profit from "anomalies", patterns in expected returns and covariances that drive optimal portfolios away from just holding the market. The entire active management and hedge fund industry sells "alpha" to such investors, and they are buying it. So where are all the investors on the other side?

Really, there are only three logical possibilities. First, the investors might be there, in which case a new industry needs to start marketing the opposite strategies: identifying investors with state-variable our outside-income exposure that means they should take portfolios that do even worse than the market on mean-variance grounds, but insure these investors against their state-variable or outside-income risks. Second, the investors might not be there, in which case the anomaly and resulting portfolio advice will die out quickly. Third, we might have deeply mis-modeled the average investors' utility function and state-variable risks, so the advice being spun out by portfolio calculations and implemented by hedge funds is inappropriate to the average investor to which it is being sold. In which case, we should just go back to the market portfolio and understand why the average investor should hold it.

Catch 22

This conundrum illustrates a deeper Catch-22: portfolio theory only really works if you can't use it. If more than measure zero agents should change their portfolios based on your advice, then prices and returns will change when they do so, changing the optimal portfolio.

Of course, it is possible that the "anomalies" do not represent "equilibrium compensation for risk." They might represent "mispricing." If so though, that mispricing will quickly be eliminated once investors learn about it. The resulting optimal portfolio exercise will only be useful for the fleeting moment before the mispricing is wiped away.

The anomalies might instead represent "mispricing" that is not easy to arbitrage because market frictions or institutional constraints make it difficult to trade on the anomaly. But then by definition one can't trade on the anomalies, so an optimal portfolio calculation to take advantage of them cannot work.

The anomalies may represent mispricings that need institutional changes to overcome trading costs. This is a plausible story for the small - firm premium. That premium was stronger before 1979 when the small firm funds were founded. It was very hard before 1979 to hold a good diversified portfolio of small - cap stocks. But finding such an anomaly requires invention of a new institution (the small-cap fund) not simple portfolio calculations.

Of course, most investors think they are simply smarter than the average, like the children of Lake Woebegone who are all above average. At least half of them are wrong.

Consumption

How did we escape the equity premium? We didn't. Every portfolio calculation includes a prescription for consumption, and such prescriptions almost always involve orders of magnitude more consumption volatility, and much higher correlation of consumption growth with portfolio returns, than we observe in practice for any individual investor.

For example our portfolio advice for the power utility investor in (58) sounds perfectly sensible: $\alpha = (\mu - r)/\gamma\sigma^2$. For standard numbers $\mu - r = 8\%$ and $\sigma = 16\%$, $0.08/0.16^2 = 3.125$ so this value of risk aversion invests 100% in equities, while the standard 60/40 advice corresponds to $\gamma = 0.08/(0.6 \times 0.16^2) = 5.21$. If, as I argued in the equity premium chapter, that mean is overstated so perhaps $\mu - r = 4\%$, then the fully-invested investor has risk aversion $\gamma = 0.04/0.16^2 = 1.56$ and 60/40 corresponds to $\gamma = 0.04/(0.6 \times 0.16^2) = 2.6$.

That portfolio advice came together with *consumption* advice, however, in (57). The consumption advice is to set consumption proportional to wealth, for *any* value of risk aversion. Thus the volatility of log consumption growth is equal the volatility of wealth, and from $dW/W = rdt + \alpha(dR - rdt)$, the volatility of wealth is equal to α times the volatility of stock returns. If the investor accepts portfolio advice to be fully invested in equities, he should also accept the advice that consumption should vary by 16 percentage points per year! (This is the flow of consumption services, not durables purchases.) If he accepts 60/40 advice, consumption should vary by 10 percentage points per year.

There is no logical reason to accept the portfolio advice but ignore the consumption advice. Another good question to ask of any portfolio calculation is, "what does the optimal consumption stream look like?" If we don't accept the consumption advice, that means either the environment is wrongly specified (perhaps there is a lot of mean-reversion in stock returns, so that it makes sense for consumption to ignore "transitory variation in wealth") or perhaps the utility function is wrongly specified (habits, for example). But if that is true, the portfolio advice changes as well, and just as drastically.

Payoffs vs. portfolios

I have deliberately contrasted the “payoff” approach with the more traditional “portfolio” approach. Obviously, the former has an attractive simplicity, elegance, and connection to the rest of modern asset pricing.

The most obvious practical limitation of the “payoff” approach is that it leaves you with optimal payoffs, but no strategy for implementing them. I think this limitation is a virtue. Much of the success of the original mean-variance analysis consisted of stopping and declaring victory just before the hard part began. Markowitz showed us that optimal portfolios were mean-variance efficient, but we still really don’t know how to calculate a mean-variance efficient portfolio. Hedge funds are basically selling different ideas for doing that, at very high prices. This is not a defect: understanding the *economic characterization* of the optimal portfolio without having to solve the engineering problem of its construction is a great success. The payoff strategy does the same thing: it stops and characterizes the optimal payoffs without solving the engineering problem of supporting those payoffs by a trading strategy in a particular set of markets.

This observation also suggests a natural organization of the money management industry. Investors are, in the end, interested only in the payoffs. Perhaps the money-management industry should provide the payoffs and not involve the investor too deeply in the portfolio that hedges them.

In some sense, this is what happens already. Why do stocks pay dividends, and why do bonds pay coupons? Consumers could in principle synthesize such securities from dynamic trading of non-dividend-paying stocks and zero-coupon bonds. But date- and state-contingent streams are the securities consumers want in the end, it is less surprising that these are the basic marketed securities. Similarly, swaps which provide desired date- and state-contingent flows are more popular than duration hedges.

It’s also a convenience to solve problems in a way that does not require specifying an implementation. Often, there are lots of different ways to implement an optimal payoff, for example purchases of explicit options vs. dynamic trading. If one chooses portfolio weights, one has to start all over again for every implementation, solving both the economic and engineering questions together. The portfolio approach solves for a lot of things we really aren’t always interested in (dynamic trading strategies) on the way to giving the economic characterization of the payoffs.

Of course, this is all a matter of point of view rather than substance. The two solutions are equivalent. One can always find an implementation for an optimal payoff. I know of no problem that can be fully solved by one method but not by the other.

The other main limitation of the payoff approach is in dealing with incomplete markets. My section on a mean-variance approximation is meant to advance this cause somewhat, by pointing out a nice characterization that is valid in incomplete markets, just as conventional mean-variance analysis is an important benchmark for one-period portfolio problems. No matter how dynamic, intertemporal or incomplete the market, the investor splits his payoffs between the indexed perpetuity and the market, i.e. the claim to the aggregate consumption stream.

Where did all the dynamic trading go? It’s in the market portfolio. Again, portfolio theory, added up across people, is a theory of the composition of the market portfolio. If $\mu - r$ rises, each investor wants to invest more in risky assets, and collectively they can do so by the implicit assumption of linear technologies. The market portfolio becomes more invested in risky assets. By saying "buy the market portfolio" we are simply saying "do what everyone else does." For example, when $\mu - r$ rises in these models, there will be a wave of share issues; your "market index" will buy these new shares, so you will end up doing the same "dynamic trading" as everyone else. You just don't have to solve any portfolio optimization to do it.

Outside income, tailored portfolios, and non-priced factors.

I have emphasized outside income in this treatment, even though it is rarely discussed in the modern portfolio theory literature. I think it's the most important and most overlooked component of portfolio theory, and that paying attention to it could change academic theory and the practice of the money management industry in important ways.

Almost all investors have substantial outside, i.e. fixed or nontradeable, labor or business income. The first thing investors should do is hedge this income stream. If you can't short your own company stock, at least don't buy more than you have to, or short an industry or correlated portfolio. This is just simple insurance. Actually, the first thing investors should do is buy home insurance, even though it's a terrible mean-variance investment. Hedging labor income risk is the same idea.

This fact reopens the door to a modern version of "tailored portfolios." The famous two-fund theorem of mean-variance analysis dealt a serious blow to the once-common idea that investors needed professional advice to pick a stock portfolio appropriate for their individual characteristics – and investors paid large fees for that advice. In the two-fund world, one only needse market portfolio and the risk free rate. Low-cost stock index funds and money market mutual funds were born, and we are all better off for them.

Once we reintroduce labor income or preference shocks, the famous two-fund theorem is not true. There will be many additional "funds," corresponding to typical outside income risks. Figuring out what the funds should be, and matching investors to those funds is not a trivial task. Thus, we have a need for academic research to identify portfolios that match typical outside income risks, and an industry to help investors choose the right ones, and should be able to charge a fee for that service. In this sense, portfolio theory with outside income resurrects tailored portfolios.

In fact, the natural industrial organization of the money management industry might split the two functions. One set of advisers hedges outside income risk. A second set promises mean-variance efficient investment or "alpha." As the formulas separate these two functions, so can the industry structure.

In constructing outside-income hedge portfolios, *nonpriced* factors are just as, or more interesting than the *priced* (or pricing) factors on which most current research focuses. For example, a set of low-cost easily shorted industry portfolios might be very useful for hedging outside income risk, even though they may conform perfectly to the CAPM and provide no "alpha" whatsoever. Buying insurance for no premium is just what you want to do. It's interesting that academic research has focused so exclusively on finding priced factors, which are only interesting to the one remaining jobless mean-variance investor.

One impediment to doing all this is that we don't observe the present value of labor income or outside businesses, so the usual approach that focuses on returns and return covariances is hard to apply. We do observe labor income flows, however, so the payoff-focused approach highlighted in the dynamic mean-variance approximation may prove a useful way to examine outside income hedge portfolios.

10 Problems

1. You can invest in a stock which currently has price \$100. It will either go up to \$130 or down to \$90, with probability 1/2 of each event. (Call the two states u and d.) You can also invest in a bond, which pays zero interest–A \$100 investment gives \$100 for sure.

   (a) Find a discount factor m_{t+1} that prices stock and bond.

   (b) A one-period investor with log utility $u(W_{t+1}) = \ln(W_{t+1})$ has initial wealth \$100. Find this investor's optimal allocation to the stock and bond.

   Hint: first find optimal wealth in the two states tomorrow. Then figure out how to obtain this optimal wealth by investing in the h shares of stock and k bonds.

2. Take the payoff in the one period Black-Scholes example,

$$\hat{x} = \frac{W}{E(m^{1-\frac{1}{\gamma}})} m^{-\frac{1}{\gamma}} = W e^{(1-\alpha)\left(r+\frac{1}{2}\alpha\sigma^2\right)} R_T^{\alpha}$$

   Suppose instead of supporting this payoff by dynamic trading, you choose to support it by a portfolio of put and call options at time zero.

   (a) Find the number of put and call options to buy/sell as a function of strike price. Hint: Graph the payoff of buying a call with strike k, selling a call with strike $k+\Delta$, selling a put with strike $k+\Delta$, and buying a put with strike $k+2\Delta$. Take the limit of this payoff as $\Delta \to 0$ in such a way that the integral of the payoff is one. First express your payoff in terms of how many of these portfolios you buy at each k, and then in terms of how many of the underlying put and call options you want to buy.

   (b) Similarly, implement the payoff for habit utility $(c-h)^{1-\gamma}$ using put and call options.

3. CARA utility and normal. If utility is

$$U(W_T) = e^{-\alpha W_T}$$

   and x_T and hence $W_T = w' x_T$ are normally distributed, then

$$E\left[U(W_T)\right] = e^{-\alpha E W_T + \frac{\alpha^2}{2}\sigma^2(W_T)} = e^{-\alpha w' E(x_T) + \frac{\alpha^2}{2} w' V w}$$

   the first order condition gives

$$w = \frac{1}{\alpha} V^{-1} E(x_T)$$

   Again, the optimal portfolio is mean-variance efficient.

4. Suppose a quadratic utility investor with a one-year horizon (no intermediate consumption) wakes up in the lognormal iid world.

   (a) Mirroring what we did with power utility, find the return on his optimal portfolio in terms of a discount factor m and its moments. Using the definition $\left(\frac{c^b}{R^f W} - 1\right) = \frac{1}{\gamma}$, you can express this return $\hat{R} = \hat{x}/W$ in terms of the local risk aversion at initial wealth, without c^b or W We are looking here for the analogue to

$$\hat{R} = \frac{\hat{x}}{W} = \frac{m^{-\frac{1}{\gamma}}}{E(m^{1-\frac{1}{\gamma}})}$$

(b) Adding the lognormal iid assumption, i.e. $dS/S = \mu dt + \sigma dz$, $dB/B = rdt$, find the return on the investor's portfolio as a function of the stock return. We are looking here for the analogue to

$$\hat{R} = e^{(1-\alpha)\left(r+\frac{1}{2}\alpha\sigma^2\right)} R_T^{\alpha}$$

where $R_T = S_T/S_0$ denotes the stock return, and

$$\alpha \equiv \frac{1}{\gamma}\frac{\mu - r}{\sigma^2}.$$

(c) For $\gamma = 1, 3.125, 5, 20$ and $\mu = 0.09, \sigma = 0.16, r = 1$, make a plot to compare the function $\hat{R} = ...R_T$ in the log and power utility cases. How good is the quadratic as an approximation to power in the range $R_T = 1 \pm 2\sigma$, where the stock is most likely to end up? How good is the quadratic as an approximation to power in the full range, and in particular for describing demands for out of the money options?

(d) Now, solve the portfolio weight problem for the quadratic utility investor by dynamic programming, mirroring what we did with power utility.

(e) To solve the Bellman equation, you can either guess a quadratic form $V(W,t) = -\frac{1}{2}e^{2\eta(T-t)}\left(e^{-r(T-t)}c^b - W\right)^2$ and solve for η. There is a better way however. Since you have solved for the optimal payoffs above, you know the distribution of $W_T = \hat{x}_T$, so you can find the value function directly. Thus, first find the value function from the above solution, then use this as a guess, i.e. verify that it solves the Bellman equation. (This is a great idea, now that I think of it, and offers a constructive way to find value functions for difficult portfolio problems.)

(f) Simulate the stock process using a daily interval. For each value of risk aversion, plot the resulting wealth process for the power and quadratic utility investor, starting at $W_0 = 1$. For each value of risk aversion, also plot the optimal weight in the stock over time for the quadratic and power utility investor. (The power utility investor puts a constant weight in the risky asset, so that one is a horizontal line, but the quadratic utility investor's allocation to the risky asset varies over time.) Again, comment on the dimensions for which the quadratic and power solutions seem similar for a given initial risk aversion, and the dimensions for which the solutions seem quite different.

11 References

Barberis, Nicholas, 2000, "Investing for the Long Run when Returns are Predictable," *Journal of Finance* 55, 225-264.

Bawa, Vijay S., Stephen J. Brown, and Roger W. Klein, 1979, *Estimation Risk and OptimalPortfolio Choice* North-Holland, New York

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* 54, 1609-1646.

Brandt, Michael, and Pedro Santa-Clara 2006, "Dynamic Portfolio Selection by Augmenting the Asset Space" *Journal of Finance* 61, 2187-2218.

Brennan, Michael J., 1981"Optimal Portfolio Insurance," *Journal of financial and quantitative analysis*

Brennan, Michael J., 1998 "Learning about Predictability: the Effect of Parameter Uncertainty on Dynamic Asset Allocation," *European Economic Review* 1, 295-306.

Brennan, Michael J., Eduardo S. Schwartz and Roland Lagnado, 1997, "Strategic Asset Allocation," *Journal of Economic Dynamics and Control* 21, 1377-1403.

Brennan, Michael J., 1975, "The Optimal Number of Securities in a Risky Asset Portfolio When There are Fixed Costs of Transacting: Theory and Some Empirical Results," *Journal of Financial and Quantitative Analysis* 10, 483-96.

Brennan, Michael J. and Yihong Xia, 2000, "Stochastic Interest Rates and Bond-Stock Mix," European Finance Review 4, 197-210.

Brennan, Michael J. and Yihong Xia 2002, "Dynamic Asset Allocation under Inflation," *Journal of Finance* 57, 1201-1238.

Brennan, Michael J. and Yihong Xia 2002a, "Optimal Investment Policies: Fixed Mix and Dynamic Stochastic Control Models", working paper, UCLA.

Campbell, John Y. and Luis M. Viceira, 2001, "Who Should Buy Long-Term Bonds?" *American Economic Review*, 91, 99-127.

Campbell, John Y. and Luis M. Vicera, 1999, "Consumption and Portfolio Decisions when Expected returns are Time Varying," *Quarterly Journal of Economics* 114, 433-95.

Campbell, John Y. and Luis M. Viceira, 2000, "Consumption and Portfolio Decisions when Expected Returns are Time Varying: Erratum" *Quarterly Journal of Economics*,

Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors.* New York: Oxford University Press.

Cambpell, John Y, Yeung Lewis Chan and Luis Viceira, 2003, "A Multivariate Model of Strategic Asset Allocation," *Journal of Financial Economics*, 67, 41-80.

Campbell, John Y. George Chacko, Jorge Rodriguez, and Luis M. Viceira, 2004, "Strategic Asset Allocation in a Continuous-Time VAR Model," *Journal of Economic Dynamics and Control* 28, 2195-2214.

Canner, Niko, N. Gregory Mankiw and David N. Weil 1997, "An Asset Allocation Puzzle," *American Economic Review*, v87, n1, 181-191.

Cochrane, John H., 1999, "Portfolio Advice for a Multifactor World," *Economic Perspectives* Federal Reserve Bank of Chicago 23, 59-78.

Cochrane, John H., 2004, *Asset Pricing, Revised Edition,* Princeton: Princeton University Press

Cochrane, John H., 2005, "A Mean-Variance Benchmark for Intertemporal Portfolio Theory," Manuscript, University of Chicago

Cochrane, John H. Francis Longstaff and Pedro Santa-Clara, 2006 "Two Trees: Asset Price Dynamics Induced by Market Clearing," Forthcoming, *Review of Financial Studies.*

Cox, John C. and Chi-fu Huang, 1989, "Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process," *Journal of Economic Theory,*49, 33-83.

Cuoco, Domenico, 1997. "Optimal Consumption and Equilibrium Prices with Portfolio Constraints and Stochastic Income," *Journal of Economic Theory* v72, n1 (January), 33-73.

Cuoco, Domenico and Hua He, 1994, "Dynamic Equilibrium in Infinite-Dimensional Economies with Incomplete Financial Markets," Manuscript, University of Pennsylvania and University of California, Berkeley

Davis, Steven, and Paul Willen, 2002, "Risky Labor Income and Portfolio Choice," in Zvi Bodie, Olivia S. Mitchell, Brett Hammond and Stephen Zeldes, ed., *Innovations for Financing Retirement*, University of Pennsylvania Press.

Davis, Steven, and Paul Willen, 2000, "Occupation-Level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice," NBER Working Paper 7905.

Davis, Steven, and Paul Willen, 2000, "Using Financial Assets to Hedge Labor Income Risk: Estimating the Benefits," Manuscript, University of Chicago.

Davis, Steven J., Felix Kubler, and Paul Willen, 2005, "Borrowing Costs and the Demand for Equity over the Life Cycle," forthcoming, *Review of Eocnomics and Statistics.*

Davis, Steven J. and Paul Willen, 2002, "Income Shocks, Asset Returns, and Portfolio Choice," *Innovations in Retirement Financing*, pp. 20-49, Pension Research Council Publications. Philadelphia: University of Pennsylvania Press.

Dybvig, Philip H. and Jonathan E. Ingersoll Jr., 1982, "Mean-Variance Theory in Complete Markets," *Journal of Business* 55, 233-51.

Duffie, Darrell, 2001 *Dynamic Asset Pricing Theory*, 3d Edition, Princeton: Princeton University Press.

Fama, Eugene F. 1996, "Multifactor Portfolio Efficiency and Multifactor Asset Pricing," *Journal of Financial and Quantitative Analysis*, 31, 441-465.

Feldman, David, 1992, Logarithmic preferences, myopic decisions, and incomplete information, *Journal of Financial and Quantitative Analysis* 27, 619–629.

Hansen, Lars Peter, 1987, "Calculating Asset Prices in Three Example Economies," in T.F. Bewley, *Advances in Econometrics, Fifth World Congress,* Cambridge University Press.

Hansen, Lars Peter, John C. Heaton and Nan Li, 2005, "Consumption Strikes Back?" Manuscript, University of Chicago.

Hansen, Lars Peter and Thomas J. Sargent, 2004, *Recursive Models of Dyanmic Linear Economies*, Mansuscript, University of Chicago and New York Univeristy, http://homepages.nyu.edu/~ts43/

He, Hua, and Neil D. Pearson, 1991, "Consumption and Portfolio Policies with Incomplete Markets and Short-Sales Constraints: The Infinite Dimensional Case," *Journal of Economic Theory*, 54, 259-304

Heaton, John, and Deborah Lucas, 2000, "Portfolio Choice in the Presence of Background Risk," *Economic Journal*, 110, 1-26.

Gennotte, Gerard, 1986, Optimal portfolio choice under incomplete information, *Journal of Finance* 41, 733–746.

Ingersoll, Jonathan E. Jr., *Theory of Financial Decision Making*

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

Kandel, Shmuel, and Robert F. Stambaugh, 1996, "On the predictability of stock returns: An asset-allocation perspective," *Journal of Finance* 51, 385–424.

Kim, Tong Suk and Edward Omberg (1996), "Dynamic Nonmyopic Portfolio Behavior," *Review of Financial Studies* 9, 141-61.

Klein, Roger W., and Vijay S. Bawa, 1976, "The Effect of Estimation Risk on Optimal Portfolio Choice" *Journal of Financial Economics* 3, 215-231.

Ledoit, Olivier, and Michael Wolf, 2004, 11A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices," *Journal of Multivariate Analysis*, 88, 365-411.

Ledoit, Olivier, and Michael Wolf, 2003, "Improved Estimation of the Covariance Matrix of Stock Returns With an Application to Portfolio Selection," *Journal of Empirical Finance* 10, 603-621.

Ledoit, Olivier, and Michael Wolf, 2004, "Honey, I Shrunk the Sample Covariance Matrix," *Journal of Portfolio Management,* 31

Lettau, Martin and Jessica Wachter 2005, "Why is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium," Manuscript, NYU and University of Pennsylvania.

Liu, Jun, 2004, "Portfolio Selection In Stochastic Environments," Manuscript, Rady School of Business, University of California at San Diego

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