Applications of long-run mean-variance analysis

John H. Cochrane

February 22, 2008

1 Introduction

In the following sections I pursue several simple applications of the long-run mean / long-run variance portfolio theory. I do this with some trepidation. After all, one big selling point of the theory is that one can declare victory and gain the standard insights from mean-variance analysis without actually doing a lot of computation. Here, I attack that hard part. These notes are also a report of difficulties that need to be surmounted in applications.

The first exercise one tries with standard one-period mean-variance optimization is to allocate investment among a small number of ad-hoc portfolios, such as size, value or industry, and find the mean-variance frontier, using sample moments for the mean and variance of the portfolios. Of course, one tries to find interesting portfolios that one hopes capture most of the investment possibilities of larger collections of assets. Later, one tackles the hard problems of large numbers of assets (there are more securities than data points), time-varying means and correlations (especially of individual securities), the addition of dynamic strategies, the wacky weights that result from small sampling errors in mean returns, and so forth.

Here, I attempt analogous first exercises. I start with a small number of rather arbitrary, but, I hope, well-chosen and interesting, payoffs. In particular, I choose a few *portfolios* like value and industry, rather than individual securities; I choose a few (but, again I hope well-chosen and interesting) *dynamic strategies*, like an investment in the market that rises linearly with a regression forecast of the market expected excess return, and I choose some arbitrary (but, once again, I hope well-chosen and interesting) *payout rules* for converting accumulated wealth into consumption streams.

I follow two approaches to constructing a discount factor. First, I construct a discount factor from one period returns, and iterate that discount factor forwards so that it prices long-term payoffs. Second, I construct a discount factor directly from the stream of payoffs. The second approach is in some ways much more satisfying, as it gets away entirely from modeling one period returns. However, the lessons of the first approach are very useful in setting it up.

2 Iid normal returns

I start with the natural benchmark for portfolio problems: continuous time, iid returns and a constant interest rate. The investor can rebalance continuously. This familiar benchmark case gives some feeling for how long run means and variances work, and it is a natural springboard for the more complex cases in which the answers are not known.

However, this is not a good environment for applying long-run mean-variance ideas. Log-normal long-horizon returns become severely distorted from normal. Mean-variance thinking is more naturally applied in an environment with returns that are closer to normal. Fortunately, the data seem to behave more normal than lognormal, as I document below. Hence, this sec-

tion is more in the nature of a problem-set exercise that reveals some intuition about long-run mean-variance mathematics rather than a springboard for successful real-world application.

I compare the quadratic and power utility solutions in this environment. The quadratic is in no sense an "approximation" for the power utility answer. However, we can see in what sense the quadratic and power solutions resemble each other, which gives us some idea how useful a "benchmark" the quadratic solution might be in this situation.

2.1 Setup and x^*

There are N basis assets whose instantaneous excess returns follow

$$dr_t^e = \mu dt + \sigma dz_t; \quad \sigma\sigma' = \Sigma,$$

and there is a constant riskfree rate $r^f dt$. I use the notation dr to denote instantaneous returns, i.e. if dividends are paid at rate Ddt, then $dr \equiv dp/p + D/p\ dt$; $dr^e = dr - r^f dt$.

The investor creates a payoff stream x_t from initial wealth W_0 by dynamically investing,

$$dW_t = (rW_t - x_t)dt + \omega_t' dr_t^e$$

with the usual "transversality condition" that the time-zero value of wealth must eventually tend to zero $\lim_{T\to\infty} p(W_T) = 0$, and the usual limits on portfolio investments ω to rule out doubling strategies. The payoff space $\underline{X}$ consists of payoffs generated in this way.

This market is complete, so the unique payoff $x^* \in \underline{X}$ that generates prices by $p = k\mathcal{E}(x^*x)$ is characterized by

$$\frac{dx_t^*}{x_t^*} = (\rho - r)\, dt - \mu'\Sigma^{-1}\sigma dz_t \tag{46}$$

and thus

$$x_t^* = e^{\left(\rho - r - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\sigma \int_0^t dz_t} \tag{47}$$

This discount factor also prices any payoff x_t formed by *arbitrary* trading strategies in the underlying assets.

This is the discount factor scaled by $e^{\rho t}$, so we can express prices as $p = k\mathcal{E}(x^*x) = E\int e^{-\rho t}x_t^* x_t dt$. The usual unscaled discount factor is $\Lambda_t^* = e^{-\rho t}x_t^*$, which satisfies

$$\frac{d\Lambda_t}{\Lambda_t} = -rdt - \mu'\Sigma^{-1}\sigma dz_t$$

and for which $p\Lambda_0^* = E\int_0^\infty \Lambda_t^* x_t dt$. Obviously, it is only a matter of notational convenience whether we use Λ^* or x^*.

The price of x^* itself is

$$\begin{aligned} p(x^*) &= k\mathcal{E}(x^{*2}) = E\int_0^\infty e^{-\rho t} e^{2\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - 2\mu'\Sigma^{-1}\sigma\int_0^t dz_t} dt \\ p(x^*) &= \frac{1}{2r^f - \rho - \mu'\Sigma^{-1}\mu} \end{aligned} \tag{48}$$

I consider the possibility that $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$ and thus $p(x^*) = \infty$ below. From (46) (or directly from (47)) we have $E(x_t^*) = e^{(\rho - r)t}$ and hence

$$\mathcal{E}(x^*) = \rho \int_0^\infty e^{-\rho t} e^{\left(\rho - r^f\right)t} = \frac{\rho}{r^f}. \tag{49}$$

With a constant risk free rate, we have $y^f = r^f$ of course.

2.2 Consumption-portfolio problems

Since this market is complete, we can write the portfolio problem directly as

$$\max \mathcal{E}\left[u(c_t)\right] \quad s.t.\ p(c_t) = k\mathcal{E}\left(x_t^* c_t\right) = W_0.$$

The first order conditions give the solution to the portfolio problem,

$$u'(c_t) = \lambda x_t^* \quad \Rightarrow \quad c_t = u'^{-1}\left(\lambda x_t^*\right)$$

For power utility

$$u(c) = \frac{1}{1-\gamma} c_t^{1-\gamma},$$

we have

$$c_t^p = \left(\lambda x_t^*\right)^{-\frac{1}{\gamma}}. \tag{50}$$

where the superscript p in c^p reminds us that this is consumption for the power utility investor. For quadratic utility

$$u(c) = -\frac{1}{2}\left(c_t^b - c_t\right)^2,$$

we have

$$c_t^q = c_t^b - \lambda x_t^* \tag{51}$$

As before, the investor purchases the bliss point, financed by shorting the "most expensive" payoff x^*.

Equations (50) and (51) give us the essence of the difference between power and quadratic utility, i.e. long-run mean-variance efficient portfolios. The power utility payoff is a *nonlinear* declining function of x^*. The quadratic utility payoff, which corresponds to the long-run mean-variance frontier, is a *linear* declining function of x^*, and thus the payoff is itself mean-variance efficient.

To standardize the portfolios to initial wealth, we find Lagrange multipliers from the budget constraints. For power utility the result is[2]

$$\frac{c_t^p}{W_0} = \frac{1}{\gamma}\left[\rho + (\gamma - 1)\left(r^f + \frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right)\right] {x_t^*}^{-\frac{1}{\gamma}}. \tag{52}$$

[2]

$$\begin{aligned}
W_0 &= p(c_t) = \mathcal{E}\left[x_t^* \left(\lambda x_t^*\right)^{-\frac{1}{\gamma}}\right] \\
W_0 &= \lambda^{-\frac{1}{\gamma}} E \int_0^\infty e^{-\rho t} e^{\left(\frac{\gamma-1}{\gamma}\right)\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \left(\frac{\gamma-1}{\gamma}\right)\mu'\Sigma^{-1}\int_0^t dz_t} \\
&= \lambda^{-\frac{1}{\gamma}} \int_0^\infty e^{-\rho t} e^{\left(\frac{\gamma-1}{\gamma}\right)\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t + \frac{1}{2}\left(\frac{\gamma-1}{\gamma}\right)^2 \mu'\Sigma^{-1}\mu}
\end{aligned}$$

We have already worked out the quadratic utility case. From Proposition 2 equation (24), the nicest expression is

$$\frac{c_t^q}{W_0} = r^f + \frac{1}{\gamma}\left(r^f - y_t^*\right) \tag{53}$$

where γ represents a local risk aversion coefficient, controlled by the bliss point c^b.

Figure 1 contrasts the power utility case (52) with the quadratic or mean-variance case (53). Both consumption decision rules slope downward. Lower returns mean higher marginal utility x^* or y^* and hence lower consumption. The power utility investor is particularly anxious not to suffer consumption declines; he therefore consumes more in bad times (right hand side of the graph) funding that consumption in good times (middle of the graph) He is also more price-sensitive; if consumption is really cheap on the left side of the graph, he can increase consumption without limit. The quadratic utility investor does not exceed the bliss point, even if the price of consumption (x^*) falls to zero.

Clearly, the difference between power and quadratic is "small" for "small" returns and thus movements in x^* , and vice versa. Figure 1 includes the 1,5, and 10 year densities of x^* to gives some sense of how large the plausible range is for typical parameters. The major difference in the likely states is on the left side of the graph, times of good returns, where the power utility consumption increases without bound.

Figure 2 plots a simulation of power and quadratic consumption from this model. In this particular draw, good and bad luck mostly balance, so x^* mostly stays in the range that quadratic and power utility give quite similar answers. Again, the major difference comes if returns are very *good*, meaning the discount factor x^* *declines* a great deal. (Left hand side of Figure 1.) Here, the quadratic utility investor stops increasing consumption as it approaches the bliss point, $c^b = 0.75$ in this case, while the power utility investor keeps going. Very bad luck leads to increases in x^*. Here, the power utility consumption will stay positive, while quadratic utility consumption continues to fall and can even fall below zero. We see this kind of event near year 10 of the simulation (I picked a simulation that showed both tails).

The general conclusion I draw from the analysis is that quadratic utility can be a useful "benchmark" in that the portfolios are qualitatively similar. However, it is certainly not an "approximation", in that the answers are different in quantitatively important ways.

$$\begin{aligned} &= \lambda^{-\frac{1}{\gamma}} \int_0^\infty e^{-\frac{1}{\gamma}\left[\rho+(\gamma-1)\left(r^f+\frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right)\right]t} \\ &= \frac{\lambda^{-\frac{1}{\gamma}}}{\frac{1}{\gamma}\left[\rho+(\gamma-1)\left(r^f+\frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right)\right]} \end{aligned}$$

Substitute this expression for λ in (50).

If $\gamma < 1$ and enough so that

$$\rho + (\gamma - 1)\left(r^f + \frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right) < 0$$

it appears that $p(c)$ is infinite. In this case, the answer is $c = 0$. The consumer endlessly puts off consumption since investment is so much more attractive.

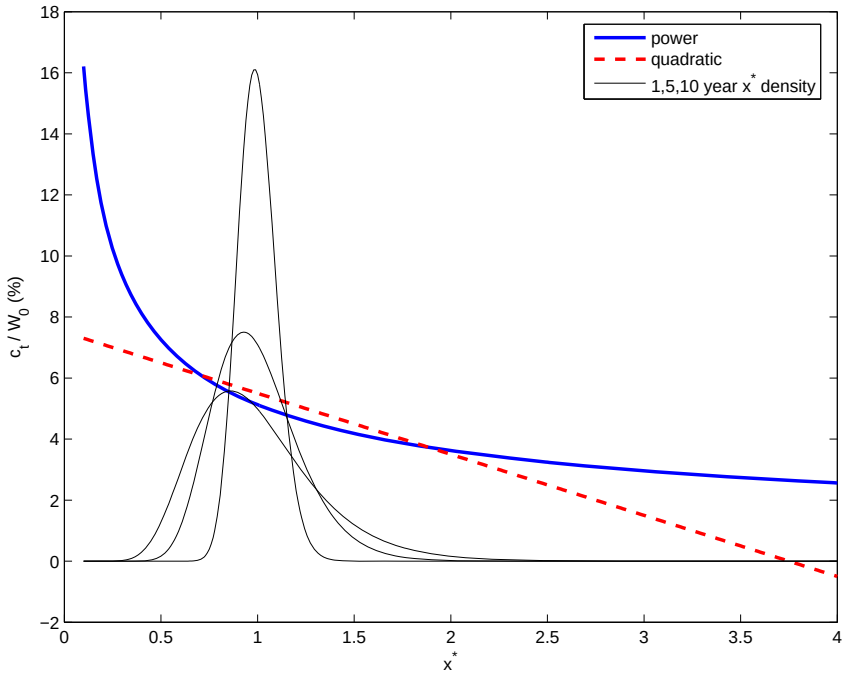

Figure 1: Consumption payoffs from power utility and quadratic utility in the iid return case. The x axis is the pricing payoff or marginal utility process x^*. The y axis gives consumption relative to initial wealth as a function of x^*. Parameters are risk aversion $\gamma = 2$, discount factor $\rho = 0.01$, riskfree rate $r^f = 0.05$ excess return mean $\mu = 0.04$ and standard deviation $\sigma = 0.20$.

2.3 Long-run mean-variance frontier

Since z^* defines the mean-variance frontier, the long-run Sharpe ratio or slope of the long-run mean-variance frontier is given by $\mathcal{E}(z^*)/\tilde{\sigma}(z^*)$. Evaluating in the iid lognormal model with a riskfree rate, we have

$$\frac{\mathcal{E}(z_t^*)}{\tilde{\sigma}(z_t^*)} = \sqrt{\frac{(r^f - \rho)^2 + \rho\mu'\Sigma^{-1}\mu}{r^{f2} - (r^f - \rho)^2 - \rho\mu'\Sigma^{-1}\mu}}. \tag{54}$$

Derivation We evaluate

$$\begin{aligned}
\mathcal{E}(z_t^*) &= \mathcal{E}\left(1 - \frac{y_t^*}{r^f}\right) = 1 - \frac{\mathcal{E}(x^*)}{r^f p(x^*)} = 1 - \frac{\rho}{r^{f2} p(x^*)} \\
&= 1 - \frac{\rho}{r^{f2}}\left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right) \\
&= \frac{(r^f - \rho)^2 + \rho\mu'\Sigma^{-1}\mu}{r^{f2}}.
\end{aligned}$$

The first line follows from (49). We can find the variance from

$$\tilde{\sigma}^2(z^*) = \mathcal{E}(z^{*2}) - [\mathcal{E}(z^*)]^2 = \mathcal{E}(z^*) - [\mathcal{E}(z^*)]^2 = \mathcal{E}(z^*)\left[1 - \mathcal{E}(z^*)\right]$$

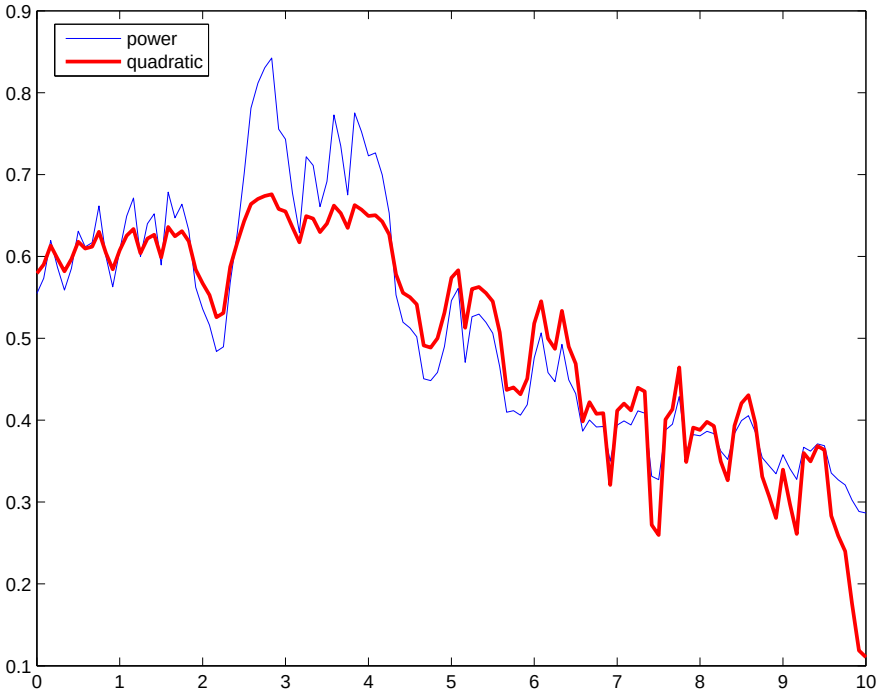

Figure 2: Simulated consmption from power and quadratic utility in iid return case.

The second equality follows from the defining property $\mathcal{E}(z^* z) = \mathcal{E}(z)$. Then we have

$$\frac{\mathcal{E}(z^*)}{\tilde{\sigma}(z^*)} = \frac{\mathcal{E}(z^*)}{\sqrt{\mathcal{E}(z^*)\left[1-\mathcal{E}(z^*)\right]}} = \sqrt{\frac{\mathcal{E}(z^*)}{1-\mathcal{E}(z^*)}} \tag{55}$$

and (54) follows.

The formula (54) is complex because the investor cares about volatility over time as well as across states of nature. To understand it, suppose first that there is no instantaneous Sharpe ratio, $\mu'\Sigma^{-1}\mu = 0$, or equivalently that the risky asset is simply absent. Then there still is a long-run mean-variance frontier,

$$\frac{\mathcal{E}(z_t^*)}{\tilde{\sigma}(z_t^*)} = \sqrt{\frac{(r^f-\rho)^2}{r^{f2}-(r^f-\rho)^2}}. \tag{56}$$

The only investment is the riskfree rate r^f, but there still is the question of how fast to take money out of wealth placed in the riskfree investment, and thus a trade-off between average level and volatility over *time*. If the consumer takes consumption $c_t = \alpha W_t$, then wealth grows at $dW_t/W_t = \left(r^f - \alpha\right) dt$, and consumption follows $c_t = W_0 \alpha e^{(r^f-\alpha)t}$. If the consumer chooses $\alpha = r^f$, he obtains constant consumption and wealth $W_t = W_0$, $c_t = r^f W_0$. This path has zero long-term variance, and long-run mean $\mathcal{E}(c) = r^f W_0$. However, if the consumer chooses a lower consumption/wealth ratio $\alpha < r^f$, he will start with lower consumption initially, but he

will then obtain a rising consumption path. The rising path represents a source of long-term variance, since that measure prizes stability over time as well as over states of nature.

If $\rho = r^f$, of course, the investor chooses a constant consumption path. Here, we view that choice as coming from a constant long-run mean, independent of the choice α of consumption-wealth ratio, so the investor chooses the minimum long-run variance portfolio. If $r^f < \rho$, the consumer chooses a declining consumption path. Here, we see in this case that that lower initial consumption with higher later growth raises the long-run mean as well as the variance. In this way, the choice of *payout rate* gives rise to a long-run mean-variance trade-off, captured by formula (56), even when there is no risk at all. The long-run Sharpe ratio in this case expresses the attractiveness of distorting consumption away from a constant *over time* in order to raise its overall level.

Now, suppose there is risk, but $\rho = r^f$ so there is no interesting substitution over time. The long-run Sharpe ratio simplifies to

$$\frac{\mathcal{E}(z_t^*)}{\tilde{\sigma}(z_t^*)} = \sqrt{\frac{\mu'\Sigma^{-1}\mu}{r^f - \mu'\Sigma^{-1}\mu}} \tag{57}$$

This function is increasing in the instantaneous sharpe ratio $\sqrt{\mu'\Sigma^{-1}\mu}$ as we might expect. Now the investor will get more mean, and variance, by taking a larger investment in the risky assets.

As a result of the denominator, typical numbers for the long-run Sharpe ratio will be a good deal larger than those for the instantaneous Sharpe ratio $\sqrt{\mu'\Sigma^{-1}\mu}$. That fact simply reflects different units. 10 year Sharpe ratios are about $\sqrt{10}$ larger than one-year Sharpe ratios, and the long-run frontier in essence characterizes a weighted sum of long-run returns.

2.4 Portfolio strategies

The whole point of the contingent-claim approach is to forget about trading strategies and to focus on final payoffs. However, we can easily find the trading strategies in this case, so it's interesting to investigate what they are, and how the quadratic and power utility portfolio weights compare.

We generate yields or payoff streams by varying the portfolio weights ω_t in the risky assets and by taking the payoff as a dividend. Wealth then follows

$$dW_t = \left(r_t^f W_t - x_t\right) dt + \omega_t' dr_t^e. \tag{58}$$

x_t and ω_t describe dollar payouts and dollar positions in the excess returns. I use $\alpha = x/W$ and $w = \omega/W$ to describe portfolio weights as fractions of wealth.

The power-utility consumption stream has a familiar portfolio characterization. The investor holds an instantaneously mean-variance efficient portfolio with constant weights and consumes a constant fraction of wealth,

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

The quadratic utility investor also holds a mean-variance efficient portfolio. (Equivalently, the long-run mean-variance frontier results from such a portfolio)

$$w_t = \frac{\left(c^b - r^f W_t\right)}{r^f W_t}\Sigma^{-1}\mu = \frac{1}{\gamma_t}\Sigma^{-1}\mu. \tag{59}$$

The *composition* of the portfolio is instantaneously mean-variance efficient. However, the *size* of the risky asset portfolio changes over time. We can interpret the term before Σ^{-1} as the inverse of a local risk-aversion coefficient. As wealth rises to the value $r^f W_t = c^b$ that would support bliss-point consumption, the quadratic-utility investor becomes more risk averse.

The payout rules for power c^p and quadratic c^q consumption are

$$c_t^p = \frac{1}{\gamma}\left[\rho + (\gamma - 1)\left(r^f + \frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right)\right] W_t.$$

$$c_t^q = c^b - \left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)\left(\frac{c^b}{r^f} - W_t\right) \tag{60}$$

Both consumptions rise proportionally to wealth, but quadratic utility consumption also has an intercept.

In sum, we can view the difference between the power and quadratic utility investor is not in the *composition* of the risky portfolio – both investors hold mean-variance efficient portfolios – but in the *payout rate* relating consumption to wealth and the *portfolio weight* or *allocation* between risky and riskfree assets. The quadratic utility / long-run mean-variance portfolio adds a constant to the rules of the power utility case.

In this traditional environment, the power-utility investor holds a constantly-rebalanced fraction of wealth in each asset. There is another representation of the quadratic utility investor's portfolio that also holds a constantly-rebalanced constant fraction of wealth in a single riskfree asset.

Start with the diffusion representation for x^* (46) and then rewrite it as a value process in the form (58),

$$\begin{aligned}
\frac{dx^*}{x^*} &= \left(\rho - r^f\right)dt - \mu'\Sigma^{-1}\sigma dz_t \qquad (61)\\
\frac{dx^*}{x^*} &= \left[r^f - \left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)\right]dt - \mu'\Sigma^{-1}\left(\mu dt + \sigma dz_t\right)\\
\frac{dx^*}{x^*} &= \left[r^f - \frac{1}{p(x^*)}\right]dt - \mu'\Sigma^{-1}\left(\mu dt + \sigma dz_t\right)\\
dx_t^* &= \left(r^f x_t^* - y_t^*\right)dt - x_t^*\mu'\Sigma^{-1}dr_t^e.
\end{aligned}$$

In this way, x^* is value process that generates the payoff $y^* = x^*/p(x^*)$ as its dividend stream, as well as being a discount factor, it is a payoff, and it is a value process that generates the payoff

From (61), then, the portfolio that generates y^* is a constantly-rebalanced constant-weight *short* position in a mean-variance efficient investment, $w = -\mu'\Sigma^{-1}dr_t^e$, that pays out a constant fraction $\left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)$ of its value. As we saw, the quadratic utility investor consumes

$$c^q = \hat{x} = c^b - \left[c^b/y^f - W\right]y^*$$

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

In sum, the quadratic utility investor *shorts* a portfolio that is *short* a constant fraction of wealth in a mean-variance efficient portfolio.

More generally, to create a long-run mean-variance efficient yield

$$y^{mv} = r^f + w \times (r^f - y^*)$$

you do *not* simply hold a constant fraction of wealth in a mean-variance efficient instantaneous return $r^f + a\mu'\Sigma^{-1}dr^e$ for some positive a and pay out at some fixed rate. Instead, you hold a *short* position in a security (y^*) that itself cumulates a *short* position $-\mu'\Sigma^{-1}dr^e$ in a mean-variance efficient portfolio. While the double shorting action cancels itself in instantaneous returns, it does not cancel itself in cumulated returns, since the cumulation process is nonlinear.

Uniting (58), (59), and (60), we can express the evolution of wealth as

$$\left(\frac{c^b}{r^f} - W_t\right) = \left(\frac{c^b}{r^f} - W_0\right) e^{\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\mu \int_0^t dB_s}$$

Thus, $\left(\frac{c^b}{r^f} - W_t\right)$ is lognormally distributed. Wealth never exceeds c^b/r^f.

Derivation

$$\begin{aligned}
dW_t &= \left(r^f W_t - \left[c^b - \left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)\left(\frac{c^b}{r^f} - W_t\right)\right]\right)dt + \left(\frac{c^b}{r^f} - W_t\right)\mu'\Sigma^{-1}dr_t^e \\
dW_t &= \left(r^f W_t - c^b + \left(2r^f - \rho\right)\left(\frac{c^b}{r^f} - W_t\right)\right)dt + \left(\frac{c^b}{r^f} - W_t\right)\mu'\Sigma^{-1}\sigma dB \\
dW_t &= \left(r^f - \rho\right)\left(\frac{c^b}{r^f} - W_t\right)dt + \left(\frac{c^b}{r^f} - W_t\right)\mu'\Sigma^{-1}\sigma dB \\
d\left(\frac{c^b}{r^f} - W_t\right) &= \left(\rho - r^f\right)\left(\frac{c^b}{r^f} - W_t\right)dt - \left(\frac{c^b}{r^f} - W_t\right)\mu'\Sigma^{-1}\sigma dB
\end{aligned}$$

The "return on the market portfolio" is price gain plus dividend,

$$\begin{aligned}
dr^m &= \frac{dW_t}{W_t} + \frac{c_t}{W_t}dt \\
dr^m &= r^f dt + \left(\frac{c^b}{r^f W_t} - 1\right)\mu'\Sigma^{-1}dr^e \\
dr^m &= \left[r^f + \left(\frac{c^b}{r^f W_t} - 1\right)\mu'\Sigma^{-1}\mu\right]dt + \left(\frac{c^b}{r^f W_t} - 1\right)\mu'\Sigma^{-1}\sigma dB
\end{aligned}$$

The last expression follows either by substituting in the expressions for dW_t or more directly from the portfolio weights (59)

This last expression makes an important point. It doesn't really make sense to specify a lognormal diffusion for the market portfolio in a mean-variance exercise. When all investors are of this form, and a long-run CAPM emerges, investors will control the characteristics of the market portfolio so that it does not follow a lognormal diffusion. In particular, as wealth grows, the weights in the risky assets decline. This action cuts off the troublesome right tail of the lognormal diffusion. (In an intertemporal model, we think of the asset returns as the returns on linear technologies. Investors determine the composition of the market portfolio. They

collectively reduce investment in the risky technologies when wealth rises.) The Dybvig and Ingersoll (1982) puzzle does not apply, since market returns are never large enough. Together with the fact in Figure 6 that market returns are in fact missing the large troublesome right tail of the lognormal, this calculation suggests better hope for the long-run mean-variance analysis in actual data than a lognormal model suggests.

2.5 Potentially infinite price of x^*

Formula (48),

$$p(x^*) = \frac{1}{2r^f - \rho - \mu'\Sigma^{-1}\mu}$$

reveals a technical limitation of this standard setup. If $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$, the price of the payoff x^* is *infinite.* Then the yield $y^* = x^*/p(x^*)$ is undefined. These parameters are not particularly extreme, at least relative to standard statistics. For example, if one takes a 6% equity premium and an 18% standard deviation of market returns, then $\mu'\Sigma^{-1}\mu = 1/9 = 0.11$. A 6% interest rate r^f and a 1% discount rate ρ are just enough to give an infinite $p(x^*)$, and any lower interest rate or higher discount factor make matters worse. (One might argue that these standard statistics are extreme, which is the entire equity premium puzzle, but that's another issue.) The issue is similar to the "nirvana" solutions that Kim and Ohmberg (1996) find for power utility in a time-varying return environment.

Everything seems to explode at $2r^f - \rho - \mu'\Sigma^{-1}\mu = 0$. We can write the Sharpe ratio from (54) as

$$\frac{\mathcal{E}(z_t^*)}{\tilde{\sigma}(z_t^*)} = \sqrt{\frac{(r^f-\rho)^2 + \rho\mu'\Sigma^{-1}\mu}{r^{f2} - (r^f-\rho)^2 - \rho\mu'\Sigma^{-1}\mu}} = \sqrt{\frac{r^{f2}}{\rho\left[2r^f - \rho - \mu'\Sigma^{-1}\mu\right]} - 1},$$

so the slope of the long-run mean-variance frontier rises to infinity. The minimum second-moment yield $y_t^* = x_t^*/p(x^*)$ is collapses to zero – zero mean and zero standard deviation. A payoff that costs one and delivers zero is a great short opportunity; shorting this payoff and investing in the risk free rate gives higher and higher long-run Sharpe ratios. The consumption decision rule (60) goes simply to consumption at the bliss point, $c_t = c^b$, for any initial wealth.

These pathologies all seem to indicate an arbitrage opportunity at or beyond the limit $2r^f - \rho - \mu'\Sigma^{-1}\mu = 0$. They do not. The key restriction on dynamic trading is that the time-zero value of wealth must be zero, $\lim_{T\to\infty} p(W_T) = 0$. Instead, the strategy that seems to provide bliss point consumption and the strategy that seems to provide $-y^* = 0$ despite an initial investment of -1 both involve rolling over debts forever, in such a way that the present value of the debt remains constant for $2r^f - \rho - \mu'\Sigma^{-1}\mu = 0$ and grows with horizon past that point, violating the no-arbitrage limitation on dynamic trading.

Underlying it all, at the limit the discount factor x_t^* ceases to be a traded payoff – it exits the payoff space $\underline{X}$ which is limited to payoffs generated by trading strategies that do not violate arbitrage. To see this, compute from the x^* process (**??**) $E(x_t^{*2}) = e^{\left(\rho - 2r^f + \mu'\Sigma^{-1}\mu\right)t}$ If this grows faster than $e^{\rho t}$, x_t^* no longer obeys the transversality condition $\lim_{T\to\infty} p(x_T^*) = \lim_{T\to\infty} E(x_T^{*2}) = 0$. At the same time, it leaves the space of square-integrable $\mathcal{E}(x_t^2) < \infty$ payoffs.

How did x^* leave the space of traded payoffs? Didn't the law of one price guarantee a discount factor in the payoff space? The answer is that the payoff space in this infinite-horizon model does

not meet the full set of assumptions required for the existence theorem. It is not "complete:" If you take a sequence of payoffs resulting from positive payout rates, but with a payout rate tending to zero, each of those payoffs is in $\underline{X}$ – the time-zero value of their limiting wealth is zero – but the limit *point* is not in the payoff space. For $2r^f - \rho - \mu'\Sigma^{-1}\mu > 0$ we are "lucky". Even though we have no guarantee that there will be a discount factor in the payoff space, there happens to be one. Beyond, we are not so lucky.

Therefore, at $2r^f - \rho - \mu'\Sigma^{-1}\mu = 0$ and beyond, this *solution* technique is no longer valid. It gives answers to the long-run mean-variance frontier and to the quadratic utility investor's problem that violate arbitrage limits on trading strategies. The discount factor x^* is no longer a traded asset, so y_t^* is undefined (x_t^* has no price to divide by) and we cannot find the long-run mean-variance frontier or the solution to the quadratic utility investor's problem by taking linear functions of y^*.

To solve quadratic problems for $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$, in this setup, then, one must impose arbitrage constraints on trading, such as a borrowing constraint and a portfolio constraint (to rule out doubling strategies), and one must expect them to bind. With constraints imposed, we still will not guarantee a traded discount factor, since now the assumption of arbitrary portfolio formation is violated. If payoff x grazes the borrowing constraint, payoff $2x$ violates that constraint. In practice, it may be easier instead to limit the investor to the choice of a finite number of trading strategies – payout rules and portfolio weights – and all linear combinations of those trading strategies. This space is complete, and bounds the investor away from arbitrage strategies.

There is "almost" an arbitrage opportunity as we approach the limit, and the slope of the long-run mean-variance frontier becomes arbitrarily high. However, the strategies that give these results increasingly rely on wealth that "almost" explodes, rolling over debt that grows astronomically if not quite explosively. While technically correct, such solutions are less and less interesting.

Why not....

Why does power utility not run in to similar problems? For $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$, x_t^* is still outside the payoff space and its price is still infinite. However, the power utility solution is $u'(c_t^p) = \lambda x_t^*$; $c_t^p = (\lambda x_t^*)^{-\frac{1}{\gamma}}$. Even though marginal utility λx_t^* has an infinite price, $(x_t^*)^{-\frac{1}{\gamma}}$ has a finite price, so consumption $(\lambda x_t^*)^{-\frac{1}{\gamma}}$ has a finite price. It is not always necessary that marginal utility have a finite price.

This is a "complete market" so the issue has nothing to do with market incompleteness. The payoff space can fail to be "complete" (Cauchy sequences converge) even when the market is "complete" (all contingent claims are traded.)

One might think that this situation can be resolved by giving the quadratic utility investor a growing bliss point, perhaps even driven by shocks dz, so that he wants growth as does the power utility investor. This modification will also help to make the quadratic utility setup a better approximation to power or other utility results. Alas, this modification does not solve the problem of potentially infinite $p(x^*)$. From (23) we can write the optimal portfolio $c^q = -\lambda x^* + \hat{c}^b - \hat{e}$. The investor purchases claims to the bliss point (and labor income) hedge portfolios, by selling claims to x^*. The nature of this operation does not change if $\hat{c}^b$ grows (deterministically or stochastically) over time; if $p(x^*)$ becomes infinite, the shadow value of

wealth λ seems to go to zero and it seems the investor can finance any $\hat{c}^b$. If c^b grows so much (in mean or variance) that its value too becomes infinite, we might overcome the problem, but alas two infinities do not cancel.

One might think that the situation can be resolved by raising the discount rate in the $\mathcal{E}$ operator to quash long-run growth. Alas, the situation becomes *worse* if one chooses a larger (arbitrary) ρ in the definition of $\mathcal{E}$, disconnecting it from the subjective discount factor. Directly, $2r^f - \rho - \mu'\Sigma^{-1}\mu > 0$ requires higher r^f or lower $\mu'\Sigma^{-1}\mu$ for higher ρ. More deeply, the pricing operator $\Lambda_t = e^{-\rho t}x_t^*$ must be the same for any ρ. Thus as one chooses to represent prices with a larger ρ, the corresponding x_t^* must grow *faster* so that $e^{-\rho t}x_t^*$ is unaffected.

Intuitively, the lognormal return model gives rise to infrequent "disasters" of extremely low return, low consumption, and thus very high marginal utility. When we *price* marginal utility, taking $E(x^{*2})$, these high marginal utility states also have very high prices (x^* plays both roles). Since a quadratic utility investor does not care that much about consumption declines, being willing even to tolerate negative consumption with finite marginal utility, he actively sells consumption in these very high-priced states to finance a great deal of consumption in other states. More fundamentally, the result comes from the probability structure imposed by lognormal returns. A probability model with less frequent or less highly valued disaster states will not generate infinite prices.

Finite horizons

To understand these limiting issues more deeply, we can look at the corresponding finite-horizon problems (from $t = 0$ to T). These models are well-behaved. By construction, terminal wealth is always zero $W_T = 0$, and they do not produce solutions with apparent arbitrage opportunities. All the discount factors and consumption streams remain square-integrable. We can then take the limit as $T \to \infty$ and understand how the *limit point* differs from the limit.

Since the general case is rather complex, I focus on two special cases that get at the central intuition. By examining the case $\mu = 0$, we understand what happens to the allocation of consumption *over time* as the boundary, now $\rho = 2r^f$ is crossed. By examining the case $\rho = r^f$, $\mu \neq 0$, we then understand the effects of risky assets without the mess caused by the terms involving allocation over time. All the derivations are presented below.

Allocation over time, $\mu = 0$. In this case the only decision is how to allocate consumption over time. The optimal consumption profile in this case is

$$\begin{aligned} c_t &= c^b - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W\right] \frac{\left(2r^f - \rho\right)}{1 - e^{(\rho - 2r^f)T}} e^{(\rho - r^f)t}; \quad \text{if } \rho \neq 2r^f \\ c_t &= c^b - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W\right] \frac{1}{T} e^{(\rho - r^f)t}; \text{ if } \rho = 2r^f \end{aligned} \tag{62}$$

To understand this profile over time for fixed T, we can write it as

$$c^b - c_t = \left(c^b - c_0\right) e^{(\rho - r^f)t} \tag{63}$$

Figure 3 illustrates these paths. As we expect, if $\rho = r^f = 0.05$, the investor wants a constant

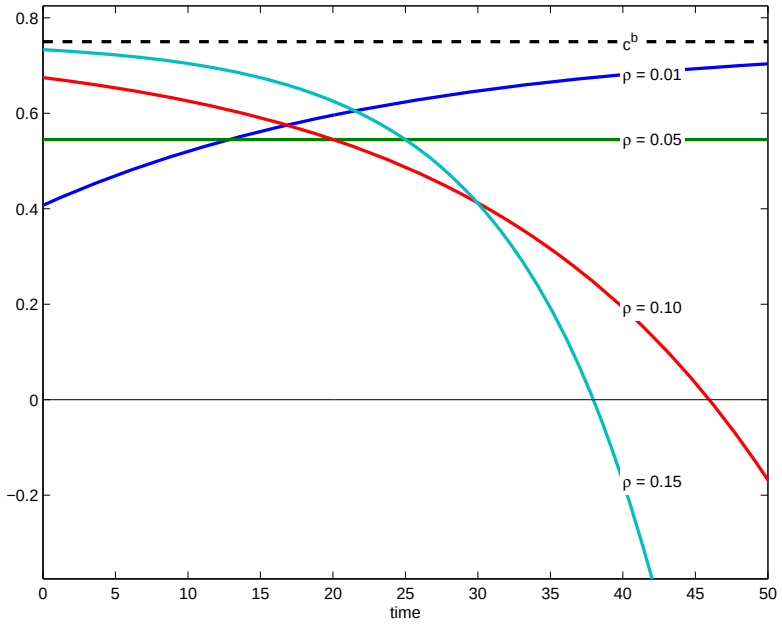

Figure 3: Consumption paths over time. $r = 0.05$, $\gamma = 2$, $T = 50$, ρ as indicated.

consumption path that spends the annuity value of wealth,

$$c_t = c_0 = W\frac{r^f}{1 - e^{-r^f T}}.$$

If $\rho < r^f$, the investor wants a rising consumption profile. $(\rho - r^f)$ in (63) is negative, so consumption starts below the bliss point c^b (assuming, as always inadequate wealth to finance c^b directly), and then exponentially approaches c^b over time. The investor saves early in life and builds up wealth to finance larger consumer later on. $\rho > r^f$ means the consumer is impatient, and wants a *declining* consumption path. $(\rho - r^f)$ in (63) is positive, so consumption begins above the annuity value of wealth, and diverges exponentially downward from c^b. This consumer is borrowing early in life and then repaying it by dramatic consumption declines later in life. There is nothing wrong with negative consumption under quadratic utility, and the extreme parameter values $\rho = 10\%$, $r^f = 5\%$ and above produce such values.

Figure 4 illustrates this point by showing how the consumption paths vary as we change *horizon.* I do not plot a $\rho < r^f$ case; the paths rise with time and quickly approach their limit, which also rises with time. For $\rho = r^f$, in the top left panel, consumption is flat. It is a bit larger for smaller horizons T, as wealth can support a slightly larger consumption for a short time. The top right panel presents the first interesting case, $r^f < \rho < 2r^f$. Here, consumption is declining with horizon, but still approaches its sensible infinite-period limit. $\rho = 2r^f$ in the bottom left panel is the knife-edge case in which the $T \to \infty$ limit is $c = c^b$ for any initial wealth. We see that for any fixed t, consumption steadily rises as T rises and approaches its limit c^b. However, the larger consumption at any t gets paid for more and more dramatic consumption declines later in life. The declines in consumption are dramatic as the horizon T approaches. For $\rho = 0.10$, the final consumption is -0.2 in the $T = 50$ case, -6.5 in the $T = 100$ case, -5,366 in the T=250 case and -720,048,993 in the $T = 500$ case. These are very impatient investors,

they borrow tremendously against a far-off future in order to consume a bit more now.

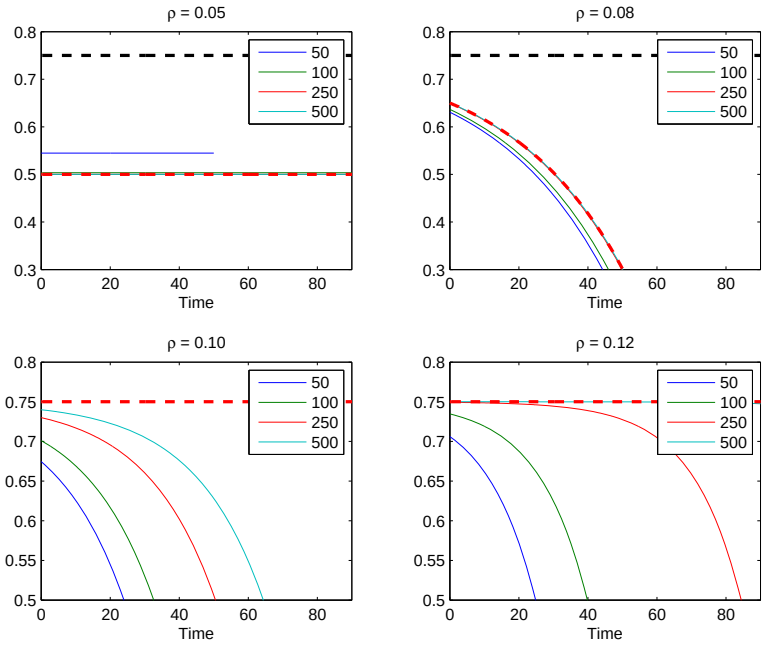

Figure 4: Consumption paths as horizon T is increased. In each case, $r = 0.5$, and ρ as indicated. The dashed horizontal line plots c^b. The curved (red) horizontal line plots the limit as $T \to \infty$ of the consumption path.

Examining Figure 4 or equation (62), we see that the limit $T \to \infty$ is well-behaved if $\rho < 2r^f$. When $\rho \geq 2r^f$, however, we see that for any *finite* T, the increase in consumption c_t towards c^b is always paid for by dramatic declines in consumption further out. If we fix t and let $T \to \infty$, consumption c_t approaches c^b. Thus, indeed, $\lim_{T\to\infty} c_t = c^b$. However, this consumption is supported at each T in each case by borrowing, which is paid back by dramatic consumption declines later in life, which is indefinitely postponed at the limit *point*. For each *finite* value of the horizon T, this is a perfectly well-specified model. Wealth is exactly exhausted in the finite lifetime, so $W_T = 0$, and naturally $p(W_T) = 0$ However, though $p(W_T) = 0$ for each *finite* value of T, the limit *point* violates the transversality condition, with $\lim_{T\to\infty} p(W_T) \neq 0$. Thus, the limit *point* is not a valid solution of the infinite-period model.

The "problem" is then clear in this case. The investor is borrowing more and more as T increases. To generate a sensible infinite-period model, we have to impose a borrowing constraint $W > -\bar{W}$, in order to avoid this arbitrage opportunity in the infinite-period model that is not present in any of the finite-horizon solutions. We did not need to impose that constraint for $\rho < 2r^f$, since our solution method that ignored the borrowing constraint produced an infinite-period solution that did not violate that constraint. Alas, when $\rho \geq 2r^f$, this happy state of affairs no longer holds. To solve a quadratic model with those parameters, we must explicitly introduce the borrowing constraint, and we must expect it to bind. In the presence of a constraint, the investor follows a declining consumption profile that hits a lower bound when the constraint binds. At that point the shadow value of the constraint adds to the interest rate, leading to an optimal constant consumption value.

Risk premium, $\mu \neq 0$, $\rho = r^f$

This much is pretty standard from the permanent income model, featuring constant interest rates and no asset markets. The presence of risky assets is novel. It is also the ingredient that causes the quantitative trouble. We don't need $\rho > 2r^f$ to fit the world, but the squared sharpe ratio of 0.10-0.25 is much higher than the riskfree rate. Having understood the issues raised by $\rho \neq r^f$, we can focus on $\rho = r^f$ to simplify the formulas for $\mu \neq 0$.

In this case, we have

$$\begin{aligned} c_t &= c^b - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W\right] \left(\frac{r^f - \mu'\Sigma^{-1}\mu}{1 - e^{\left[-r^f + \mu'\Sigma^{-1}\mu\right]T}}\right) e^{-\frac{1}{2}\mu'\Sigma^{-1}\mu t - \mu'\Sigma^{-1}\sigma \int_0^t dz_t}; \text{ if } 0 \neq r^f - \mu'\Sigma^{-1}\mu \\ c_t &= c^b - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W\right] \left(\frac{1}{T}\right) e^{\left(-\frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\sigma \int_0^t dz_t}; \text{ if } 0 = r^f - \mu'\Sigma^{-1}\mu \end{aligned}$$

Again, to understand the evolution of consumption over time for fixed T, it's easier still to write as in (63)

$$c^b - c_t = \left(c^b - c_0\right) e^{-\frac{1}{2}\mu'\Sigma^{-1}\mu t - \mu'\Sigma^{-1}\sigma \int_0^t dz_t} \tag{64}$$

Now, consumption is not *expected* to drift up or down, $E(c_t) = c_0$. However, it will do so stochastically depending on the outcomes of the risky assets. If the risky assets do well, consumption will drift up towards the bliss point, and vice versa. The essential issue is the same however – since consumption is an exponential function of shocks, there is a big *left tail*. As T increases, c_0 approaches c^b, so the first term in (64) gets smaller. But the variance of the shock rises. The result is a distribution of long-horizon consumption that bunches up against c^b, but with a very long left tail. When $r^f < \mu'\Sigma^{-1}\mu$, the investor again finances large consumption by borrowing and investing in stocks. He then waits for a sufficiently large stock return to pay off the loan. If stocks go against him, he increases his portfolio position "doubling up." In a finite-time model, he faces the danger that this stock return does not happen; in that case he accepts a dramatic consumption plunge late in life in order to pay back the loan. In the limit *point*, he waits forever for stocks to bail him out, again producing an apparent arbitrage opportunity by violating standard trading limits.

To see this behavior, Figure 5 plots consumption in periods $t = 5$, $t = 10$ in models with horizon $T = 5$ and $T = 10$. Comparing the two $t = 5$ lines, we see as expected that the distribution of consumption approaches $c_5 = c^b$ as the horizon T increases. However, we see also that the left tail of final period $T = 10$ consumption crosses that of $T = 5$ consumption. As the horizon increases, this pattern continues: an increasingly small but increasingly disastrous end-of life consumption in the event of a stock market decline pays for consumption above the annuity value of wealth early in life. This time of paying the piper disappears at the limit *point* $T = \infty$.

To give a more concrete picture, we can examine wealth and the asset market position as well. I found it most intuitive to look at the consumption decision rule in terms of wealth:

$$\begin{aligned} c^b - c_t &= \frac{\mu'\Sigma^{-1}\mu - r^f}{e^{\left(\mu'\Sigma^{-1}\mu - r^f\right)(T-t)} - 1} \left(W_t^b - W_t\right) \text{ if } r^f \neq \mu'\Sigma^{-1}\mu \\ c^b - c_t &= \frac{1}{T-t} \left(W_t^b - W_t\right) \text{ if } r^f = \mu'\Sigma^{-1}\mu \end{aligned}$$

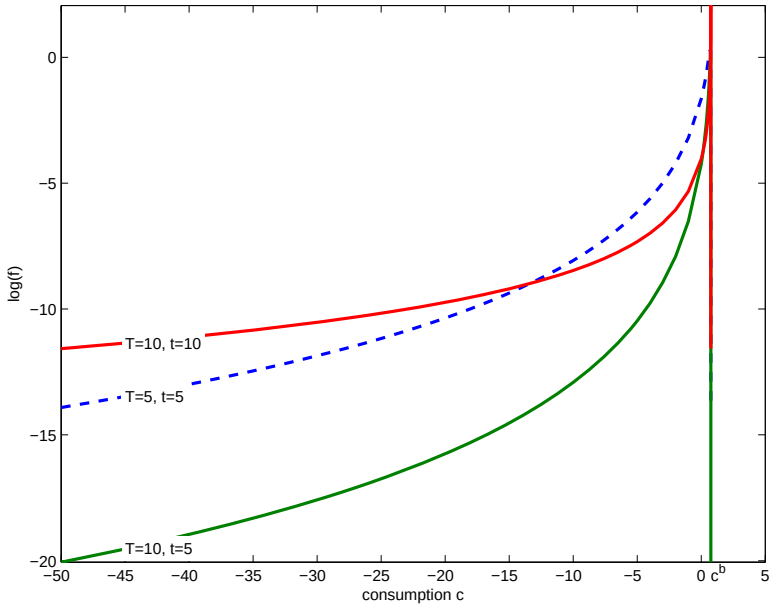

Figure 5: Log density of consumption at the last moment of a 5-year ($T = 5$) and 10-year ($T = 10$) model. Parameters are $r = \rho = 0.5$, $\mu = 0.08$, $\sigma = 0.16$, $\gamma = 2$, $W = 10$ and hence $c^b = 0.75$.

$$W_t^b \equiv c^b \frac{1 - e^{-r^f(T-t)}}{r^f}$$

Early in life t is small, wealth is large, and the denominator is a large number. Hence, the investor consumes very near c^b. Later in life, if stock returns have been strong, W_t is still large and this strategy works. If stocks have not done so well, however, W_t is now small or negative, and possibly very negative. As t rises, the denominator approaches 0, so the investor will undertake a drastic consumption decline late in life to pay back his debts and leave $W_T = 0$. This consumption decline is pushed ever later, and is not present at the limit *point* $T = \infty$.

The weights in risky assets are

$$\omega_t = \left(W_t^b - W_t\right) \mu' \Sigma^{-1}$$

The portfolio is mean-variance efficient, as usual. However, if stocks have not done well, the investor takes larger and larger positions, "doubling up" to try to use asset markets to get out of the hole. In the limit *point*, he can always wait long enough to get out of debt by this strategy, though for each finite T, he must plan a disastrous decline in consumption late in life to pay back debts. Again, to solve the infinite-period model with these parameter values, we must impose the standard wealth constraint and constraint on the size of trades, and expect them to bind.

Derivations

For the finite-horizon model, the discount factor is still

$$x_t^* = e^{\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\sigma\int_0^t dz_t}.$$

Its price is

$$\begin{aligned} p(x^*) &= E\int_0^T e^{-\rho t}x_t^{*2}dt = E\int_0^T e^{-\rho t}e^{2\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - 2\mu'\Sigma^{-1}\sigma\int_0^t dz_t}dt \\ &= \int_0^T e^{-\rho t}e^{2\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t + 2\mu'\Sigma^{-1}\mu t}dt = \int_0^T e^{\left[(\rho - 2r^f) + \mu'\Sigma^{-1}\mu\right]t}dt \end{aligned}$$

$$\begin{aligned} p(x^*) &= \frac{1 - e^{\left[(\rho - 2r^f) + \mu'\Sigma^{-1}\mu\right]T}}{2r^f - \rho - \mu'\Sigma^{-1}\mu} \text{ if } 2r^f - \rho - \mu'\Sigma^{-1}\mu \neq 0 \\ p(x^*) &= T \text{ if } 2r^f - \rho - \mu'\Sigma^{-1}\mu = 0 \end{aligned}$$

The price $p(x^*)$ is now finite in all cases – finite horizon models are well behaved. In the troublesome cases $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$, however, $p(x^*)$ grows with horizon. Of course, we have

$$p(c^b) = c^b E\int_0^T e^{-r^f t}dt = c^b\frac{1 - e^{-r^f T}}{r^f}.$$

The optimal consumption stream is

$$c_t = c^b - \left[p(c^b) - W\right]y_t^*$$

$$c_t = c^b - \left[c^b\frac{1 - e^{-r^f T}}{r^f} - W\right]\frac{x_t^*}{p(x^*)}$$

or, explicitly,

$$\begin{aligned} c_t &= c^b - \left[c^b\frac{1 - e^{-r^f T}}{r^f} - W\right]\left(\frac{2r^f - \rho - \mu'\Sigma^{-1}\mu}{1 - e^{\left[(\rho - 2r^f) + \mu'\Sigma^{-1}\mu\right]T}}\right)e^{\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\sigma\int_0^t dz_t}; \text{ if } \rho \neq 2r^f - \mu'\Sigma^{-1}\mu \quad (6 \\ c_t &= c^b - \left[c^b\frac{1 - e^{-r^f T}}{r^f} - W\right]\left(\frac{1}{T}\right)e^{\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\sigma\int_0^t dz_t}; \text{ if } \rho = 2r^f - \mu'\Sigma^{-1}\mu \end{aligned}$$

To examine portfolio rules, we can find wealth at time t as

$$W_t = p_t(c_t) = p_t\left(c^b\right) - \left[c^b\frac{1 - e^{-r^f T}}{r^f} - W\right]\frac{p_t(x^*)}{p(x^*)}$$

where the time-t price of the remaining stream x_s^*, $t \leq s \leq T$ is

$$\begin{aligned} p_t(x^*) &= \frac{1}{e^{-\rho t}x_t^*}E_t\int_t^T e^{-\rho s}x_s^{*2}ds \\ &= \frac{1}{e^{-\rho t}x_t^*}E_t\int_t^T e^{-\rho s}x_t^{*2}e^{2\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)(s-t) - 2\mu'\Sigma^{-1}\sigma\int_t^s dz_\tau}ds \\ &= x_t^*\int_t^T e^{\left(\rho - 2r^f + \mu'\Sigma^{-1}\mu\right)(s-t)}ds \end{aligned}$$

$$\begin{aligned} p_t(x^*) &= x_t^* \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{2r^f-\rho-\mu'\Sigma^{-1}\mu}; \text{if } \; 2r^f-\rho-\mu'\Sigma^{-1}\mu \neq 0 \\ &= x_t^*\,(T-t)\,; \text{ if } \; 2r^f-\rho-\mu'\Sigma^{-1}\mu = 0. \end{aligned}$$

Thus, we have

$$\begin{aligned} W_t &= c^b \frac{1-e^{-r^f(T-t)}}{r^f} - \left[c^b \frac{1-e^{-r^fT}}{r^f} - W\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} x_t^* \text{if } \; 2r^f-\rho-\mu'\Sigma^{-1}\mu \neq 0 \quad (66) \\ &= c^b \frac{1-e^{-r^f(T-t)}}{r^f} - \left[c^b \frac{1-e^{-r^fT}}{r^f} - W\right] \frac{T-t}{T} x_t^*; \text{if } \; 2r^f-\rho-\mu'\Sigma^{-1}\mu = 0 \end{aligned}$$

or, explicitly,

$$\begin{aligned} W_t &= c^b \frac{1-e^{-r^f(T-t)}}{r^f} - \left[c^b \frac{1-e^{-r^fT}}{r^f} - W\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} e^{\left(\rho-r^f-\frac{1}{2}\mu'\Sigma^{-1}\mu\right)t-\mu'\Sigma^{-1}\sigma\int_0^t dz_t} \quad (67) \\ &= c^b \frac{1-e^{-r^f(T-t)}}{r^f} - \left[c^b \frac{1-e^{-r^fT}}{r^f} - W\right] \frac{T-t}{T} e^{\left(\rho-r^f-\frac{1}{2}\mu'\Sigma^{-1}\mu\right)t-\mu'\Sigma^{-1}\sigma\int_0^t dz_t} \end{aligned}$$

Note $W_T = 0$ – all wealth *is* exhausted.

To express consumption in terms of wealth, we substitute (67) in (65), yielding

$$\begin{aligned} c_t &= c^b - \frac{\left(2r^f-\rho-\mu'\Sigma^{-1}\mu\right)}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}} \left(c^b \frac{1-e^{-r^f(T-t)}}{r^f} - W_t\right); \quad \text{if } \rho \neq 2r^f-\mu'\Sigma^{-1}\mu \\ c_t &= c^b - \frac{1}{T-t}\left(c^b \frac{1-e^{-r^f(T-t)}}{r^f} - W_t\right); \text{ if } \rho = 2r^f-\mu'\Sigma^{-1}\mu \end{aligned}$$

To find portfolios, we look at the diffusion component of dW. From (66),

$$dW_t = ...dt - \left[c^b \frac{1-e^{-r^fT}}{r^f} - W\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} dx_t^*$$

$$dW_t = ...dt + \left[c^b \frac{1-e^{-r^fT}}{r^f} - W\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} x_t^* \mu'\Sigma^{-1}\sigma dz_t$$

Thus the portfolio weight on risky assets is.

$$w_t = \left[c^b \frac{1-e^{-r^fT}}{r^f} - W\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} x_t^* \mu'\Sigma^{-1}$$

It is useful to express this weight with wealth W_t as the state variable. Substituting from (67) above

$$w_t = \left(c^b \frac{1-e^{-r^f(T-t)}}{r^f} - W_t\right) \mu'\Sigma^{-1}$$

We can also write this in a very familiar form as

$$w_t = \frac{\left(c^b - W_t \frac{r^f}{1-e^{-r^f(T-t)}}\right)}{W_t \frac{r^f}{1-e^{-r^f(T-t)}}} W_t \mu'\Sigma^{-1} = \frac{1}{\gamma_t} W_t \mu'\Sigma^{-1}$$

with a local risk aversion coefficient γ_t However, since most of the analysis concerns $W_t < 0$ where this definition of γ doesn't make any sense, this expression is not so useful in studying limits.

2.6 Lognormal returns and mean-variance frontiers

The strange behavior of the long-run mean-variance frontier in the iid lognormal environment really has nothing to do with long-run mean-variance frontiers per se; it is mirrored in the traditional T period arithmetic mean-variance frontier generated by lognormal returns. The apparently strange portfolio of the quadratic utility investor – either shorting a portfolio that is short the risky asset, or adopting a time-varying weight which becomes more risk averse as wealth rises – are exactly how you must produce a $T-$period mean-variance efficient portfolio in a lognormal environment. The message is simply that lognormal returns and mean-variance analysis do not work well together. The long positive tail of a lognormal adds variance without adding much mean. Given a choice via dynamic trading, an investor interested in mean and variance adopts a dynamic strategy that reduces this right tail, by scaling back investment in good states, and expands the limited left tail of a lognormal, by expanding investment in bad states. (Campbell and Viceira (2006) look at the mean-variance properties of long-run *log* returns, which are well-behaved, but not the solution to a portfolio problem.)

To understand this claim, let us examine the mean-variance properties of T period returns, in the standard lognormal iid setup. Suppose there is a single risky asset that follows

$$\frac{dS}{S} = \mu dt + \sigma dB$$

and a constant risk free rate $r^f dt$. We have the following facts, all derived below:

1. The mean, standard deviation and Sharpe ratio of the risky asset are given by

$$\begin{aligned} E\left(\frac{S_T}{S_0}\right) &= e^{\mu T} \\ \sigma\left(\frac{S_t}{S_0}\right) &= e^{\mu T}\sqrt{e^{\sigma^2 T} - 1} \end{aligned}$$

2. The Sharpe ratio of the risky asset is

$$SR = \frac{1 - e^{-\left(\mu - r^f\right)T}}{\sqrt{e^{\sigma^2 T} - 1}}$$

   The Sharpe ratio first rises, then declines with horizon. The limiting Sharpe ratio is zero

$$\lim_{T\to\infty} SR = 0$$

3. The maximum attainable Sharpe ratio from trading continuously in the risky and risk free asset, however, is

$$SR_{\max} = \sqrt{e^{\frac{(\mu - r)^2}{\sigma^2}T} - 1}$$

4. This quantity unambiguously rises with horizon, and

$$\lim_{T\to\infty} SR_{\max} = \infty.$$

5. Two interesting portfolios on the mean-variance frontier, which attain this maximum Sharpe ratio are

(a) First, a constantly–rebalanced *short* position in the risky asset, a portfolio that is constantly *short* the growth-optimal portfolio.

$$\frac{dV_t}{V_t} = \left[r^f - \frac{\left(\mu - r^f\right)^2}{\sigma^2} \right] dt - \frac{\mu - r^f}{\sigma^2} \sigma dB_t$$

i.e. a portfolio weight

$$\omega = -\frac{\mu - r^f}{\sigma^2}.$$

This portfolio is on the lower portion of the time-T mean-variance frontier. It is the only constant-weight portfolio on the time-T mean-variance frontier.

(b) Second, to express a portfolio on the upper portion of the mean-variance frontier, we can either short the portfolio V that shorts the risky asset,

$$W_t = 2e^{r^f t} - \frac{V_t}{V_0}$$

or, equivalently, express the result as a portfolio that "doubles up," investing more in the risky asset as wealth declines, but investing less as wealth increases, to the point that wealth never grows faster than twice the riskfree rate,

$$dW_t = \left\{ r^f W_t + \left(2e^{r^f t} - W_t\right) \frac{\left(\mu - r^f\right)^2}{\sigma^2} \right\} dt + \left(2e^{r^f t} - W_t\right) \frac{\mu - r^f}{\sigma^2} \sigma dB_t.$$

6. A quadratic utility investor, whose objective is

$$\max E\left[\left(-\frac{1}{2}\right) \left(W^b - W_T\right)^2 \right]$$

holds a portfolio with weight in the risky asset given by

$$w = \left(W^b e^{-r(T-t)} - W_t\right) \frac{\mu - r}{\sigma^2}$$

He increases his weight in the risky asset as wealth falls, decreasing it to zero as wealth rises towards the value that guarantees bliss point terminal wealth.

Derivation

The mean, variance and Sharpe ratio of the risky asset follow from

$$\begin{aligned}
\frac{S_t}{S_0} &= e^{\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma \int_0^t dB} \\
E\left(\frac{S_t}{S_0}\right) &= e^{\mu t} \\
E\left[\left(\frac{S_t}{S_0}\right)^2\right] &= e^{\left(2\mu + \sigma^2\right)t} \\
\sigma\left(\frac{S_t}{S_0}\right) &= e^{\mu t}\sqrt{e^{\sigma^2 t} - 1}.
\end{aligned}$$

Notice the standard deviation grows faster than the mean.

The Sharpe ratio is

$$SR = \frac{e^{\mu t} - e^{r^f t}}{e^{\mu t}\sqrt{e^{\sigma^2 t} - 1}} = \frac{1 - e^{-\left(\mu - r^f\right)t}}{\sqrt{e^{\sigma^2 t} - 1}}$$

To find the maximum Sharpe ratio, we characterize the discount factor, and use the theorem that maximum Sharpe ratio is determined by the variance of the discount factor:

$$\begin{aligned}
\frac{d\Lambda}{\Lambda} &= -r^f dt - \frac{\mu - r^f}{\sigma} dB \\
\frac{\Lambda_t}{\Lambda_0} &= e^{\left(-r - \frac{1}{2}\frac{\left(\mu - r^f\right)^2}{\sigma^2}\right)t - \frac{\mu - r^f}{\sigma}\int_0^t dB} \\
E\left(\frac{\Lambda_t}{\Lambda_0}\right) &= e^{-r^f t} \\
E\left[\left(\frac{\Lambda_t}{\Lambda_0}\right)^2\right] &= e^{\left(-2r^f + \frac{(\mu - r)^2}{\sigma^2}\right)t} \\
\sigma\left(\frac{\Lambda_t}{\Lambda_0}\right) &= e^{-r^f t}\sqrt{e^{\frac{\left(\mu - r^f\right)^2}{\sigma^2}t} - 1} \\
SR_{max} &= \frac{\sigma\left(\frac{\Lambda_t}{\Lambda_0}\right)}{E(\frac{\Lambda_t}{\Lambda_0})} = \sqrt{e^{\frac{\left(\mu - r^f\right)^2}{\sigma^2}t} - 1}
\end{aligned}$$

The max sharpe ratio increases without bound!

To find the Sharpe-ratio maximizing payoff/portfolio, we can exploit the fact that Λ_T/Λ_0 is on the mean-variance frontier, and then find the self-financing strategy that yields the final payoff Λ_T/Λ_0. The answer is

$$V_t = \frac{\Lambda_t}{\Lambda_0} e^{-2\left[r - \frac{1}{2}\frac{\left(\mu - r^f\right)^2}{\sigma^2}\right](T-t)}.$$

with differential characterization

$$\frac{dV}{V} = \left[r^f - \frac{\left(\mu - r^f\right)^2}{\sigma^2}\right] dt - \frac{\mu - r^f}{\sigma} dB$$

We can find this answer either by guessing the form $V_t = e^{-\eta(T-t)}\Lambda_T/\Lambda_0$, Ito's lemma, and imposing that drift and diffusion coefficients must be of the form $dV/V = \left[r^f + \omega\left(\mu - r^f\right)\right] dt + \omega\sigma dB$, or by evaluating

$$V_t = E_t\left(\frac{\Lambda_T}{\Lambda_t}\frac{\Lambda_T}{\Lambda_0}\right) = \frac{\Lambda_t}{\Lambda_0} E_t\left[\left(\frac{\Lambda_T}{\Lambda_t}\right)^2\right]$$

To examine a portfolio on the top of the frontier, we can short V and invest in the T bill rate, producing a wealth process

$$W_t = 2e^{r^f t} - \frac{V_t}{V_0}$$

One way to do this, of course is to explicitly short V: you *short* a portfolio that is always *short* stocks in the amount $\frac{\mu - r^f}{\sigma^2}$. Or we can unwind all this shorting and express the same thing as a *time-varying* weight that is long the stock. Differentiating the last equation,

$$\begin{aligned} dW_t &= 2r^f e^{r^f t}dt - \frac{V_t}{V_0}\left(\left[r^f - \frac{\left(\mu - r^f\right)^2}{\sigma^2}\right]dt - \frac{\mu - r^f}{\sigma}dB\right) \\ &= 2r^f e^{r^f t}dt - \left(2e^{r^f t} - W_t\right)\left(\left[r^f - \frac{\left(\mu - r^f\right)^2}{\sigma^2}\right]dt - \frac{\mu - r^f}{\sigma}dB\right) \\ &= \left\{r^f W_t + \left(2e^{r^f t} - W_t\right)\frac{\left(\mu - r^f\right)^2}{\sigma^2}\right\}dt + \left(2e^{r^f t} - W_t\right)\frac{\mu - r^f}{\sigma}dB \end{aligned}$$

The quadratic utility investor's problem is

$$\begin{aligned} \max E\left[\left(-\frac{1}{2}\right)(W^b - W_T)^2\right] \text{ s.t. } W_0 &= E\left(\frac{\Lambda_T}{\Lambda_0}W_T\right) \\ W^b - W_T &= \lambda\frac{\Lambda_T}{\Lambda_0} \\ W_T &= W^b - \lambda\frac{\Lambda_T}{\Lambda_0} \end{aligned}$$

His wealth at any time t is

$$W_t = E_t\left(\frac{\Lambda_T}{\Lambda_t}W_T\right) = W^b e^{-r(T-t)} - \lambda E_t\left(\frac{\Lambda_T}{\Lambda_t}\frac{\Lambda_T}{\Lambda_0}\right) = W^b e^{-r(T-t)} - \lambda V_t \qquad (68)$$

Using time-0 wealth to eliminate λ,

$$\begin{aligned} \frac{W^b e^{-rT} - W_0}{V_0} &= \lambda \\ W_t &= W^b e^{-r(T-t)} - \left(W^b e^{-rT} - W_0\right)\frac{V_t}{V_0} \end{aligned}$$

Differentiating,

$$dW_t = rW^b e^{-r(T-t)}dt - \left(W^b e^{-rT} - W_0\right)\frac{dV_t}{V_0}$$

$$dW_t = rW^b e^{-r(T-t)}dt - \left(W^b e^{-rT} - W_0\right)\frac{V_t}{V_0}\left(\left[r - \frac{(\mu - r)^2}{\sigma^2}\right]dt - \frac{\mu - r}{\sigma}dB\right)$$

Using (68),

$$\frac{\left(W^b e^{-r(T-t)} - W_t\right)}{\left(W^b e^{-rT} - W_0\right)} = \frac{V_t}{V_0}$$

we can eliminate V,

$$dW_t = rW^b e^{-r(T-t)}dt - \left(W^b e^{-rT} - W_0\right)\frac{\left(W^b e^{-r(T-t)} - W_t\right)}{\left(W^b e^{-rT} - W_0\right)}\left(\left[r - \frac{(\mu - r)^2}{\sigma^2}\right]dt - \frac{\mu - r}{\sigma}dB\right)$$

rearranging,

$$dW_t = rW_t dt + \left(W^b e^{-r(T-t)} - W_t\right) \frac{(\mu - r)^2}{\sigma^2} dt + \left(W^b e^{-r(T-t)} - W_t\right) \frac{\mu - r}{\sigma} dB$$

We recognize an investment with dollar weights

$$w = \left(W^b e^{-r(T-t)} - W_t\right) \frac{\mu - r}{\sigma^2}$$

in the risky asset.

2.7 Lognormal returns in data?

Lognormal returns pose difficulties for long-run mean-variance analysis. This raises the natural question, how lognormal are returns in actual data? To address this issue, Figure 6 plots the cumulative distribution of CRSP value-weighted returns at one year and 10 year horizons. The axes are stretched so that a normal distribution would plot as a straight line.

At the one-year horizon (top left) we see the familiar "fat left tail" of the actual return (squares) relative to a lognormal model. Interestingly we see a less renown "thin right tail." Stock market booms are *less* frequent than the lognormal model predicts. The top right panel show that one-year returns are in fact much better modeled as *normal* than as lognormal! There is no fat left tail, and only a barely discernible thin right tail relative to a normal distribution.

We are after long-run distributions, and the bottom panels of Figure 6 paint an even more interesting picture. At a 10 year horizon, the "fat left tail" relative to a lognormal has disappeared, but the "thin right tail" is dramatic. We are only just beginning to see the normal distribution's obvious failure on the left tail: it predicts the possibility of returns less than -100%, so the data must have a thinner left tail at some point than the normal predicts. Even the normal distribution now predicts too large a right tail.

These characteristics of long-horizon returns mean that instantaneous returns are not lognormal. A normal return distribution would be generated by a price process of the form

$$dS = \mu dt + \sigma dB$$

i.e.

$$\frac{dS}{S} = \left(\frac{\mu}{S}\right) dt + \left(\frac{\sigma}{S}\right) dB$$

To generate a pure normal, in particular, volatility must increase after market downturns and decrease after good times. This qualitative pattern is, in fact, what we see in the data.

In sum, long-run mean-variance analysis will work a lot better if one models the payoff streams as *normal* rather than *lognormal*. In particular, the large predicted *right* tail of the lognormal distribution is the central source of problems for mean-variance analysis. But while we are used to theoretical models that specify lognormal returns, the data are if anything more consistent with a normal distribution, especially on the troublesome right side. In turn, the failure of lognormality means that stock returns carry some important dynamics, such as time-varying volatility. Alas, this observation precludes simple back of the envelope exercises. The nature of stochastic volatility is central to the results.

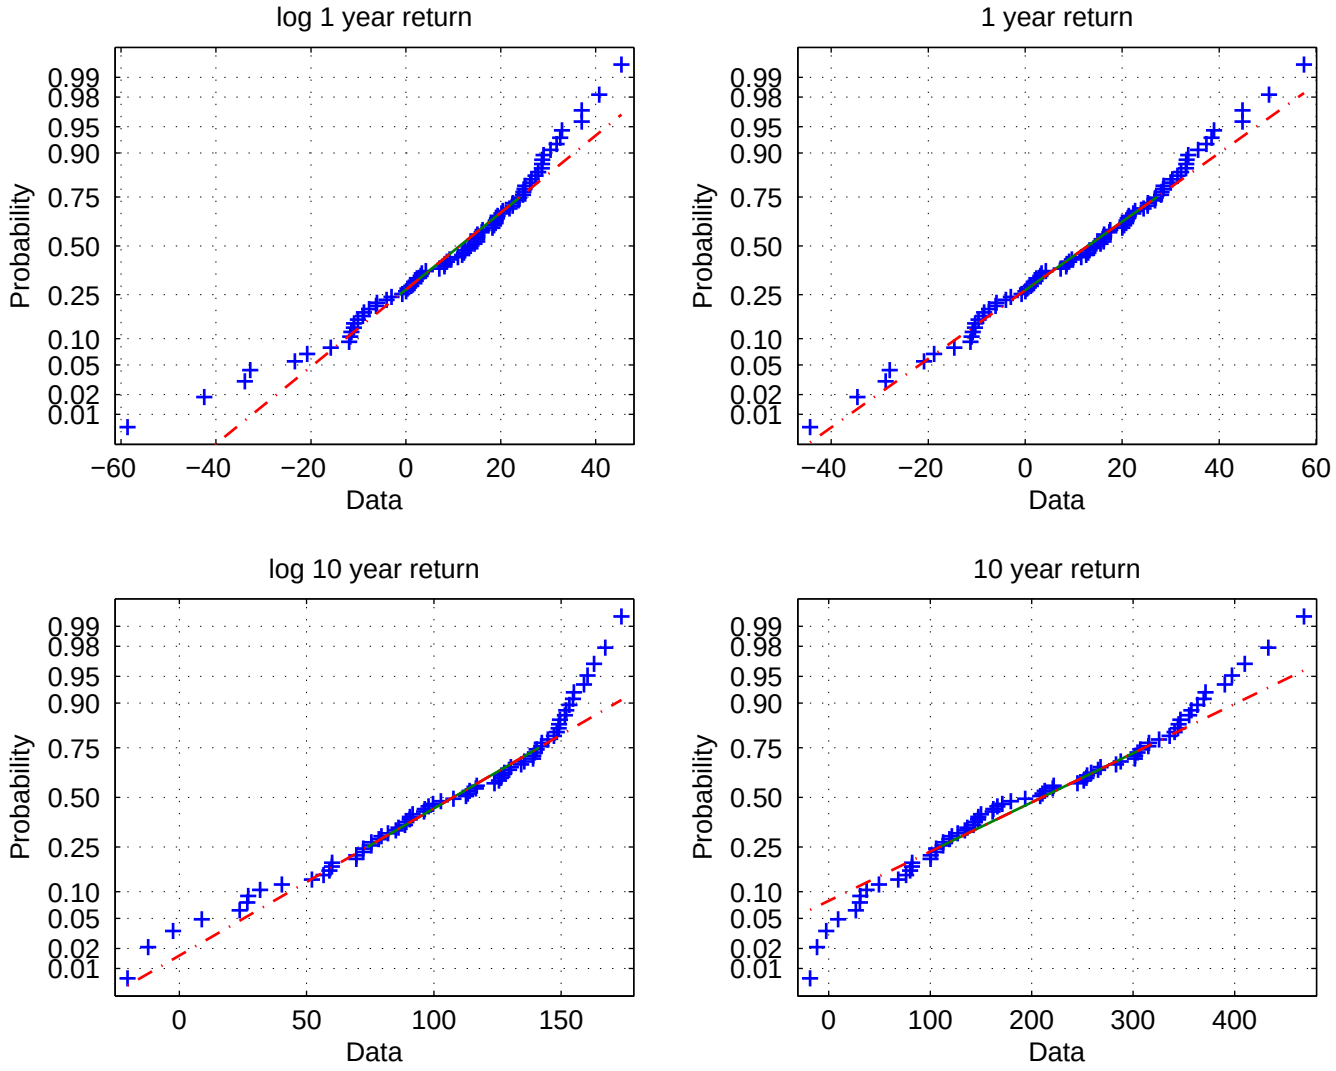

Figure 6: Distribution of CRSP value-weighted return, 1927-2006. Each plot presents the cumulative distrbibution of the data (blue cross), with the y axis scaled so that normally distributed data will fall on a line. The red dashed line is fit to the first and third quartiles to help assess linearity (matlab function normplot). The 10 year returns are overlapping annual observations. The units are percent in all cases.

3 Avoiding returns

The above analysis maintains a close link to traditional dynamic portfolio theory exercises, but in doing so takes us away from the focus on payoffs rather than one-period returns that motivates the whole approach.

Rather than model one-period returns, it is much more attractive to simply model the final payoffs. After all, in the simplest application of one-period portfolio theory, we take the mean and covariance matrix of a set of returns and find the optimal portfolio. Why not similarly

evaluate the long-run mean and long-run covariance matrix of a set of asset yields and directly calculate the optimal payoff? Any dynamics will show up in the difference between long-run and short-run moments, as (say) variance ratios capture autocorrelation. This section implements this simpler idea.

3.1 Payoff streams and payout rules

The first choice in applying the long-run mean / long-run variance idea is what streams of payoffs $\{x_t\}$ to examine.

It's tempting simply to define $\{x_t\}$ as the dividend streams corresponding to the chosen assets. Using dividends directly is not generally such a good choice for portfolio theory, however, as using individual stocks is not such a good choice for one-period portfolio theory. The investor can dynamically reallocate his investment across assets, and he can consume out of capital gains as well as dividends. The optimal portfolio will therefore typically feature a consumption stream that may have little to do with the underlying dividend streams. For example, in the analytic solution to the i.i.d. return case studied above, the investor consumes a proportion of wealth and rebalances continuously, ignoring any distinction between dividends and capital gains[3].

Given that the investor can trade dynamically, it's natural to think of portfolios in terms of their allocations across assets and their payout rates. With underlying assets that follow

$$\begin{aligned} dr_t &= r^f dt + dr_t^e; \\ dr_t^e &= \mu_t dt + \sigma_t dB_t, \end{aligned}$$

the value of an investment characterized by a dollar payout y_t and a dollar portfolio allocation ω_t follows

$$dV_t = \left(r_t^f V_t - y_t\right) dt + \omega_t' dr_t^e;\ V_0 = 1.$$

Since $V_0 = 1$, this strategy generates a yield y_t.

Our objective is to find the optimal payoff y_t^*, or equivalently the set of long-run mean / long-run variance efficient payoffs. Rather than construct a parametric model of one-period returns and the discount factor that prices those returns, as in the last section, I base the analysis on the long-run properties of a set of such yields.

I start, naturally enough with linear payout rules, i.e. generating yields from asset returns by

$$\frac{dV_t^i}{V_t^i} = \left(r_t^f - \alpha^i\right) dt + w^{i\prime} dr_t^e;\ V_0^i = 1.$$

and hence

$$y_t^i = \alpha^i V_t^i$$

[3]Examining dividend streams does make sense for long-run asset pricing as in Hansen Heaton and Li (2005), since in that case we are interested in learning how the fundamental cash flows affect value and it is more important to remove future price fluctuation from the "exogenous" variables.

It's also an attractive idea that firms choose to pay dividend streams that are optimal for consumers. After all, since dividend policy is irrelevant to first order, why not pay dividends that investors want, rather than forcing them to synthesize the dividend stream they want with dynamic trading? If this is true, using dividends would make a lot of sense. I do not use dividends here in part for practical reasons; I tried it with the size and book to market portfolios, but easy rebalancing strategies provided much better long-run mean-variance properties.

$$\frac{dy_t^i}{y_t^i} = \left(r_t^f - \alpha^i\right) dt + w^{i\prime} dr_t^e; \quad y_0 = \alpha^i$$

with α and w constant over time. The solution to the iid problem is of this form, with $\alpha = 2r^f - \rho - \mu'\Sigma^{-1}\mu$. Given a set of asset returns, we can easily find the sample yields of constant payout and constant-weight strategies and their sample long-run means and variances. One must check of course that the initial payoffs satisfy the no arbitrage condition $\lim_{T\to\infty} p(V_T) = 0$ and $\mathcal{E}(y_t^2) < \infty$ (i.e.$\lim_{T\to\infty} e^{-\rho t} E(y_t^2) = 0$). I also add interesting, though ad-hoc dynamic strategies, in which α and w are functions of information at time t.

The payoff space is then linear combinations of the payoffs formed by these ad-hoc weighted portfolios. This linear completion adds possibilities. The payoff to a portfolio that is continuously rebalanced to be half stocks and half bonds is *not* half of a portfolio that is always bonds and half of a portfolio that is always stocks. Similarly, the payoff of a portfolio in which the investor consumes 5% of wealth each year is not the linear combination of the payoff of a portfolio in which he consumes 4% and one in which he consumes 6%. Combining two different constant-payout rules gives rise to a payoff that results from a time-varying portfolio and payout rule.[4] The payout and portfolio weights adjust towards those of the more successful security over time.

Though it adds interesting possibilities, this fact also suggests that we need to start with a fairly rich set of payoffs. If the optimum is (say) a constant 55% in risky assets, then that optimum will be poorly approximated by the linear combination of a 0% and 100% constantly-rebalanced portfolio. A richer set of basis assets, perhaps 40%, 50%, 60% weighted in risky assets, may be necessary to get an adequate approximation. On the other hand, small errors may not be that vital. since we discount, tracking errors that build up over long periods of time are less important to the analysis.

With a *finite* number of basis payoffs, we can then find a discount factor payoff by simple matrix inversion, $x^* = \frac{1}{k} p'\mathcal{E}(xx')^{-1}x$. The result will be an approximation, of course, but potentially a sensible one.

This approach is exactly the analogue of the standard approach to one-period problems. Rather than directly estimate a mean-variance frontier with 6000 individual stocks, plus bonds, options, international stocks, etc., we first summarize the vast cross section by a small number

[4]Suppose two payoffs are generated by

$$\frac{dV_i}{V_i} = \left(r^f - \alpha_i\right) dt + w_i' dr^e; i = 1, 2.$$

Then consider a portfolio $V = V_1 + V_2$. That portfolio obeys

$$\begin{aligned}
dV &= \left(r^f - \alpha_1\right) V_1 dt + V_1 w_1' dr^e + \left(r^f - \alpha_2\right) V_2 dt + V_2 w_2' dr^e \\
dV &= \left[\left(r^f - \alpha_1\right) V_1 + \left(r^f - \alpha_2\right) V_2\right] dt + \left[V_1 w_1' + V_2 w_2'\right] dr^e \\
\frac{dV}{V} &= \left[r^f - \alpha_1 \frac{V_1}{V} - \alpha_2 \frac{V_2}{V}\right] dt + \left[\frac{V_1}{V} w_1' + \frac{V_2}{V} w_2'\right] dr^e \\
\frac{dV}{V} &= \left(r^f - \alpha_t\right) dt + w_t' dr^e
\end{aligned}$$

Thus, the payout rate α_t and portfolio investment w_t vary over time, tending to whichever of V_1 or V_2 grows larger than the other.

of sensible portfolios, we estimate the means and covariance matrix of the portfolios, and we find the frontier by matrix methods.

Finally, this approach can overcome or at least sidestep the problems caused by $p(x^*) = \infty$ in the iid analysis. If the set of basis payoffs x_t^i, formed by a set of arbitrary payout and portfolio rules α^i,w^i, each satisfies $\lim_{T\to\infty} p(V_T^i) = 0$ and $\mathcal{E}(x_t^{i2}) < \infty$ then linear combinations $x_t = ax_t^i + bx_t^j$ also satisfy these conditions, and the limit points of any convergent sequence of such payoffs also satisfies the conditions and so is a traded payoff. $x^* = \frac{1}{k}p'\mathcal{E}(xx')^{-1}x$ *will* be a traded payoff. Not only can we bound payout rules away from payout rates $\alpha \leq 0$ that lead to explosive wealth, we can keep the payout rules a sensible distance away to rule out uninteresting as well as illegal wealth explosions.

3.2 Fixed payout and portfolio weights in the iid world

Before going to multiple assets and real data, I check this approach in the (now) well-studied two-asset iid environment. Here we know the mean-variance frontier is generated by y^* and r^f. If we start with a set of arbitrary fixed portfolio and payout rules that do not include y^*, how close an approximation do we get?

In this environment, we can calculate the long-run mean and long-run variance of fixed-weight portfolios. (Eventually, we will take these directly from data, as we start standard portfolio calculations with sample means and a covariance matrix.) With a fixed payout rate α_i and portfolio weight w_i, value follows

$$\frac{dV_i}{V_i} = \left(r^f - \alpha_i\right)dt + w_i' dr_t^e \tag{69}$$

We generate yields (payoffs to price one securities) with

$$y_i = \alpha_i V_i; V_0 = 1.$$

Hence, the y process follows

$$\frac{dy_i}{y_i} = \left(r^f - \alpha_i + w_i'\mu\right)dt + w_i'\sigma dz;\ ; y_{i0} = \alpha_i$$

Now we can calculate the long-run mean and long-run covariance matrix of such yields[5],

$$\mathcal{E}\left(y_{it}\right) = \frac{\rho\alpha_i}{\rho - r^f + \alpha_i - w'\mu} \tag{70}$$

[5]

$$\begin{aligned}
dE\left(y_{it}\right) &= \left(r^f - \alpha_i + w_i'\mu\right)y_{it}dt \\
E\left(y_{it}\right) &= \alpha_i e^{\left(r^f - \alpha_i + w_i'\mu\right)t} \\
\mathcal{E}\left(y_{it}\right) &= \rho\alpha_i \int_0^\infty e^{-\rho t} e^{\left(r^f - \alpha_i + w_i'\mu\right)t} = \frac{\rho\alpha_i}{\rho - r^f + \alpha_i - w'\mu}
\end{aligned}$$

$$\begin{aligned}
d\left(y_i y_j\right) &= y_i dy_j + y_j dy_i + dy_i dy_j \\
dE\left(y_i y_j\right) &= y_i E\left(dy_j\right) + y_j E\left(dy_i\right) + E\left(dy_i dy_j\right) \\
&= y_i y_j\left[\left(r^f - \alpha_i + w_i'\mu\right) + \left(r^f - \alpha_j + w_j'\mu\right) + w_i'\Sigma w_j\right] \\
&= y_i y_j\left[2r^f - \left(\alpha_i + \alpha_j\right) + \left(w_i + w_j\right)'\mu + w_i'\Sigma w_j\right] \\
E(y_{it}y_{jt}) &= \alpha_i\alpha_j e^{\left[2r^f - \left(\alpha_i + \alpha_j\right) + \left(w_i + w_j\right)'\mu + w_i'\Sigma w_j\right]t}
\end{aligned}$$

$$\mathcal{E}(y_{it}y_{jt}) = \frac{\rho\alpha_i\alpha_j}{\rho - 2r^f + (\alpha_i + \alpha_j) - (w_i + w_j)'\mu - w_i'\Sigma w_j} \tag{71}$$

Not all payout and portfolio rates are legal – we need to ensure that the payoff is square integrable $\mathcal{E}(y^2) < \infty$ and we need to ensure that the time-zero value of wealth of the underlying portfolio tends to zero $\lim_{T\to\infty} p(V_T) = 0$. From (70), the former condition implies

$$\mathcal{E}(y^2) < \infty : \quad \alpha_i > r^f + w_i'\mu + \frac{1}{2}w_i'\Sigma w_j - \frac{1}{2}\rho$$

and the latter only requires[6]

$$\lim_{T\to\infty} p(V_T) = 0 : \quad \alpha_i > 0.$$

Both conditions ensure that we pull money out of the investment fast enough as wealth rises, or put money back in fast enough if wealth becomes large and negative.

If we allow all $\alpha > 0$ and $w \in R^N$, the mean-variance frontier is spanned by the risk free rate and the minimum second moment portfolio,

$$\min_{\{\alpha,w\}} \mathcal{E}(y_t^2) = \min_{\{\alpha,w\}} \frac{\alpha^2}{\rho - 2r^f + 2\alpha - 2w'\mu - w'\Sigma w}$$

which gives

$$w = -\Sigma^{-1}\mu$$
$$\alpha = 2r^f - \rho - \mu'\Sigma^{-1}\mu.$$

$$\begin{aligned}
\mathcal{E}(y_{it}y_{jt}) &= \rho\alpha_i\alpha_j \int e^{-\rho t} e^{\left[2r^f - (\alpha_i+\alpha_j) + (w_i+w_j)'\mu + w_i'\Sigma w_j\right]t} dt \\
&= \frac{\rho\alpha_i\alpha_j}{\rho - 2r^f + (\alpha_i + \alpha_j) - (w_i + w_j)'\mu - w_i'\Sigma w_j}
\end{aligned}$$

[6]The value of the strategy that delivers y_{it} is simply x_{it} with

$$\frac{dx_i}{x_i} = \left(r^f - \alpha_i\right)dt + w_i'dr^e;\ x_{i0} = 1$$

We have

$$\begin{aligned}
p(x_{it}) &= e^{-\rho t}E(x_t^* x_{it}) \\
\frac{d(x_t^* x_{it})}{x_t^* x_{it}} &= \frac{dx_t^*}{x_t^*} + \frac{dx_{it}}{x_{it}} + \frac{dx_t^* dx_{it}}{x_t^* x_{it}} \\
\frac{d(x_t^* x_{it})}{x_t^* x_{it}} &= \left(\rho - r^f\right)dt - \mu'\Sigma^{-1}\sigma dz + \left(r^f - \alpha_i + w_i'\mu\right)dt + w_i'\sigma dz - \mu' w_i dt \\
\frac{dE(x_t^* x_{it})}{x_t^* x_{it}} &= (\rho - \alpha_i)\,dt \\
e^{-\rho t}E(x_t^* x_{it}) &= e^{-\rho t}e^{(\rho-\alpha_i)t} = e^{-\alpha_i t}.
\end{aligned}$$

Thus,

$$\lim_{T\to\infty} p(x_{it}) = 0 \text{ iff } \alpha_i > 0.$$

Both conditions also hold if we add constants to the payout rules, i.e.

$$dV = \left[r^f V - (a + \alpha V)\right]dt + \left[\omega + wV\right]'dr^e,$$

i.e. only the slopes α, w on V matter for square-integrability and no arbitrage.

This is just our friend y^* with

$$\mathcal{E}\left(y_t^*\right) = \frac{\rho}{r^f}\left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)$$

$$\mathcal{E}(y_t^{*2}) = \rho\left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)$$

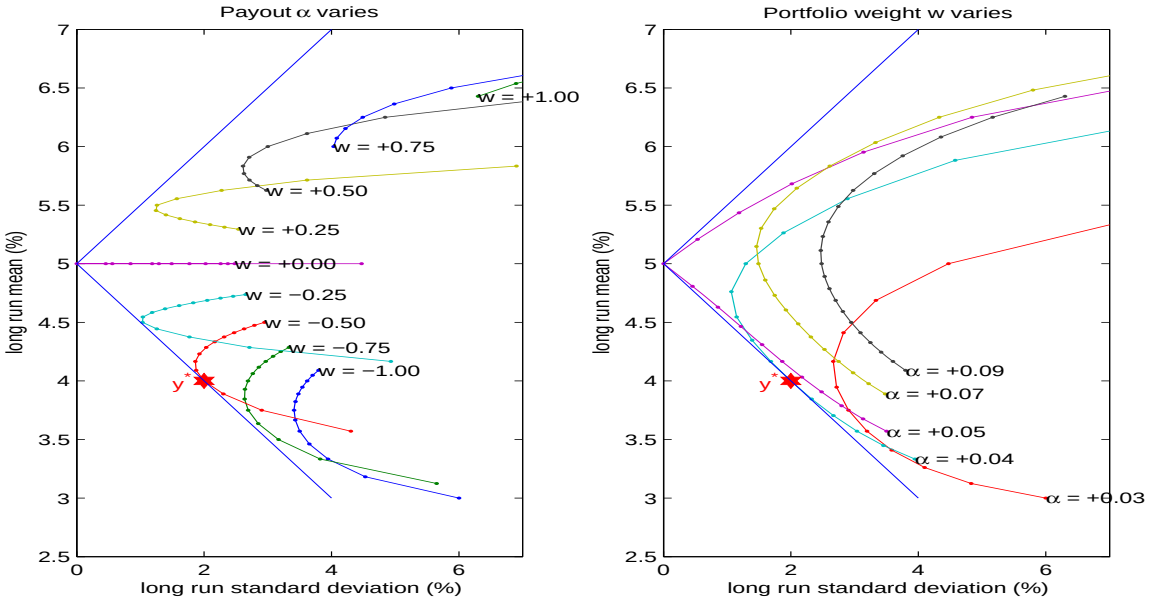

Figure 7: Long-run mean-variance frontiers. The dots are the means and variances of portfolio strategies with constant payout rates α and constantly rebalanced weights w in the risky asset. The left panel connects portfolios with the same portfolio weight w, the right panel connects portfolios with the same payout rate α. The red star gives the minimum second-moment portfolio y^*. Parameters $\mu = 0.02, \rho = r^f = 0.05, \sigma = 0.20$

Figure 7 presents long-run means and standard deviations in the iid economy, for a variety of fixed portfolio allocation w and payout rules α. The frontier in Figure 7 is generated by combinations of the risk free rate together with y^* which here occurs at $w = -0.50$, $\alpha = 0.04$.

The figure shows that many portfolios are close to mean-variance efficient. In particular, when the payout rate $\alpha \approx 0.04$, a wide variety of portfolio rules generate close-to efficient portfolios. This suggests that a relatively sparse sprinkling of ad-hoc portfolio weights can do a good job of spanning the full set.

To examine this question in more detail, suppose we start with a set of payout rules and weights that does *not* include the optimum y^*. How close to that optimum will we get with linear combinations of a few arbitrary payoffs, and then construct an approximate y^* as a linear combination of those arbitrary payoffs by

$$y^* = \frac{1'\mathcal{E}(yy')^{-1}y}{1'\mathcal{E}(yy')^{-1}1}? \tag{72}$$

Figure 8 presents a calculation. The basis assets here have 3% and 5% payout rates and portfolio weights -1, -0.75, -0.25, 0, in both cases avoiding the exact optimum of a 4% payout and -0.50 portfolio weight. The blue dots give the long run means and standard deviations of these basis payoffs. The blue triangle gives the mean-variance efficient portfolio formed from these basis assets by (72). The red line gives the exact mean-variance efficient frontier generated by a 4% payout rate and -0.5 portfolio weight. As the graph shows, the approximate calculation gives almost exactly the same mean-variance trade-off as the exact calculation. The approximate Sharpe ratio is 0.494 rather than 0.500. The numbers next to each blue dot give the weight of each basis asset in the mean-variance efficient portfolio. The two closest portfolios ($\alpha = 0.05, w = -0.25$ and $w = -0.75$) each give about 1/2 weight with smaller weights on the other portfolios.

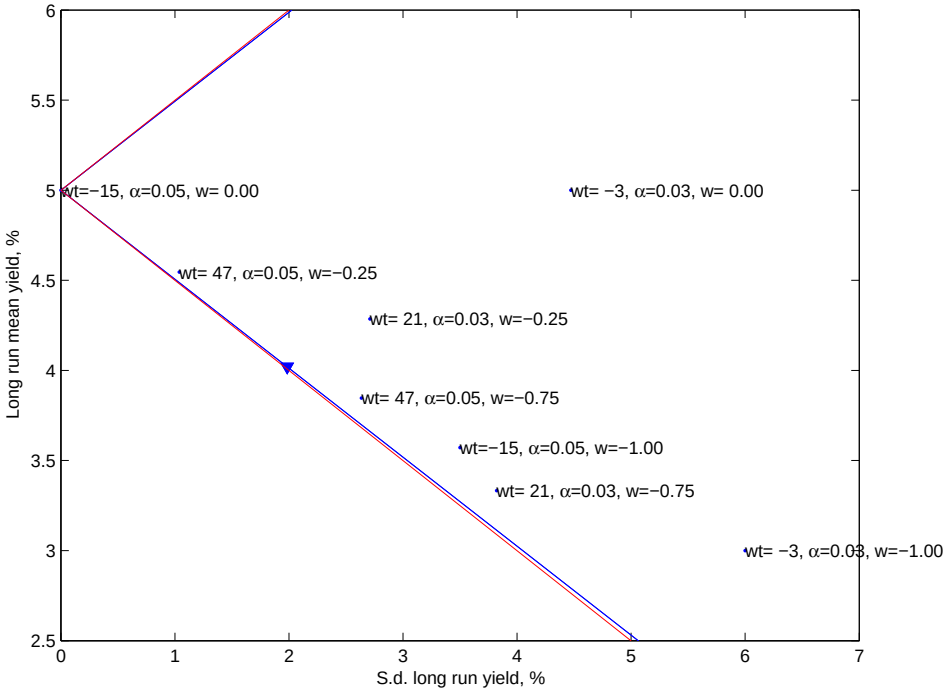

Figure 8: Long-run mean-variance frontier computed from a few fixed payout portfolios. The blue dots give the long run mean and standard deviation of the basis portfolios. The basis portfolios have payout rates of 3% and 5% and portfolio weights -1, -0.75, -0.25, 0. The blue triangle gives the mean-variance efficient portfolio formed from these basis assets by $y^* = 1'\mathcal{E}(yy')^{-1}y/\left[1'\mathcal{E}(yy')^{-1}1\right]$. The numbers give the weight of each basis asset in the mean-variance efficient portfolio. The red line gives the exact mean-variance efficient frontier generated by a 4% payout rate and -0.5 portfolio weight. Paramters $\rho = r^f = 0.05$, $\mu = 0.02$, $\sigma = 0.20$.

It's a little uncomfortable to focus so much on short positions in "maximally inefficient" portfolios on the lower half of the mean-variance frontier. That issue needs to be understood.

A first lesson of Figure 7 is how poorly portfolios do with constant positive weights and payout rates – all the portfolios in the upper half of the mean-variance frontier. Portfolios with $w = 1$, the standard "market portfolio", barely fit on the graph. The $w = 1$ portfolio that does fit has a large payout rate $\alpha = 0.09$. The most natural portfolio with $w = 1$, $\alpha = r^f = 0.05$ isn't

even in the payoff space; its conditional variance grows faster than $e^{-\rho t}$ its long-term variance is infinite. The reason is that geometric Brownian motion can explode in the positive direction, but the long-run variance criterion prizes stability in the level, not the log.

To illustrate, Figure 9 plots a simulation of the yield from three portfolios with $\alpha = 0.04$, one with risky-asset weight $w = +0.5$, one with $w = 0$ and one (the mean-variance efficient payoff y^*) with $w = -0.5$. You can see the volatility of the $w = +0.5$ payoff; it is just what a geometric Brownian should look like. $w = 0$ gives steady 1% growth, since in this case $r^f = 0.05$. The payoff that is short the market $w = -0.5$ trundles off downward, as one would expect of a short position in a positive expected return security. But trundling off *downward* has much less arithmetic variance than exploding *upward*. Thus, the last portfolio graphed in the Figure, formed by a short position in this short portfolio, $r^f + 1.67(r^f - y^*)$, which is on the upper portion of the mean-variance frontier, looks quite good. By construction (the 1.67) it matches the long-run mean of the $w = +0.5$ portfolio, but you can see the dramatically lower long-run variance since it varies much less across time as well as across states of nature.

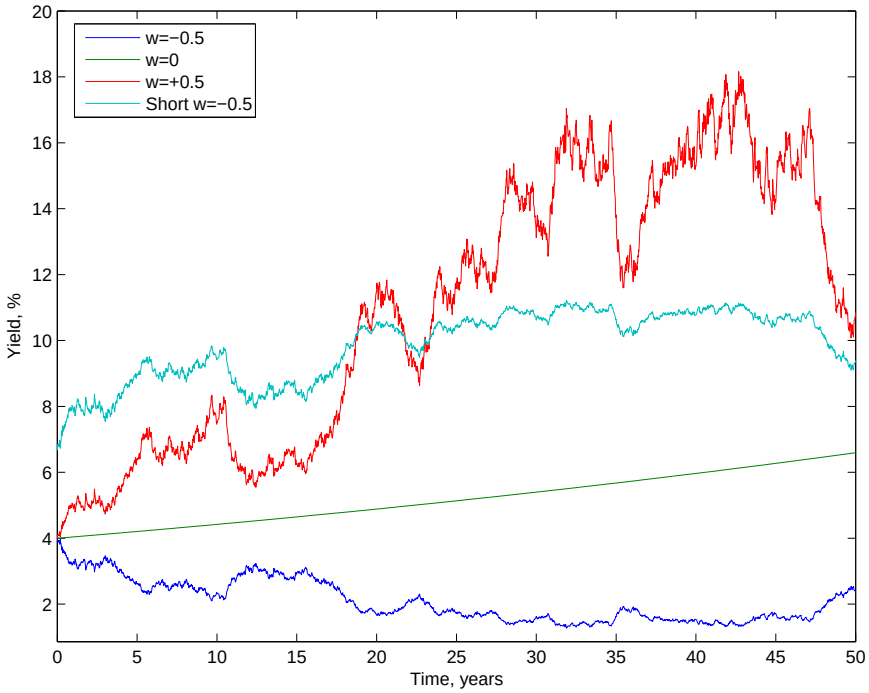

Figure 9: Draw of three portfolio-payout rules. I simulate 50 years of a random normal, then cumulate $dV = (r^f - \alpha)V dt + w'(\mu dt + \sigma dz)$. The payout is $y = \alpha V$. The first three portfolios have a constant $\alpha = 0.04$ payout and weights $w = -0.5, 0$, and 0.5 in the risky assets. The final portfolio is $r^f + 5/3(r^f - y^*)$ where y^* is the payoff with $w = -0.5$. Parameters $\rho = 0.05$, $r^f = 0.05$ $\mu = 0.02$, $\sigma = 0.20$.

Rather than understand the efficient portfolio $r^f + 1.67(r^f - y^*)$ as a short position in a portfolio that it short the market with constant weight $w = -0.5$ and payout $\alpha = 0.04$, it may make more intuitive sense to understand it as a portfolio with a time-varying payout and portfolio weight. We can instead think of it as the payout of a portfolio that starts with \$1

investment, then pays a dollar amount[7]

$$y_t = 0.0267 + 0.04W_t = 0.05W_t + 0.01 \times (2.67 - W_t)$$

and maintains a *dollar* investment in the risky excess return

$$W_t w_t = 1.33 - 0.5 \times W_t = 0.5 \times (2.67 - W_t)\,.$$

Compared to the $\alpha = 0.04$, $w = 0.5$ payoff, the portfolio is initially more aggressive; at $W_0 = 1$ paying $y = 0.0667$ rather than $y = 0.4$ and investing \$0.835 in risky assets rather than \$0.5. As wealth increases, however, the efficient portfolio becomes more cautious; at $W = 2.67$ paying out and investing only in the risk free rate.

Intuitively, if long-run variance is a concern, it makes sense for the portfolio to shift more towards riskfree assets as wealth increases.

As it is interesting to characterize the payoff *on* the frontier in this way, it can also be more intuitive to construct basis payoffs that in the upper portion of the frontier, rather than the slightly weird basis payoffs in the lower portion shown in Figure 8. At the same time, I shift to excess (price-zero) yields which are the cleanest way to express the mean-variance frontier.

The upper portion of the frontier here is generated by portfolios with payout rules and weights that include constants, which can be generated by short positions in constant-weight portfolios plus long positions in the riskless asset If one wants to see portfolios on the upper portion of the mean-variance frontier expressed as positive weights of underlying assets, then the underlying assets must also be generalized to have constants in the portfolio and payout rules, i.e.

$$dV_i = \left(r^f V_i - (a + \alpha_i V_i)\right) dt + (\omega_i + V w_i)' \, dr_t^e.$$

Mechanically, we can just take the yields calculated so far and form $y^e = k\left(r^f - y\right)$. The point however is to interpret these payoffs as the basis assets. Starting with value processes of the form (69)

$$dV^i = (r^f - \alpha)V^i dt + V^i w' dr^e; V_0^i = 1$$

the excess yields come from a value process

$$V = k(1 - V^i); V_0 = 0$$

[7]Specializing to the parameters in this case,

$$\begin{aligned} y_t^{mv} &= r^f W_t + (1 + k - W_t)\left(\mu' \Sigma^{-1} \mu\right) \\ &= 0.05W_t + (2.67 - W_t)\left(\frac{0.02^2}{0.20^2}\right) \\ &= 0.05W_t + 0.01 \times (2.67 - W_t) \\ &= 0.0267 + 0.04W_t \end{aligned}$$

$$\begin{aligned} w_t' &= (1 + k - W_t)\,\mu' \Sigma^{-1} \\ w_t' &= (2.67 - W_t)\,\frac{0.02}{0.20^2} \\ &= (2.67 - W_t)\,0.5 \\ w_t &= 1.33 - 0.5 \times W_t. \end{aligned}$$

The value process follows

$$\begin{aligned} dV &= -kdV^i = -kV^i(r^f - \alpha)dt - kV^i w' dr^e \\ dV &= (V-k)(r^f - \alpha)dt + (V-k)w' dr^e \\ dV &= \left\{ r^f V - \left[kr^f - \alpha\left(k - V\right) \right] \right\} dt - (k-V)w' dr^e \end{aligned}$$

i.e. a payout rule

$$y^e = kr^f - \alpha\left(k - V\right)$$

and a portfolio allocation rule

$$-\left(k - V\right)w.$$

Since[8] $k > V$, previously short weights are now long positions.

Figure 10 presents the results of this calculation, and should look more comfortingly familiar. Again, I delete the choices $\alpha = 0.04$ and $w = 0.5$ so we can see how well approximating the payoff space with a smaller number of assets works.

The basis payoffs represented by blue dots have positive payouts and portfolio weights, though a somewhat different interpretation. They are zero -cost portfolios. They have payouts $y^{ei} = r^f - y^i = \left[r^f - \alpha^i \left(1 - V^i\right) \right]$ and portfolio positions $\omega^i = (1 - V^i)w^i$, where V_i is the value of each underlying portfolio. Value starts at $V_0^i = 0$ and accumulates by $dV^i = \left\{ r^f V^i - y^{ei} \right\} dt - \omega^{i\prime} dr^e$. Thus a "4% payout rate" now means that the payout *changes* by 4% of the value of the portfolio; if the portfolio rises by \$100, you get \$4 *more*. Similarly, a 50% portfolio weight means that if value increases by \$100, \$50 more will be put in the risky assets. But both payout and portfolio rules have constants.

The blue triangle gives the mean-variance efficient payoff formed from these basis assets by $r^f y^{e*} = r^f \mathcal{E}(y^e)' \mathcal{E}(y^e y^{e\prime})^{-1} y^e$. The nearly indistinguishable red square gives the exact mean-variance payoff computed from $r^f y^{e*} = r^f - y^*$ using all assets and $\alpha = 4\%$, $w = 0.5$ in particular, and the lines give the corresponding frontiers. The numbers wt give the weight of each basis asset in the mean-variance efficient portfolio, i.e. the values $r^f \mathcal{E}(y^e y^{e\prime})^{-1} \mathcal{E}(y^e)$. These do not sum to 100 because as zero cost portfolios they do not have to do so.

---

[8] (We can also write

$$\begin{aligned} d(V-k) &= \left(r^f - \alpha\right)(V-k)dt + (V-k)w' dr^e \\ &= \left(r^f - \alpha + w'\mu\right)(V-k)dt + (V-k)w'\sigma dz \end{aligned}$$

$V - k$ follows a Geometric Brownian motion starting at a negative value, so it remains negative. The resulting excess (price-zero) yields are

$$y = kr^f + \alpha\left(V - k\right)$$

Thus they obey

$$dy = \alpha dV$$

$$\begin{aligned} dy &= \alpha\left\{ r^f V - \left[kr^f + \alpha\left(V - k\right) \right] \right\} dt + \alpha(V-k)w' dr^e \\ &= \alpha\left(r^f - \alpha\right)\left(V - k\right)dt + \alpha(V-k)w' dr^e \\ dy &= \left(r^f - \alpha\right)\left(y - kr^f\right)dt + \left(y - kr^f\right)w' dr^e; y_0 = k(r^f - \alpha) \end{aligned}$$

i.e, $\left(y - kr^f\right)$ also follows a Geometric Brownian motion.

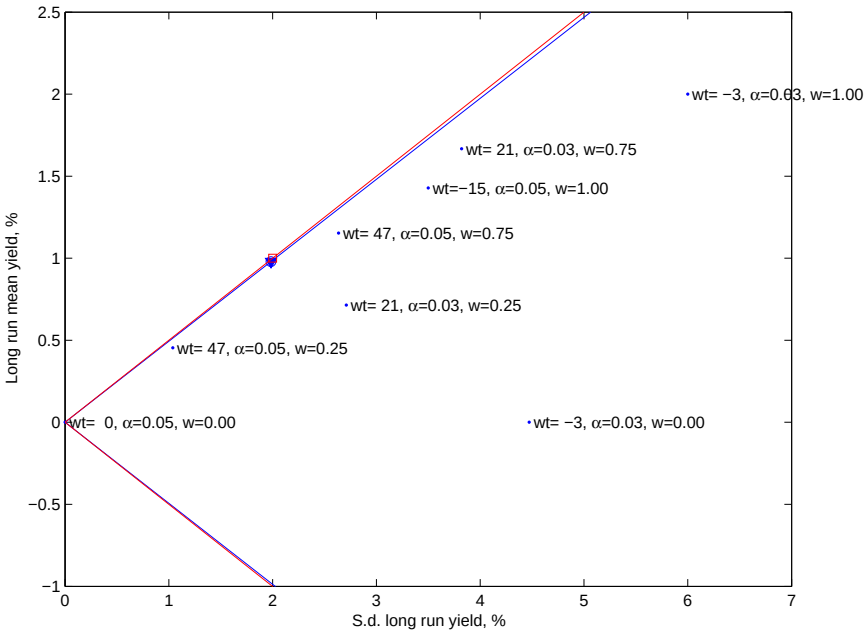
wt= −3, α=0.03, w=1.00
wt= 21, α=0.03, w=0.75
wt=−15, α=0.05, w=1.00
wt= 47, α=0.05, w=0.75
wt= 21, α=0.03, w=0.25
wt= 47, α=0.05, w=0.25
wt= 0, α=0.05, w=0.00
wt= −3, α=0.03, w=0.00
Long run mean yield, %
S.d. long run yield, %
2.5
2
1.5
1
0.5
0
−0.5
−1
0
1
2
3
4
5
6
7

Figure 10:

4 Size and Value

Now, for a real problem. How much should a mean-variance investor tilt towards small stocks, value stocks, and momentum strategies? Perhaps dynamics in the portfolios give a different answer to this question for long horizon investors than for one-period investors. To this end, I examine the value-weighted market, and the Fama-French smb, hml, and umd factor portfolios. Larger collections such as the 25 size and book/market portfolios are well described by the factor portfolios, so there is little point in expanding the set of underlying portfolios.

Table 1 presents statistics on annual returns, and Figure 11 presents the usual one-year horizon mean-variance frontier for excess returns. As usual, the size, value and momentum strategies appear to improve markedly on simply holding the market portfolio. Our objective is to do the analogous simple mean-variance exercise for a *long-run* investor. In particular, *dynamics* in the returns have the potential to change the patterns. For example, a negatively autocorrelated return can have much lower long-term variance than its short term variance.

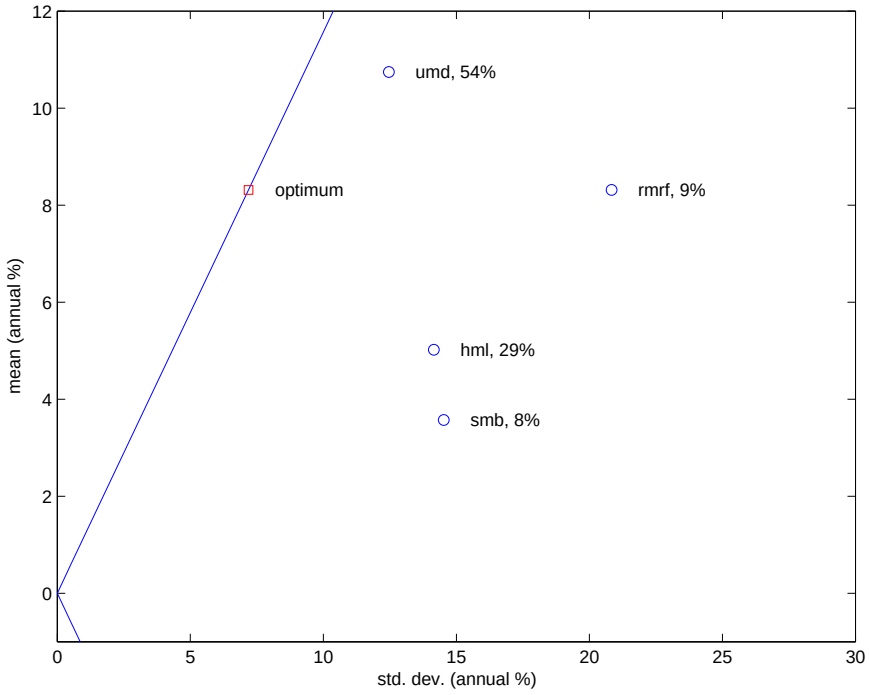

Figure 11: Mean-variance frontier for annual excess returns on market (rmrf) small (smb) value (hml) and momentum (umd) portfolos. "Optimum" is the mean-variance efficient portfolio at the mean market return. Percentages give each portfolio's weight in the optimum portfolio.

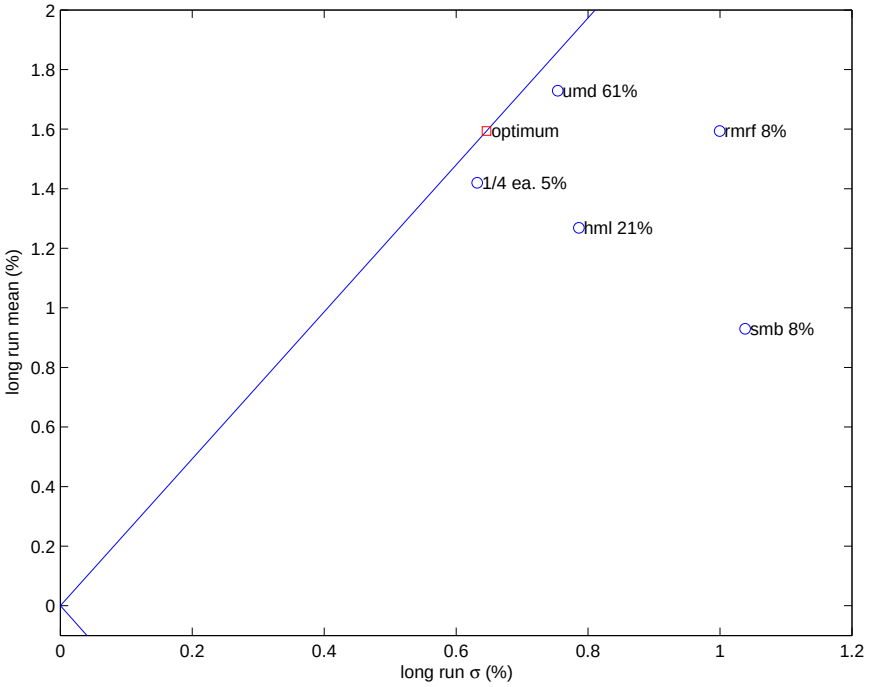

Figure 12: Long-run mean-variance frontier of excess yields. The percentages give the weight of each portfolio in the optimum.

	rmrf	smb	hml	umd	$\frac{1}{4}$ea.	opt.
	One-year returns					
mean (%)	8.3	3.6	5.0	10.7		8.3
std. dev. (%)	20.8	14.5	14.1	12.5		7.2
Sharpe	0.40	0.25	0.36	0.86		1.16
α (%)	0	1.1	4.6	10.9		7.3
β	1	0.30	0.05	-0.02		0.12
$\alpha/\sigma(\varepsilon)$		0.08	0.33	0.88		1.09
weight in optimal (%)	9	8.5	28.7	54.2		
	Long-run yields					
mean $\mathcal{E}(y^e)$ (%)	1.59	0.93	1.27	1.73	1.42	1.59
std. dev. $\tilde{\sigma}(y^e)$ (%)	1.00	1.04	0.79	0.75	0.63	0.64
Sharpe	1.59	0.90	1.61	2.29	2.25	2.47
long-run α (%)	0	0.56	0.79	1.13	0.60	0.92
long-run β	1	0.23	0.30	0.38	0.51	0.42
weight in optimal (%)	8.3	8.4	20.7	60.9	4.8	

Table 1. Mean-variance statistics for excess returns, 1927-2004. "Optimal" is the mean-variance efficient portfolio with the same mean as the value weighted market return rmrf

The first choice in applying the long-run mean / long-run variance idea is what streams of

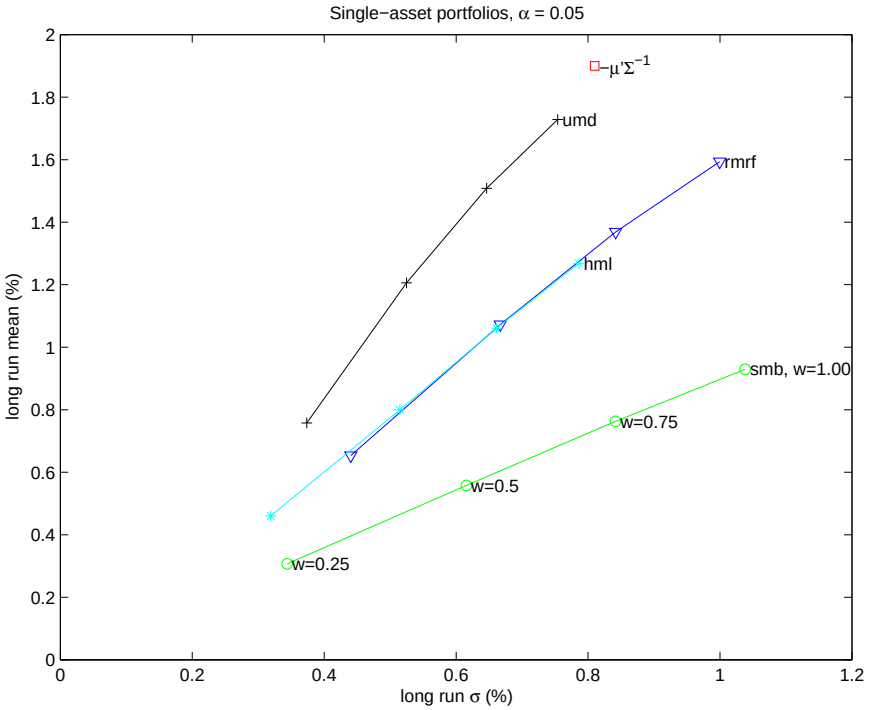

Figure 13: Long run means and standard deviations of constant-weight, constant-payout portfolios. The payout rate is $\alpha = 0.05$ in all cases.

payoffs $\{x_t\}$ to examine. It's tempting simply to define $\{x_t\}$ as the dividend streams corresponding to the chosen assets such as the size and book/market portfolios. Using dividends directly is not generally such a good choice for portfolio theory, however. The investor can dynamically reallocate his investment across assets, and he can consume out of capital gains as well as dividends. The optimal portfolio will therefore typically feature a consumption stream that has almost nothing to do with the underlying dividend streams. For example, in the i.i.d. return case, we know the solution analytically and in this case the investor consumes a proportion of wealth and rebalances continuously, ignoring any distinction between dividends and capital gains.

Given my approach of maximizing over a few representative portfolios rather than attempting the full (infinite-dimensional) exact dynamic optimization over assets and dynamic strategies, therefore, I start by constructing payoffs x_t from returns alone, ignoring the dividend/capital gain split, via sensible but ad-hoc payout rules.

I start with constant payout-portfolio rules, cumulating an initial one dollar investment at the three-month treasury bill rate. For given payout α and weight w in [rmrf, smb, hml, umd], I construct

$$\begin{aligned} \frac{dV}{V} &= \left(r_t^f - \alpha\right) dt - w' dr^e; \; V_0 = 1 \\ y_t &= \alpha V_t \end{aligned}$$

I construct one such yield with a zero weight on all assets,

$$\frac{dV^0}{V^0} = \left(r_t^f - \alpha\right) dt;\ y_t^0 = \alpha V_t^0$$

Then I construct excess yields by

$$z_t = y_t^0 - y_t$$

Again, as above, we can view these as short positions in portfolios that are short the risky assets, or as portfolios that are long the risky assets with a constant in the dollar payout rule; the iid model suggests that they will perform better in a long-run mean-variance sense than standard constant weight portfolios.

The yields so constructed respond sensibly to the fortunes of the various investments. Figure 14 presents the yield and excess yield so constructed starting with a one dollar investment in 1926. The lines represent the yields from using a 5% payout rate and a weight in rmrf of 0, 0.25, 0.5, 0.75 and 1.00 respectively. The smooth line is a 5% payout with a constant investment in treasury bills. The value slowly declines because interest rates were typically below 5% through this period. The other yields cumulate short positions and pay out proportionally to value, so you see them suffer in the runup of the late 1920s, recover a bit in the crash of 1929, and then decline swiftly as the market recovers. The *excess* yield is a short position in these yields, and a long position in the smooth yield from the top graph. Thus, they do well and poorly as we would expect through the market rise in the late 20s, the crash, and then recovery, and so forth.

It's tempting to worry a lot about tail behavior of these portfolios, but remember that the tails all get multiplied by β^t, so matter little to long-run means and variances.

To estimate long run means and variances, I recreate these value and yield processes at each date. Thus denote $y_{t,t+j}$ the yield at time $t+j$ due to a one dollar investment at time t. To estimate long-run means and variances, I start by estimating $E_0(y_j)$ by averaging all the j-step ahead yields, $E_0(y_j) = \frac{1}{T-j}\sum_{t=1}^{T-j}\left(y_{t,t+j}\right)$ and similarly for second moments $E(y_j^i y_j^k)$. I test that $\beta^j E_0(y_j)$ and $\beta^j E_0(y_j^2)$ are converging to zero and report infinity for those cases that are not converging. Then I calculate $\mathcal{E}(y) = \sum \beta^j E_0(y_j)$ and similarly for the long-run second moments. I calculate $\mathcal{E}(y) = \frac{1-\beta}{\beta(1-\beta^T)}\mathcal{E}(y)$. Standard deviations are $\tilde{\sigma}(y) = \left[\mathcal{E}(y^2) - \mathcal{E}(y)^2\right]^{0.5}$. Denoting by $\mathbf{y}^e$ the vector of all excess yields in consideration, the mean-variance efficient excess yield is y^{e*} given by[9] $y^{e*} = \mathcal{E}(\mathbf{y}^e)'\mathcal{E}(\mathbf{y}^e\mathbf{y}^{e\prime})^{-1}\mathbf{y}^e$.

Figure 13 presents the long-run means and variances for a variety of such portfolios. I standardize on a payout rate $\alpha = 0.05$ corresponding to the discount factor $\beta = 0.95$ throughout. Figure 13 then presents long-run means and variances of portfolios that put weight 0, 0.25, 0.5, 0.75, and 1.0 in each of rmrf, smb, hml, and umd respectively with zero weight in the other assets. I also present a portfolio with weight $\mu'\Sigma^{-1}$, where μ and Σ are the mean and covariance matrix of one period returns. This portfolio is long-run mean-variance efficient in the iid world, so it's a good one to try here as well.

Figure 13 says that the choice of *scale* does not matter much. As we increase the weight in a given risky asset portfolio, long-run means and standard deviations move out along a ray of

[9]With large payoff spaces of potentially very highly correlated assets, second moment matrices can be close to singular. This is particularly a problem when one generates assets as I do here by varying parameters on a fairly fine grid – the yield generated from a 5% payout is likely to be very correlated with the yield genrated from a 4.9% payout! I use a pseudo-inverse to address this problem, raising the tolerance level until the results are no longer sensitive to that choice.

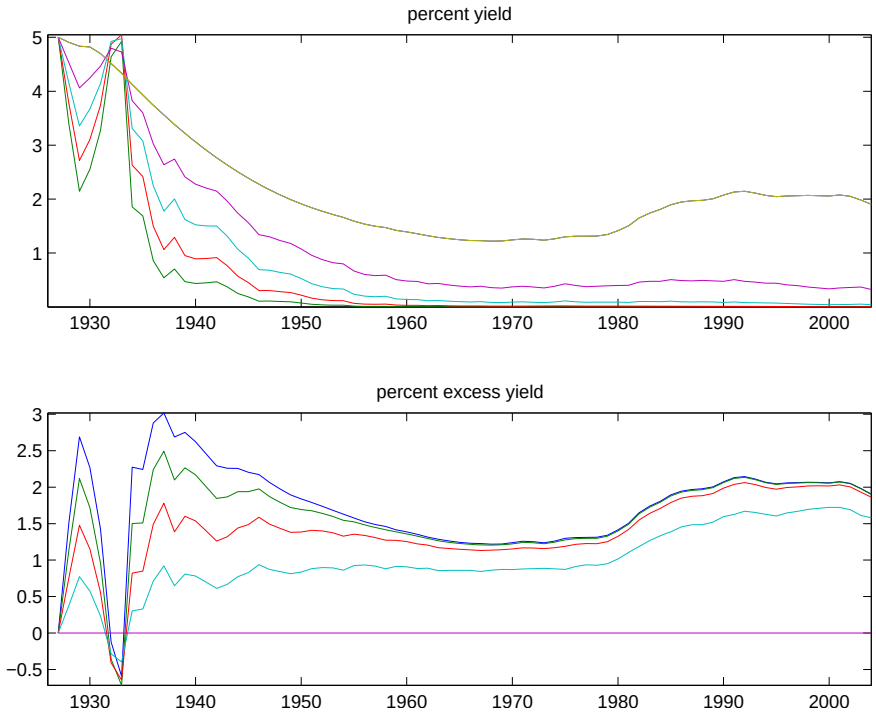

Figure 14: Yield (top) and excess yield (bottom) resulting from one dollar investment in 1926. All portfolios use a 5% payout rate, and weights 0, 0.25, 0.5, 0.75 and 1.0 in the market excess return rmrf.

constant Sharpe ratio. This fact means we can standardize on one weight, pruning considerably the number of basis portfolios that need to be included. Interestingly the weight $\mu'\Sigma^{-1}$ does not do much better than the umd portfolio alone.

Since scale does not seem to matter, Figure 12 and Table 1 consider portfolios with weight one in each of rmrf, smb, hml, and umd, along with a portfolio that places even weight $w' = [0.25\ 0.25\ 0.25\ 0.25]$ across all four assets. Again, a portfolio of a yield that is 100% invested in rmrf and one that is 100% invested in hml is not the same thing as a portfolio that is constantly rebalanced to 50%-50%, so it's interesting to include such portfolios.

Comparing the long-run mean-variance frontier in Figure 12 with its short-run counterpart in Figure 11, and comparing the corresponding numbers in Table 1, we see a lot of commonality but also some interesting differences between short and long-run mean-variance frontiers. The relative positions are about the same; hml gives about the same Sharpe ratio as the market at a lower mean and variance, smb decidedly less, and umd substantially more.

The main difference is that the improvement in Sharpe ratio is considerably less in the long-run case. The one-year mean-variance optimum gives the same mean return with only one third the standard deviation (7%, not 20%) resulting in a three times higher 1.16 Sharpe ratio rather than "only" the 0.4 Sharpe ratio that launched the equity premium puzzle. By contrast, the long-run optimal portfolio has a yield standard deviation of 0.64% rather than the 1% of the market portfolio, so the Sharpe ratio only increases by 50% from 1.59 to 2.47. That's still a big

increase in Sharpe ratio, but nothing like a factor of three!

Underlying the relatively smaller benefits of moving away from the market portfolio are larger betas. The long-run betas of hml and umd (calculated from long-run covariance divided by long-run variance) are 0.3 and 0.4 rather than essentially zero. This greater correlation results in less diversification benefit.

5 Hedging labor income

To investigate the possibilities for hedging labor or outside business income, I try two easily available data sources. Figure 15 presents real employee compensation per employee, and Figure 16 presents real proprietor's income. Both series are far from perfect of course. Compensation of all employees ignores life-cycle variation in the compensation of a given employee, and more importantly it ignores the risks to an individual employee of becoming unemployed. Also, one can see the suspicious change in the behavior of the time series in the late 1960s. Proprietor's income is an aggregate, and hence even less representative of the income to an individual proprietor. Nonetheless, these series show important cyclical variation and they are an easy place to start in order to see if there is any hope for hedging labor income.

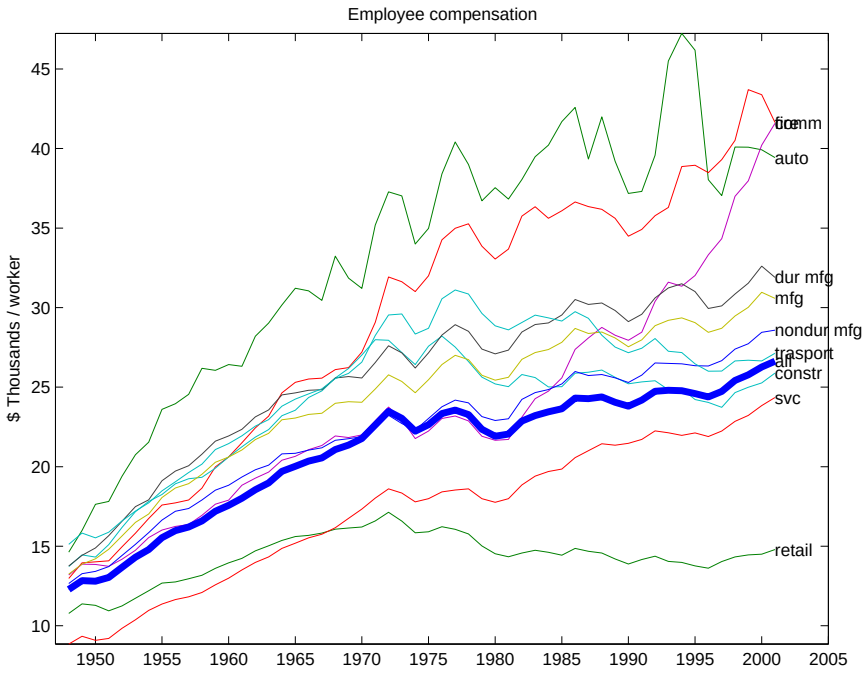

Figure 15: Employee compensation. This is "compensation of employees" divided by number of full-time equivalent emplyees and deflated by the consumer price index.

How well can we replicate these time series using asset yields? To what extent do the hedge portfolios load on priced assets, leading to a distortion of the asset allocation decision?

I first construct the labor income hedge portfolio. I start with the familiar treasury bill

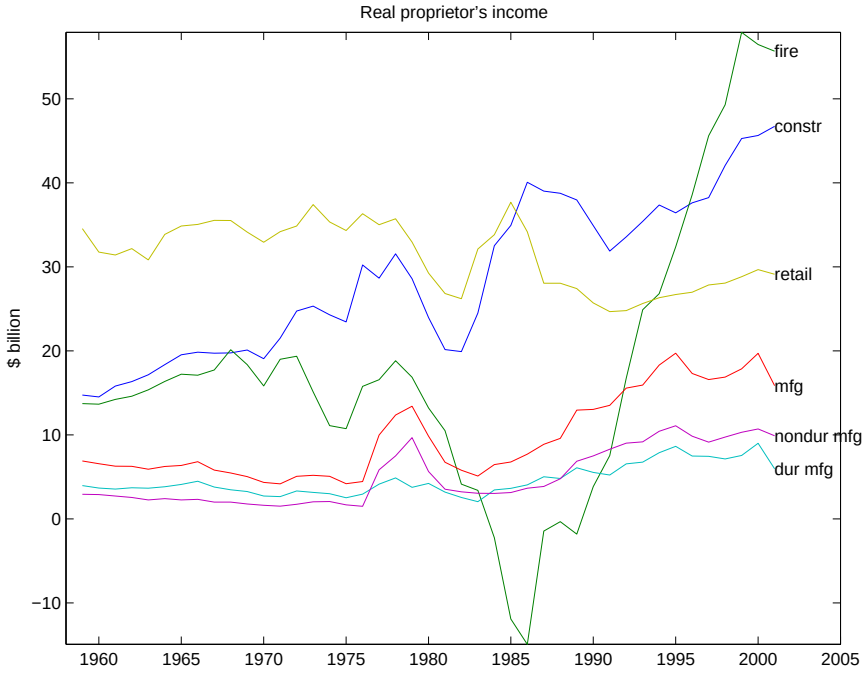

Figure 16: Real proprietor's income.

(rf), market (rmrf), size (smb), book/market (hml) and momentum (umd) portfolios. I form 5 payoffs with a 4% payout rate, one simply invested in treasury bills and the others invested -50% in each risky asset in turn. These portfolios do a good job of capturing the spectrum of *priced* assets, though not necessarily a good job of capturing the spectrum of *unpriced* assets, and of course the latter are just as important as the former if not more so for hedging. Later, I include industry portfolios into the analysis to capture the latter dimension.

Using all employee compensation (the "all" line of Figure 15), I find the long-run regression coefficient $b = \mathcal{E}(yy')^{-1}\mathcal{E}(ey)$ where e denotes employee compensation and y denotes the vector of yields on the basis assets (rf,rmrf,smb,hml,umd). I estimate $\mathcal{E}(ey)$ in the same way I estimate $\mathcal{E}(yy')$ and other long-run moments above. I estimate $E_0(e_j y_j)$ by simulating forward value processes starting with a one dollar investment at each date, producing a yield series $y_{t,t+j}$. Then I average $y_{t,t+j}e_{t+j}/e_t$ across initial dates t, $E(e_j y_j) = \frac{1}{T-t}\sum_{t=0}^{T-j} y_{t,t+j}e_{t+j}/e_t$ and finally I sum to form $\mathcal{E}(ey) = \frac{1-\beta}{\beta(1-\beta^T)}\sum_{j=1}^{T}\beta^j E(e_j y_j)$

Figure 17 presents the first observation of the best hedge payoff $b \times y_{1,t}$ along with the original employee compensation series. The long run R^2 of this regression – $\mathcal{E}\left[(b'y)^2\right]/\mathcal{E}(e^2)$ – is 89%. While this seems remarkably good, of course a large part of the good fit is the match to the change in trend from before 1970 to afterwards, which is perhaps not a particularly interesting feature of the data.

The fitted value of this regression by is a payoff closest to the income process e. Dividing $\hat{b} = b/\sum b_i$ we obtain the yield of the hedge payoff. We can characterize this yield by its long-run mean-variance decomposition,

$$\hat{e} = \hat{b}'y = y^* + w_e y^{e*} + \varepsilon$$

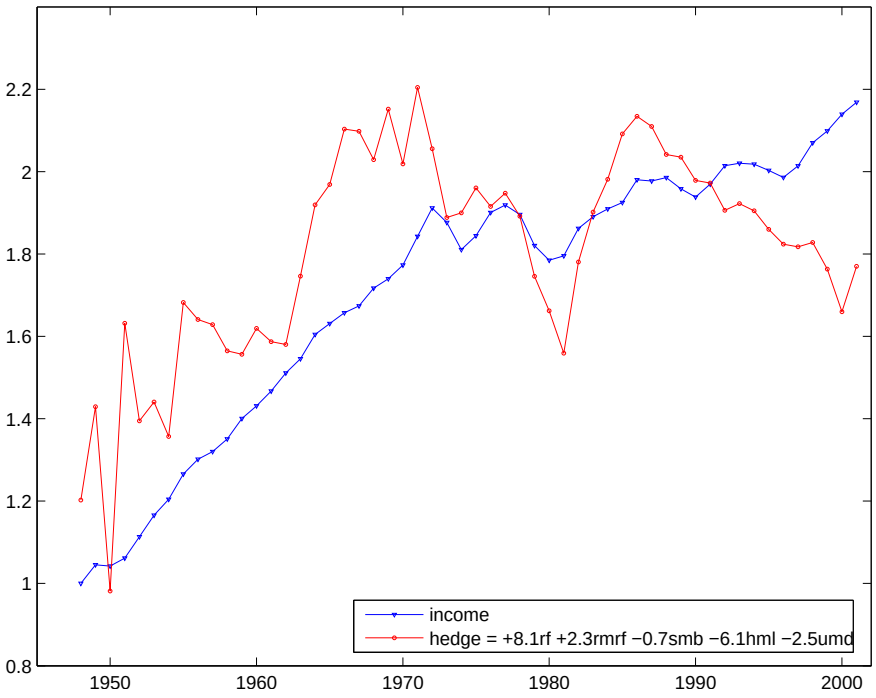

Figure 17: Real employee compensation and payoff of hedge portfolio. The hedge portfolio is a long-run regression of the employee compensation series on yields corresponding to investments in the treasury bill, market return rmrf, small smb, value hml and momentum umd portfolios.

$1 \times y^*$ establishes scale; this is a yield or price-1 payoff. $w_e y^{e*}$ is the systematic or priced component. The residual ε is a zero price, zero expected yield payoff, the truly idiosyncratic part of the labor income payoff.

Figures 18 and 19 present the mean-variance frontier for yields and for excess yields respectively, including the income hedge portfolio. Excess yields are the yield for the riskfree investment less the risky asset yields.

It's initially surprising that the minimum variance payoff in Figure 18 is so large, and that the variance of the payoff marked "'rf" is so large. rf is the real payoff of a portfolio that invests in treasury bills, and pays out a constant 4% of its value. Both the nominal risk free rate and inflation change a lot over time, and the payout rate is not engineered to match the long-run growth rate of the portfolio, so it has both a trend and substantial variation over time. Synthesizing an indexed perpetuity is not so easy! Clearly, though, one can do a better job of this synthesis and that will extend the payoff space to the left. At a minimum, 0 payoff is always available!

The hedge portfolio in both figures has a huge systematic component. It costs a lot to form the payoff graphed in Figure 17; That payoff loads strongly on assets such as hml and umd that pay very high average returns. If this is reality, investors have good reasons to avoid asset market investments with similarly high premia.

The hedge portfolio has a much smaller idiosyncratic component, so there is little an investor

with this outside income series can do to hedge without paying a premium or affecting his or her optimal asset portfolio. However, the set of basis payoffs here focuses on *priced* payoffs; adding industry portfolios and other dimensions of *unpriced* risk could substantially increase the idiosyncratic component of the hedge portfolio.

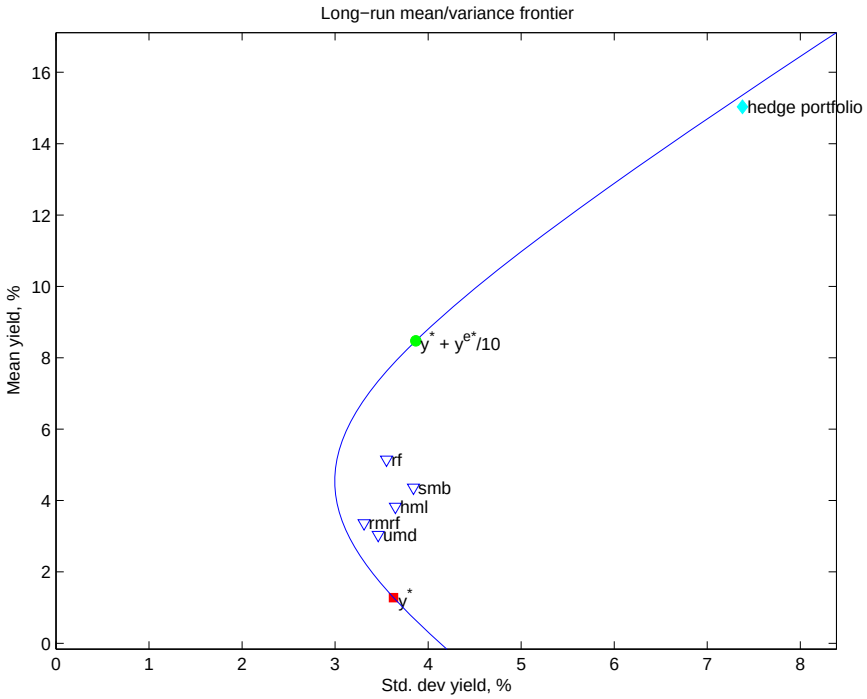

Figure 18: Long run mean-variance frontier. The underlying yields are formed from investments in the riskfree rate (rf) the market excess return (rmrf) and small (smb) value (hml) and momentum (umd) portfolios. Each payoff is formed from a 4% payout rate and a 50% short postionin the irisky assets. Payoffs are real.

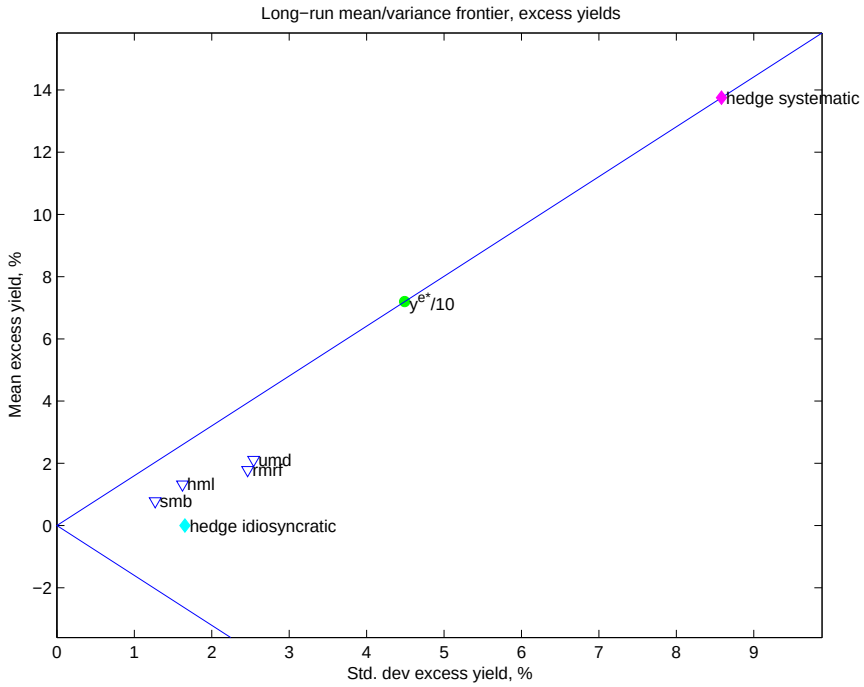

Figure 19: Long run mean-variance frontier for excess yields. The underlying yields are formed from investments in the riskfree rate (rf) the market excess return (rmrf) and small (smb) value (hml) and momentum (umd) portfolios. Each payoff is formed from a 4% payout rate and a 50% short postionin the irisky assets. Payoffs are real.

6 References

Campbell John Y., and Luis Viceira, 2005, The Term Structure of the Risk-Return Tradeoff , *Financial Analysts Journal* 61

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