# Long-Run Mean-Variance Analysis in a Diffusion Environment

John H. Cochrane

December 27, 2012

## 1 Introduction

This note explores long-run mean-variance analysis as described in Cochrane (2012a) "A Mean-Variance Benchmark for Intertemporal Portfolio Theory," in the standard environment of continuous time, i.i.d. lognormal returns and a constant interest rate. In this case markets are complete, so the discount factor method $u'(c) = \lambda m_t$ can easily solve portfolio problems, and the central points of long-run mean-variance analysis – the ability to handle incomplete markets and return dynamics – will not show. Still, we must understand this environment before proceeding to more complex and interesting environments, and that is the purpose of this note.

I compare the quadratic and power utility solutions. The quadratic solution is qualitatively similar to the power utility solution. However, at standard values for the equity premium, stock market volatility, and riskfree rates, the quadratic is not a useful "approximation" for the power utility answer, or vice versa. 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.

I characterize the long-run mean-variance frontier, and show how it captures volatility over time, captured by the difference between riskfree rate and discount rate, and volatility over states of nature, captured by the maximum Sharpe ratio.

In this dynamic trading environment, payoffs are constructed by a portfolio strategy, which describes the composition of the instantaneous portfolio, and a payout strategy, which describes how quickly one takes funds out of the investment portfolio. Both quadratic and power utility investors hold portfolios with mean-variance efficient instantaneous returns, and pay out as an increasing function of wealth. The quadratic utility investor becomes more risk averse as wealth rises, and has a constant in the consumption-wealth function.

The quadratic policy can also be described more elegantly in terms of the long-run mean-variance frontier. He invests some initial wealth in an indexed perpetuity. He invests his remaining wealth in a short position in a strategy that itself shorts a mean-variance efficient portfolio. This strategy maintains the same constant weights over time as in the power utility case. The double-short position generates the long-run mean-variance efficient payoff without further dynamic trading.

Described either way, these examples make clear that the long-run mean-variance efficient payoff is *not* formed from a constantly rebalanced allocation between riskfree rate and mean-variance efficient returns, $r^f + \frac{1}{\gamma}\mu'\Sigma^{-1}dr^e$ as is the case for power utility.

Total wealth, which becomes the market portfolio in equilibrium has a negative lognormal distribution. Wealth never grows above the value $c^b/r^f$ that can support bliss point consumption forever, but it can decline arbitrarily in the negative direction. Though the underlying returns – technologies really – are lognormal, the investor becomes more risk averse and lowers risky investment as wealth rises, so the market return does not inherit this property.

This example is important to separate the character of underlying technological opportunities (a constant riskfree investment and a set of lognormal diffusions here) from the character of observed asset returns, which include the possibility of dynamically rebalancing in and out of the underlying technologies.

This investigation also resolves Dybvig and Ingersoll's (1982) classic puzzle about the CAPM and quadratic utility in a mean-variance environment. They show arbitrage opportunities result *if* the market return can attain sufficiently high values. In a market of quadratic utility investors, the market return never achieves those values.

A long section tracks down why the formulas blow up for parameter values $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$, which are in fact quite reasonable. The answer is that an impatient quadratic utility investor exploits the lognormal opportunities. With a finite lifetime he consumes more today, driving wealth down, but then repaying with a burst of very negative consumption late in life. Similarly, he supports current consumption with a double-or-nothing trading strategy, allowing a burst of huge negative consumption in the very rare event that the risky investments do not recover. In the limit, the burst of negative consumption is forever delayed. This investigation solves the puzzle, but the fact remains then that quadratic utility, long horizons and parameter configurations $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$ are just not a very interesting combination.

This fact should not really surprise us. Lognormal 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. The pathologies of extreme long horizon lognormal returns and mean-variance portfolios are well known, and I survey them.

This is not necessarily fatal for the quadratic model to become more than an interesting conceptual benchmark. I look at the data, and market return data seem to behave more normal than lognormal. The long right tail predicted by the lognormal is missing, while the fat left tail of short horizon returns also disappears. Even at a 10 year horizon, index returns are better described by a normal rather than lognormal distribution. This finding is not that surprising: we know that there is some mean-reversion in returns, and that volatility decreases when the market rises. Both effects cut off the large troublesome right tail of the lognormal. However, it does mean that a quantitatively realistic calculation (one that violates $2r^f - \rho - \mu'\Sigma^{-1}\mu > 0$) must incorporate at least stochastic volatility and potentially mean-reversion, to say nothing of additional state variables, exceeding by far the back of any envelope.

Quadratic utility with a fixed bliss point leads to consumption that does not grow over time, and wealth always below a fixed value. I show in the last section that specifying bliss points that grow over time is a potential resolution to this difficulty. I investigate a geometrically growing bliss point; I show how to construct a stochastic bliss point so that the quadratic and power solutions match, which may be a useful starting point for approximate solutions, and I show how temporal nonesparabilties such as habits or durable goods can induce similar behavior in the bliss point and allow growth.

## 1.1 Setup and $x^*$

There is a constant riskfree rate $r^f dt$. There are $N$ basis assets whose instantaneous excess returns follow

$$dr_t^e = dr_t - r^f dt = \mu dt + \sigma dB_t; \sigma\sigma' = \Sigma,$$

I use the notation $dr$ to denote instantaneous returns, i.e. if dividends are paid at rate $Ddt$, then $dr \equiv dp/p + D/p\ dt$.

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

$$dW_t = (r^f W_t - x_t)dt + \omega_t' dr_t^e$$

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 $\omega$ to rule out doubling strategies. The payoff space $\underline{X}$ consists of payoffs generated in this way by different choices of payout strategy $x_t$ and portfolio strategy $\omega_t$.

The payoff $x^* \in \underline{X}$ that generates prices by $p = k\tilde{E}(x^* x)$ is characterized by

$$\frac{dx_t^*}{x_t^*} = \left(\rho - r^f\right) dt - \mu'\Sigma^{-1}\sigma dB_t \tag{75}$$

and thus

$$x_t^* = e^{\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\sigma \int_0^t dB_t} \tag{76}$$

Prices are given by

$$p(x) = k\tilde{E}(x^* x) = E\int_{t=0}^{\infty} e^{-\rho t} x_t^* x_t dt.$$

The price of $x^*$ itself is

$$p(x^*) = k\tilde{E}(x^{*2}) = E\int_{t=0}^{\infty} e^{-\rho t} e^{2\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - 2\mu'\Sigma^{-1}\sigma\int_{\tau=0}^{t} dz_\tau} dt \tag{77}$$

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

I consider the possibility that $2r^f - \rho - \mu'\Sigma^{-1}\mu \leq 0$ and thus $p(x^*) = \infty$ below. From (75) (or directly from (76)) we have $E(x_t^*) = e^{(\rho - r)t}$ and hence

$$\tilde{E}(x^*) = \rho\int_0^{\infty} e^{-\rho t} e^{\left(\rho - r^f\right)t} = \frac{\rho}{r^f}. \tag{78}$$

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

## 1.2 Consumption-portfolio problems

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

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

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

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

For power utility

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

we have

$$c_t^p = \left(\lambda x_t^*\right)^{-\frac{1}{\gamma}}. \tag{79}$$

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

$$c_t^q = c_t^b - \lambda x_t^* \tag{80}$$

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

Equations (79) and (80) 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[3]

$$\frac{c_t^p}{W_0} = \frac{1}{\gamma}\left[\rho + (\gamma-1)\left(r^f + \frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right)\right] x_t^{*-\frac{1}{\gamma}}. \tag{81}$$

3

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

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

$$\frac{c_t^q}{W_0} = r^f + \frac{1}{\gamma}\left(r^f - y_t^*\right) \tag{82}$$

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

Figure 1 contrasts the power utility case (81) with the quadratic or mean-variance case (82). Both consumption decision rules slope downward. Lower returns mean higher marginal utility $x^*$ or $y^*$ and hence lower consumption, so consumption declines as we go from left to right.

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.

By contrast, the quadratic utility investor consumes a bit more in normal times (near 1), financing that by consuming even negative amounts if the price $x^*$ were to rise past about 3.5 (in equilibrium, it never does), and with consumption never exceeding the bliss point (about 7% of wealth here) even if the price of consumption ($x^*$) falls to zero.

Clearly, the difference between power and quadratic is "small" for "small" returns and thus small 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 yields, where the power utility consumption increases without bound but quadratic utility (with fixed bliss point) consumption stops growing at the bliss point.

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).

---

Substitute this expression for $\lambda$ in (79).

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.

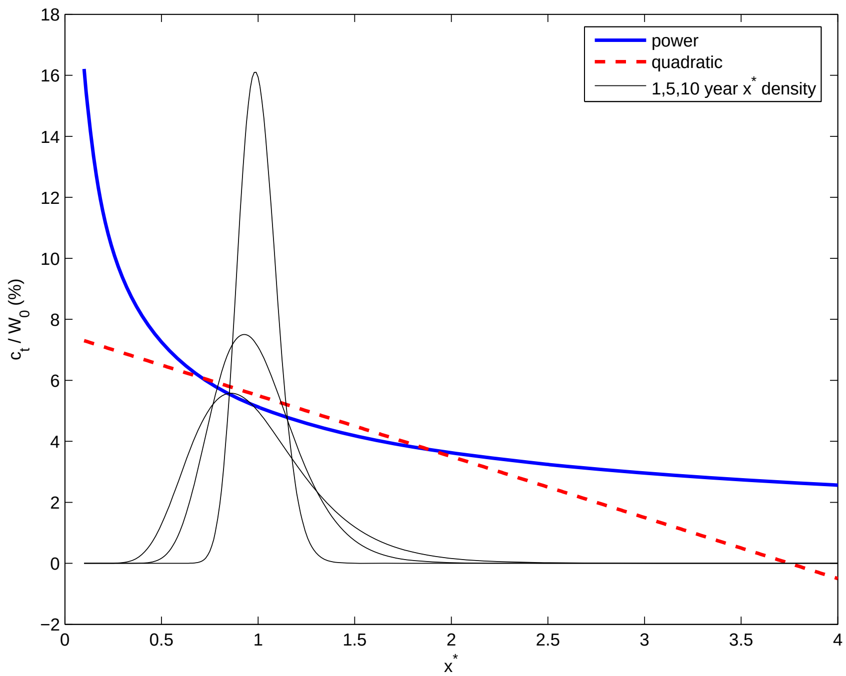


Figure 1: Consumption payoffs from power utility and quadratic utility in the iid return case. The x axis is the pricing payoff or marginal utility process $x^*$. The $y$ axis gives consumption relative to initial wealth as a function of $x^*$. Parameters are risk aversion $\gamma = 2$, discount factor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

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 for a calibration to standard values of equity market volatility.

## 1.3 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 $\tilde{E}(z^*)/\tilde{\sigma}(z^*)$. Evaluating in the i.i.d. lognormal model with a riskfree rate, we have

$$\frac{\tilde{E}(z_t^*)}{\tilde{\sigma}(z_t^*)} = \sqrt{\frac{(r^f-\rho)^2+\rho\mu'\Sigma^{-1}\mu}{r^{f2}-(r^f-\rho)^2-\rho\mu'\Sigma^{-1}\mu}}. \tag{83}$$

*Derivation* We evaluate

$$\tilde{E}(z_t^*) = \tilde{E}\left(1-\frac{y_t^*}{r^f}\right) = 1-\frac{\tilde{E}(x^*)}{r^f p(x^*)} = 1-\frac{\rho}{r^{f2}p(x^*)}$$

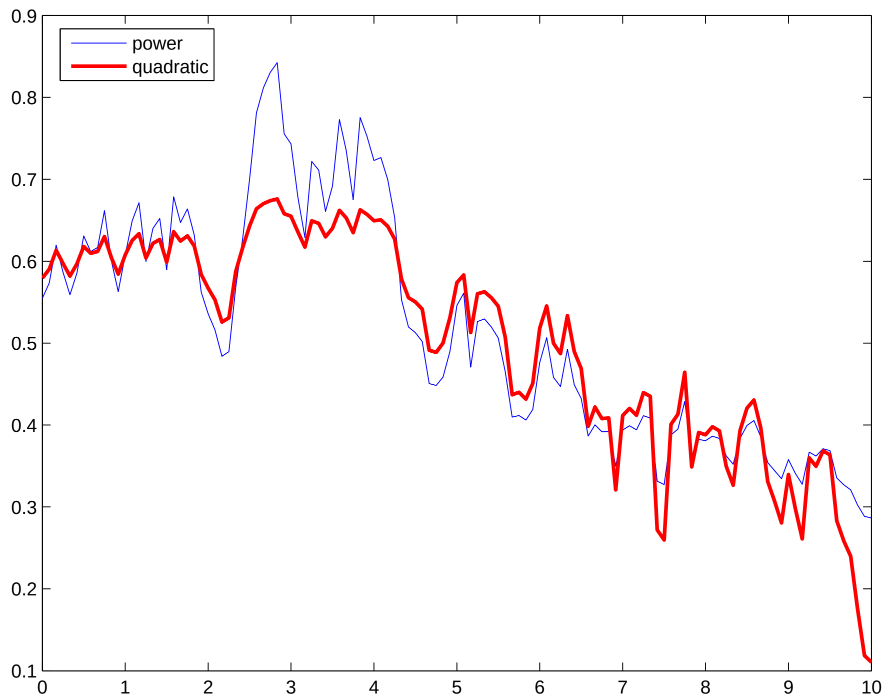


Figure 2: Simulated consmption from power and quadratic utility in iid return case.

$$
\begin{aligned}
&= 1 - \frac{\rho}{r^{f2}}\left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right) \\
&= \frac{(r^f-\rho)^2 + \rho\mu'\Sigma^{-1}\mu}{r^{f2}}.
\end{aligned}
$$

The first line follows from (78). We can find the variance from

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

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

$$
\frac{\tilde{E}(z^*)}{\tilde{\sigma}(z^*)} = \frac{\tilde{E}(z^*)}{\sqrt{\tilde{E}(z^*)\left[1-\tilde{E}(z^*)\right]}} = \sqrt{\frac{\tilde{E}(z^*)}{1-\tilde{E}(z^*)}} \tag{84}
$$

and (83) follows.

The formula (83) 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,

$$
\frac{\tilde{E}(z_t^*)}{\tilde{\sigma}(z_t^*)} = \sqrt{\frac{(r^f-\rho)^2}{r^{f2} - (r^f-\rho)^2}}. \tag{85}
$$

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 payout level and its 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-run variance, and long-run mean $\tilde{E}(c) = r^f W_0$. 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-run variance, since that measure prizes stability over time as well as over states of nature.

If $\rho = r^f$, the investor chooses the constant consumption path. In view of (85), the long-run mean excess return is zero for any level of long-run volatility, so the investor chooses the minimum long-run variance portfolio. If $r^f < \rho$, the consumer chooses a declining consumption path. 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 (85), 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 and we focus exclusively on risk. The long-run Sharpe ratio simplifies to

$$\frac{\tilde{E}(z_t^*)}{\tilde{\sigma}(z_t^*)} = \sqrt{\frac{\mu'\Sigma^{-1}\mu}{r^f - \mu'\Sigma^{-1}\mu}} \tag{86}$$

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.

The possibility $r^f - \mu'\Sigma^{-1}\mu \leq 0$ is not outlandish. With a typical 0.5 market Sharpe ratio, $\mu'\Sigma^{-1}\mu = 0.25$. I return to this issue below.

## 1.4 Portfolio and payout 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 $\omega_t$ in the risky assets and by taking the payoff as a dividend. Wealth then follows

$$dW_t = \left(r_t^f W_t - x_t\right) dt + \omega_t' dr_t^e. \tag{87}$$

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

The power-utility investor holds a constantly-rebalanced instantaneously mean-variance efficient portfolio with constant weights and consumes a constant fraction of wealth,

$$w_t = \frac{1}{\gamma}\Sigma^{-1}\mu, \tag{88}$$

$$c_t^p = \frac{1}{\gamma}\left[\rho + (\gamma - 1)\left(r^f + \frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right)\right] W_t. \tag{89}$$

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

$$w_t = \frac{\left(c^b - r^f W_t\right)}{r^f W_t}\Sigma^{-1}\mu = \frac{1}{\gamma_t}\Sigma^{-1}\mu. \tag{90}$$

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 $\Sigma^{-1}$ as the inverse of a local risk-aversion coefficient, and I have defined $\gamma_t$ in that way in the right-hand equality. 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 rule for the quadratic-utility investor is

$$c_t^q = c^b - \left[\left(2r^f - \mu'\Sigma^{-1}\mu\right) - \rho\right]\left(\frac{c^b}{r^f} - W_t\right) \tag{91}$$

Both power (89) and quadratic (91) payout rules rise proportionally to wealth, but quadratic utility consumption also has an intercept.

### 1.4.1 Constant-weight portfolios; portfolios for the frontier; and shorting the short

There is another representation of the quadratic utility portfolio (90) with constant portfolio weights, rather than the time-varying weight $1/\gamma_t$ described by (90). This expression also shows one way to dynamically construct the mean-variance frontier.

We start by understanding the dynamic portfolio underlying $y^*$. Start with the diffusion representation for $x^*$ (75) and then rewrite it as a value process in the form (87),

$$\frac{dx^*}{x^*} = \left(\rho - r^f\right) dt - \mu'\Sigma^{-1}\sigma dz_t \tag{92}$$

$$
\begin{aligned}
\frac{dx^*}{x^*} &= \left[r^f - \left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)\right] dt - \mu'\Sigma^{-1}\left(\mu dt + \sigma dz_t\right) \\
\frac{dx^*}{x^*} &= \left[r^f - \frac{1}{p(x^*)}\right] dt - \mu'\Sigma^{-1}\left(\mu dt + \sigma dz_t\right) \\
dx_t^* &= \left(r^f x_t^* - y_t^*\right) dt - x_t^* \mu'\Sigma^{-1} dr_t^e. \qquad (93)
\end{aligned}
$$

The last equation shows us that $x^*$ is value process that generates the payoff $y^* = x^*/p(x^*)$ as its dividend stream. $x^*$ wears many hats: It is a discount factor, it is a payoff, and it is a value process that generates its own yield as a payoff!

Now, we understand how to create $y^*$: From (93), the portfolio that generates the payoff $y^*$ 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 $1/p(x^*) = \left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)$ of its value. The point: unlike the quadratic utility portfolio (90), but like the power utility portfolio (88), these weights are *constants.*

The quadratic utility investor consumes a constant linear function of $y^*$

$$c^q = \hat{x} = c^b - \left[c^b/y^f - W_0\right] y^*$$

and his portfolio yield is a constant linear function of $y^*$ :

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

Thus, in place of the time-varying weights in (90), we can say that the quadratic utility investor *shorts* a portfolio that is *short* a constant fraction of wealth in a mean-variance efficient portfolio.

There are two lessons here. First, the portfolio strategy underlying a given payout is not unique or uniquely characterized, even in this simple example. Second, to create a long-run mean-variance efficient yield

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

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. It is formed either with the dynamic strategy described by (90) and (91), or by the constant-weight strategy described here: 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.

## 1.5 Consumption and the wealth (market) portfolio

We solve next for the actual values of consumption and wealth, not just the rules relating them.

With quadratic utility, wealth follows

$$\left(\frac{c^b}{r^f} - W_t\right) = \left(\frac{c^b}{r^f} - W_0\right) e^{-\left(r^f + \frac{1}{2}\mu'\Sigma^{-1}\mu - \rho\right)t - \mu'\Sigma^{-1}\sigma\int_0^t dB_s} \tag{94}$$

The return of the wealth portfolio, which is the market portfolio for an identical-agent model, is

$$dr^m = \left[r^f + \left(\frac{c^b}{r^f W_t} - 1\right)\mu'\Sigma^{-1}\mu\right]dt + \left(\frac{c^b}{r^f W_t} - 1\right)\mu'\Sigma^{-1}\sigma dB. \tag{95}$$

Consumption itself follows

$$c^b - c_t = \left(c^b - c_0\right) e^{-\left(r^f + \frac{1}{2}\mu'\Sigma^{-1}\mu - \rho\right)t - \mu'\Sigma^{-1}\sigma\int_0^t dz_t} \tag{96}$$

where

$$c^b - c_0 = \left(\frac{c^b}{r^f} - W_0\right)\left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right).$$

(I assume here $\left(2r^f - \rho - \mu'\Sigma^{-1}\mu,\right) > 0$ which I discuss below.)

The corresponding expressions for power utility are

$$W_t = W_0 e^{\frac{1}{\gamma}\left[\left(r^f + \frac{1}{2}\mu'\Sigma^{-1}\mu - \rho\right)t + \mu'\Sigma^{-1}\sigma\int_0^t dB_s\right]} \tag{97}$$

$$c_t = c_0 e^{\frac{1}{\gamma}\left[\left(r^f + \frac{1}{2}\mu'\Sigma^{-1}\mu - \rho\right)t + \mu'\Sigma^{-1}\sigma\int_0^t dB_s\right]} \tag{98}$$

$$c_0 = \frac{1}{\gamma}\left[\rho + (\gamma - 1)\left(r^f + \frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right)\right]W_0$$

In the power utility model in the power utility model (97), $W_t$ is lognormally distributed. In the quadratic utility model, equation (94) shows that $\left(\frac{c^b}{r^f} - W_t\right)$, not $W_t$ itself as is lognormally distributed. Wealth never exceeds the value needed to support the bliss point, $c^b/r^f$. It can decline to arbitrarily large *negative* values.

Expression (94) and (95) make an important point. Based on power utility experience, you might think of forming portfolios of riskfree rate and a lognormal diffusion, specifying the latter as the "market" portfolio. But it isn't, and it can't be. The lognormal diffusion is not long-run mean-variance efficient. It doesn't make sense to specify a lognormal diffusion for the market portfolio in a long-run mean-variance exercise.

The way to think of the computations here is that we have stared with two underlying *technologies*, riskfree and risky, the latter following a lognormal diffusion. Investors will then control the characteristics of the average or market portfolio so that the market portfolio does not follow a lognormal diffusion. As wealth grows, the weights in the risky assets decline. In aggregate, as I have not specified any adjustment costs, this means investors cut back physical investment in the risky technology. This action by investors cuts off the troublesome (to quadratic utility) right tail of the lognormal diffusion. Though an uncontrolled lognormal diffusion will produce wealth past the bliss point, quadratic utility investors eat up the wealth before they get to the bliss point. Power utility investors, especially with $\gamma \neq 1$, do a similar aggregate rebalancing, to control $W_t \geq 0$. But since that rebalancing does not change the character of the distribution it's harder to notice.

Dybvig and Ingersoll (1982) criticized CAPM economies in discrete time models with dynamic trading allowing market completion, pointing out that the discount factors would imply negative marginal utility and hence arbitrage opportunities for the non-quadratic investors. (The unique discount factor supporting the CAPM is $m_{t+1} = a - bR^m_{t+1}$, so if $R^m_{t+1}$ is sufficiently large, $m_{t+1}$ must be negative)

That puzzle does not apply here, since *market* returns are never large enough to drive marginal utility negative. Dybvig and Ingersoll specify the *market* return process exogenously, which is how they obtain a puzzle. I specify the *technological opportunities* exogenously, but allow investors to collectively rebalance out of risky assets. The resulting market portfolio does not have the lognormal diffusion character of the underlying technological opportunity.

Dybvig and Ingersoll's main theorem (p. 237) states:

> Theorem 1: Suppose that (i) mean-variance pricing holds for all assets....(ii) markets are complete so that any payoff across states can be purchased as some portfolio of marketed securities; and (iii) the market portfolio generates sufficiently large returns in some state(s), that is, prob $(a - bR^m_{t+1} < 0) > 0$ [my notation]. Then there exists an arbitrage possibility.

The third condition does not obtain, because a market of quadratic utility investors does not allow the market return to become sufficiently high. The value of the market may become negative $W_t < 0$, but marginal utility is proportional to the discount factor $\left(c^b - c_t\right) = \lambda x^*_t$, and marginal utility is always positive.

*Derivation*

Uniting (87), (90), and (91),

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

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 (90)

## 1.6 Potentially infinite price of $x^*$

Formula (77),

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

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. Since $p(x^*) = k\tilde{E}(x^{*2})$, at these parameter values at least the payoff $x^*$ and possibly others violates the square integrability condition $\tilde{E}(x^{*2}) < \infty$. (Hansen and Sargent 2004 p. 211 warn us, "the requirement that consumption processes in C have unambiguously finite prices is a nontrivial restriction that is not satisfied in general")

At these parameter values, the infinite-horizon quadratic utility consumption-investment problem, solved directly or with the $x^*$ machinery, suffers a technical problem. The solutions we have studied suggest $c_t = c^b$, but that solution is infeasible. Yet any solution $c_t < c^b$ can be show to be suboptimal.

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 $\rho$ 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 (83) as

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

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 (91) goes simply to consumption at the bliss point, $c_t = c^b$, for any initial wealth.

I examine this issue in this section, by studying finite horizon problems, and problems with a wealth constraint or a limit on how negative consumption can go. All three of these environments solve the technical problems – $p(x^*)$ is finite, the long-run mean-variance frontier is well defined, and the quadratic utility consumption-portfolio problem is well behaved – for any parameter values. In each case the *limit*, as the horizon grows, the wealth constraint or consumption constraint decline, is $c_t = c^b$, but this *limit point* is invalid.

However, the well behaved finite or constrained problems are not very interesting. In each case, the investor finances consumption very near the bliss point by selling off a huge lump of negative consumption, either very late in life, or in one extreme state of nature. In the limit, this negative lump disappears, which is the technical problem of truly infinite horizons. But close to the limit, this description of optimization is just not very interesting.

## 1.7 Finite horizons and the limit

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.

With a finite horizon, all the discount factors and consumption streams remain square-integrable. The discount factor payoff 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

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

The price $p(x^*)$ is now finite in all cases – finite horizon models are *technically* well behaved. In the troublesome cases $\rho \geq 2r^f - \mu'\Sigma^{-1}\mu \leq 0$, $p(x^*)$ grows with horizon, and, as we will see, the quadratic utility consumption-portfolio problems in this environment are not that interesting. The investor finances consumption near the bliss point by borrowing a huge amount, and then paying it off with a huge negative consumption late in life, or a huge negative consumption in the rare state of nature that borrowing cannot be paid off by stock market investments. To the extent that one finds such portfolio analysis unrealistic, then this is an *uninteresting* set of parameter values or a poor environment to consider.

With the finite horizon cases in hand, we can then take the limit as $T \to \infty$ and understand how the *limit point* potentially differs from the limit.

Constraints on wealth or borrowing, $W_t \geq \bar{W}$, or a constraint on consumption $c_t \geq \bar{c}$ such as $c_t \geq 0$ produce well-behaved problems just as the finite-horizon constraint $W_T = 0$ does. In the presence of wealth or consumption constraints, the impatient 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. Here too, as we relax the constraints the solutions approach $c_t = c^b$, but the limit point is not a valid solution. For practical problems, imposing a wealth or consumption constraint produce slightly more reasonable solutions. However, solving for the lagrange multiplier is slightly more algebraically involved, since you have to find the time when the constraint starts to bind. For that reason, I only present the finite-horizon cases.

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.

### 1.7.1 Allocation over time, $\mu = 0$.

Consider the finite-horizon problem

$$\max\left(-\frac{1}{2}\right)\int_0^T e^{-\rho t}(c^b - c_t)^2 dt \text{ s.t.} dW_t = \left(r^f W_t - c_t\right)dt;\; W_0; W_T = 0.$$

The optimal consumption profile in this case is

$$c^b - c_t = \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W\right]\frac{\left(2r^f - \rho\right)}{1 - e^{(\rho - 2r^f)T}} e^{\left(\rho - r^f\right)t}; \text{if } \rho \neq 2r^f \tag{99}$$

$$c^b - c_t = \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W\right]\frac{1}{T} e^{\left(\rho - r^f\right)t};\; \text{if } \rho = 2r^f$$

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

$$c^b - c_t = \left(c^b - c_0\right)e^{\left(\rho - r^f\right)t} \tag{100}$$

Figure 3 illustrates these paths.

As we expect, if $\rho = r^f = 0.05$, the investor wants a constant consumption path that spends the annuity value of wealth,

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

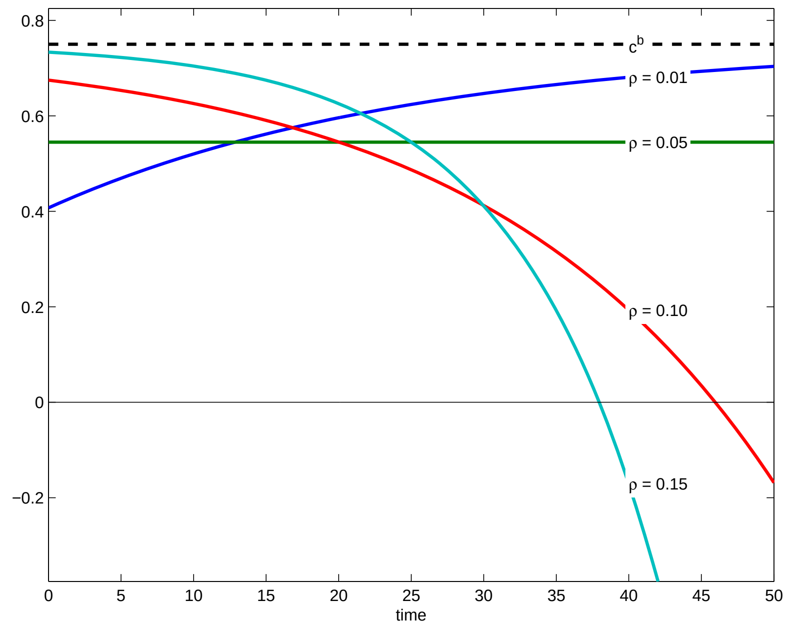


Figure 3: Consumption paths over time. $r = 0.05$, $\gamma = 2$, $T = 50$, $\rho$ as indicated.

If $\rho < r^f$, the investor wants a rising consumption profile. $(\rho - r^f)$ in (100) 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.

If $\rho > r^f$, the consumer is impatient, and wants a *declining* consumption path. $(\rho - r^f)$ in (100) 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.

Now, let's examine the troublesome case $\rho = 2r^f$, and consider how this finite-horizon economy behaves as we drive $T$ larger. Figure 4 presents the evolution of consumption, equation (99), for this case, and Figure 5 presents the evolution of the investor's wealth,

$$\frac{c^b}{r}\left(1 - e^{-r(T-t)}\right) - W_t = \left(\frac{c^b}{r}\left(1 - e^{-rT}\right) - W_0\right)\left(1 - \frac{t}{T}\right)e^{rt}.$$

We see the investor consume initially above the annuity value of wealth, spending down his wealth and then borrowing (negative wealth) to do so. Later in life, he radically reduces consumption in order, initially, to slow the rate of borrowing, and finally he reduces consumption so drastically that he pays back the debt very late in life and ends up satisfying the constraint $W_T = 0$.

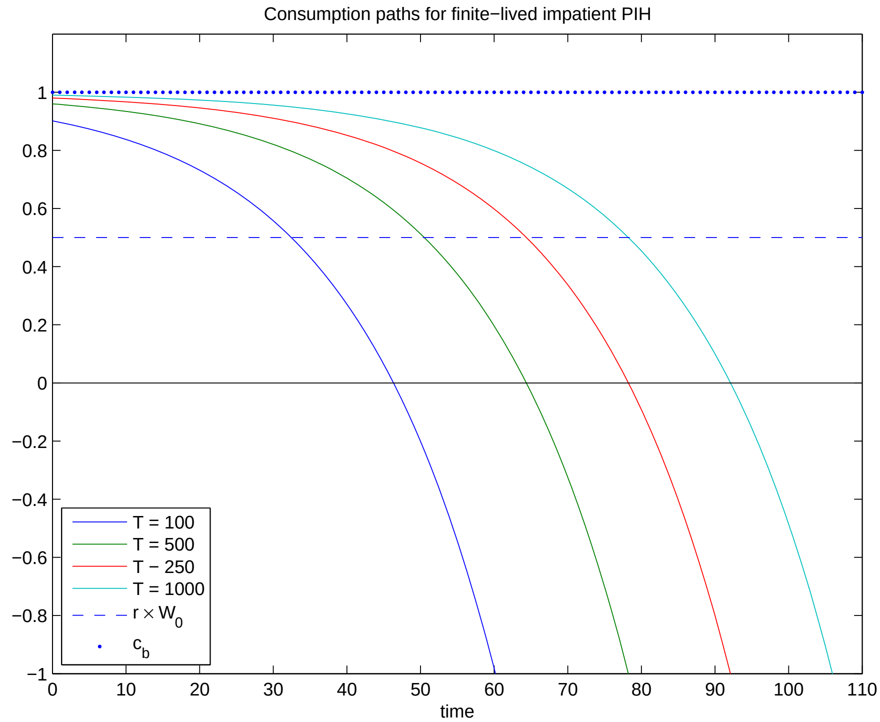


Figure 4: Consumption paths for a finite-lived $T$ impatient investor with quadratic utility. $c^b = 1$, $r = 0.05$, $rW_0 = 0.5$, $\rho = 2r = 0.10$.

Now, as $T$ grows, the investor pushes off the day of reckoning further and further, paying for early consumption closer and closer to the bliss point by longer and longer borrowing, promising a larger and larger negative consumption episode later in life. Each step of the limit is well defined: For each finite horizon, $W_T = 0$, $e^{-rT}W_T = 0$, and $W_0 = \int_{t=0}^{T} e^{-rt} c_t dt$. And the limit is well defined *for fixed* $t$. For each $t$, as $T \to \infty$, we have $c_t \to c_b$, and $W_t \to \frac{c^b}{r} - \left(\frac{c^b}{r} - W_0\right) e^{rt}$

In the limit, he promises an infinitely large negative consumption event, which is however infinitely postponed. so it never shows up.

But this limit *point* $c_t = c^b$ is invalid, and it is not the solution to the infinite-period problem. The first order condition remains

$$c^b - c_t = (c^b - c_0) e^{(\rho - r)t}.$$

If we choose the solution $c_0 = c_t = c^b$, the present value of consumption is

$$\int_{t=0}^{\infty} e^{-rt} c^b dt = \frac{c^b}{r} > W_0,$$

and

$$\lim_{t \to \infty} e^{-rt} W_t = -\left(\frac{c^b}{r} - W_0\right) < 0.$$

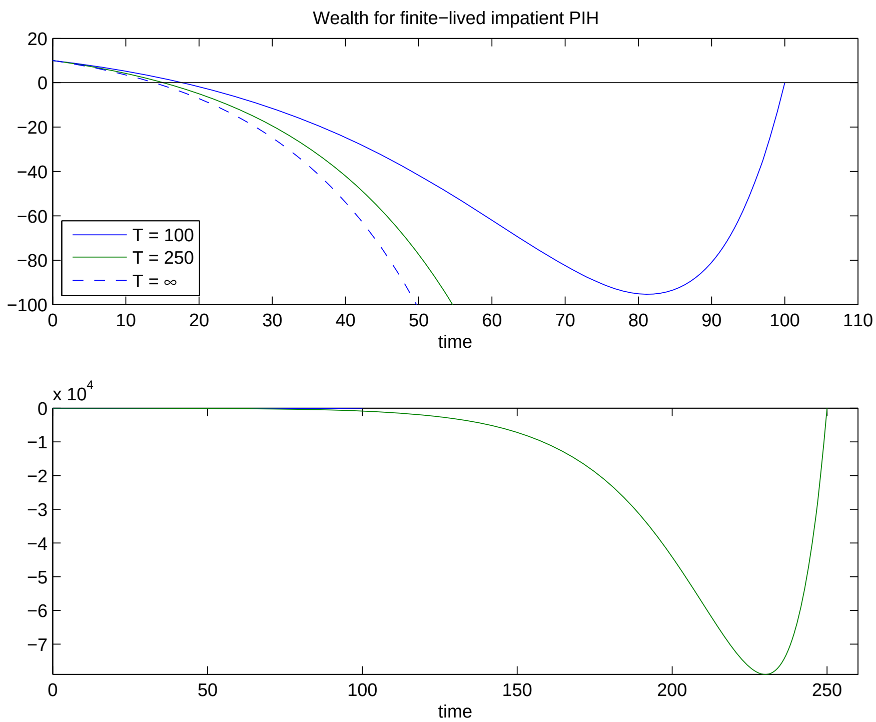


Figure 5: Evolution of wealth for a finite-lived $T$ impatient investor with quadratic utility.

This solution more than exhausts the initial wealth. The lump of infinitely negative consumption, infinitely delayed, which repays debt, vanishes from the limit point. Thus, $c_0 = c^b$ produces an infeasible solution. However, any $c_0 < c^b$ produces a suboptimal solution. Any $c_0 < c^b$ which obeys the first order condition produces wealth that grows *positively* at rate $e^{rt}$, and thus the consumer can do better by raising $c_0$. The terminal condition $\lim_{t\to\infty} e^{-rt} W_t$ is a discontinuous function of initial consumption $c_0$, and it is not equal to $W_0$ for any value of $c_0$.

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

Now consider a finite-horizon problem which focuses on risk bearing, simplifying away the intertemporal terms,

$$\max \left(-\frac{1}{2}\right) E \int_0^T e^{-\rho t}(c^b - c_t)^2 dt \text{ s.t.} dW_t = \left(r^f W_t - c_t\right) dt + \omega_t dr_t^e;\ W_0; W_T = 0; \rho = r^f.$$

The above intertemporal model is well studied in the permanent income literature. The presence of risky assets here is more 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 $\rho \neq r^f$, I focus on $\rho = r^f$ to simplify the formulas for $\mu \neq 0$.

In this case, the optimal consumption paths are

$$c^b - c_t = \left(c^b \frac{1-e^{-r^f T}}{r^f} - W_0\right)\left[\frac{r^f - \mu'\Sigma^{-1}\mu}{1-e^{-\left(r^f-\mu'\Sigma^{-1}\mu\right)T}}\right] e^{-\frac{1}{2}\mu'\Sigma^{-1}\mu t - \mu'\Sigma^{-1}\sigma\int_0^t dz_t}; \text{if } r^f \neq \mu'\Sigma^{-1}\mu \quad (101)$$

$$c^b - c_t = \left(c^b \frac{1-e^{-r^f T}}{r^f} - W_0\right)\left(\frac{1}{T}\right) e^{-\frac{1}{2}\mu'\Sigma^{-1}\mu t - \mu'\Sigma^{-1}\sigma\int_0^t dz_t}; \text{ if } r^f = \mu'\Sigma^{-1}\mu \quad (102)$$

Again, to understand the evolution of consumption over time for fixed $T$, it's easier still to write the answer in terms of initial consumption rather than solve out the budget constraint, as in (100): The consumption path is

$$c^b - c_t = \left(c^b - c_0\right) e^{-\frac{1}{2}\mu'\Sigma^{-1}\mu t - \mu'\Sigma^{-1}\sigma\int_0^t dz_t}, \quad (103)$$

and initial consumption is

$$c^b - c_0 = \left(c^b \frac{1-e^{-r^f T}}{r^f} - W_0\right)\left[\frac{r^f - \mu'\Sigma^{-1}\mu}{1-e^{-\left(r^f-\mu'\Sigma^{-1}\mu\right)T}}\right]; \text{if } r^f \neq \mu'\Sigma^{-1}\mu; \quad (104)$$

$$c^b - c_0 = \left(c^b \frac{1-e^{-r^f T}}{r^f} - W_0\right)\left[\frac{1}{T}\right] \text{ if } r^f = \mu'\Sigma^{-1}\mu. \quad (105)$$

Now, consumption is not *expected* to drift up or down, $E(c_t) = c_0$. However, (103) shows that consumption will drift up and down stochastically depending on the outcomes of the risky assets. If the risky assets do well, $\sigma \int_0^t dz_t > 0$, then consumption will drift up towards the bliss point, and vice versa. $c_t$ follows a lognormal distribution, bounded above by $c^b$ but with a long tail that extends to $-\infty$.

The investor's initial consumption $c_0$ is greater than the annuity value of initial wealth,

$$c_0 > \frac{r^f}{1-e^{-r^f T}} W_0.$$

(To see this, write (104)-(105) as

$$c^b - c_0 = \left(c^b - \frac{r^f}{1-e^{-r^f T}} W_0\right) \frac{1-e^{-r^f T}}{1-e^{-\left[r^f-\mu'\Sigma^{-1}\mu\right]T}} \left(\frac{r^f - \mu'\Sigma^{-1}\mu}{r^f}\right); \text{if } r^f \neq \mu'\Sigma^{-1}\mu;$$

$$c^b - c_0 = \left(c^b - \frac{r^f}{1-e^{-r^f T}} W_0\right) \frac{1-e^{-r^f T}}{r^f} \left(\frac{1}{T}\right) \text{ if } r^f = \mu'\Sigma^{-1}\mu.$$

The terms after the first parenthesis are positive and smaller than one; $c_0$ is closer to $c^b$ than is $\frac{r^f}{1-e^{-r^f T}} W_0$. ) He finances that greater consumption by undertaking negative consumption in extremely bad *states of nature*, where the $\rho \geq 2r$, $\mu = 0$ investor undertook negative consumption much later in *time* in order to finance higher consumption $c_0$.

The essential issue is the same as with the allocation over time problem. For fixed $t$, as $T$ increases, the consumption rules (101) and (102) approach $c_t = c^b$. That *limit point*, however, is an invalid solution of the infinite-period problem. At the limit point $T = \infty$, the negative consumption states never happen.

As $T$ increases, $c_0$ approaches $c^b$, so the first term in (103) 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 $r^f \leq \mu'\Sigma^{-1}\mu$, 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 display this behavior, Figure 6 plots the distribution of consumption in periods $t = 5$, $t = 10$ in models with horizon $T = 5$ and $T = 10$. Comparing the two $t = 5$ lines (green and blue), we see 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 (red) 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 cumulative 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* $T = \infty$.

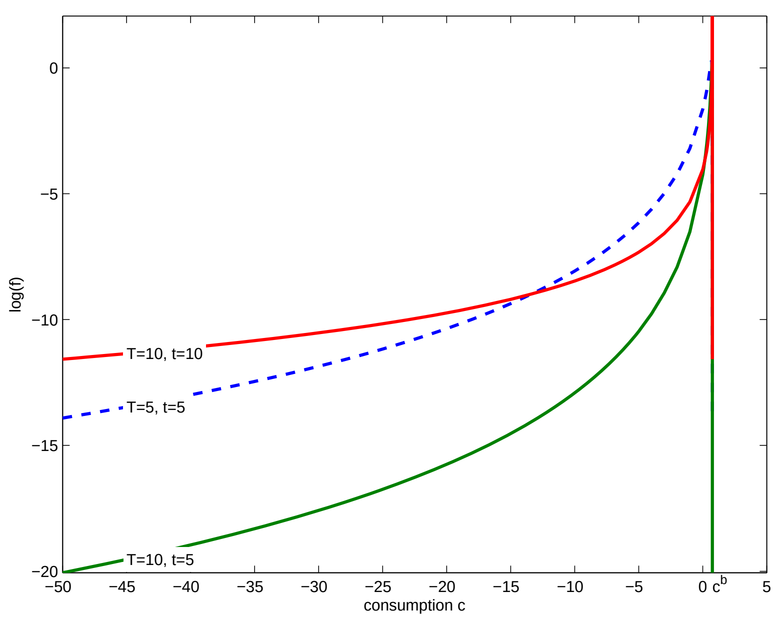


Figure 6: Log density of consumption at the last moment of a 5-year ($T = 5$) and 10-year ($T = 10$) model. Parameters are $r = \rho = 0.5$, $\mu = 0.08$, $\sigma = 0.16$, $\gamma = 2$, $W = 10$ and hence $c^b = 0.75$.

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:

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

where

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

denotes wealth that can support bliss point consumption. 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.

### 1.7.3 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

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

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

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

$$c_t = c^b - \left[p(c^b) - W_0\right] y_t^*$$

$$c_t = c^b - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W_0\right] \frac{x_t^*}{p(x^*)}$$

or, explicitly,

$$c_t = c^b - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W_0\right] \left(\frac{2r^f - \rho - \mu'\Sigma^{-1}\mu}{1 - e^{\left[(\rho - 2r^f) + \mu'\Sigma^{-1}\mu\right]T}}\right) e^{\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\sigma \int_0^t dz_t}; \text{if } \rho \neq 2r^f - \mu'\Sigma^{-1}\mu \quad (106)$$

$$c_t = c^b - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W_0\right] \left(\frac{1}{T}\right) e^{\left(\rho - r^f - \frac{1}{2}\mu'\Sigma^{-1}\mu\right)t - \mu'\Sigma^{-1}\sigma \int_0^t dz_t}; \text{ if } \rho = 2r^f - \mu'\Sigma^{-1}\mu$$

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

$$W_t = p_t(c_t) = p_t\left(c^b\right) - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W_0\right] \frac{p_t(x^*)}{p(x^*)}$$

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

$$W_t = c^b \frac{1 - e^{-r^f(T-t)}}{r^f} - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W_0\right] \frac{1 - e^{\left(\rho - 2r^f + \mu'\Sigma^{-1}\mu\right)(T-t)}}{1 - e^{\left(\rho - 2r^f + \mu'\Sigma^{-1}\mu\right)T}} x_t^* \text{if} 2r^f - \rho - \mu'\Sigma^{-1}\mu \neq 0 \quad (107)$$

$$W_t = c^b \frac{1 - e^{-r^f(T-t)}}{r^f} - \left[c^b \frac{1 - e^{-r^f T}}{r^f} - W_0\right] \frac{T - t}{T} x_t^*; \text{if} 2r^f - \rho - \mu'\Sigma^{-1}\mu = 0$$

or, explicitly,

$$W_t = c^b \frac{1-e^{-r^f(T-t)}}{r^f} - \left[c^b \frac{1-e^{-r^f T}}{r^f} - W_0\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} e^{\left(\rho-r^f-\frac{1}{2}\mu'\Sigma^{-1}\mu\right)t-\mu'\Sigma^{-1}\sigma\int_0^t dz_t} \quad (108)$$

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

Note $W_T = 0$ – all wealth *is* exhausted, but the investor dies without leaving debts, and $W_0 = W_0$, these solutions respect the initial condition.

To express consumption in terms of wealth, we substitute (108) in (106), yielding

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

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

To find portfolios, we look at the diffusion component of $dW$. From (107),

$$dW_t = ...dt - \left[c^b \frac{1-e^{-r^f T}}{r^f} - W_0\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} dx_t^*$$

$$dW_t = ...dt + \left[c^b \frac{1-e^{-r^f T}}{r^f} - W_0\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} x_t^* \mu'\Sigma^{-1}\sigma dz_t$$

Thus the portfolio weight on risky assets is.

$$w_t = \left[c^b \frac{1-e^{-r^f T}}{r^f} - W_0\right] \frac{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)(T-t)}}{1-e^{\left(\rho-2r^f+\mu'\Sigma^{-1}\mu\right)T}} x_t^* \mu'\Sigma^{-1}$$

It is useful to express this weight with wealth $W_t$ as the state variable. Substituting from (108) 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 $\gamma_t$ However, since most of the analysis concerns $W_t < 0$ where this definition of $\gamma$ doesn't make any sense, this expression is not so useful in studying limits.

## 1.8 Lognormal returns and mean-variance frontiers

The trouble revealed by the $p(x^*) = \infty$ investigation is that 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, nor will it generate this opportunity for quadratic utility investors to sell off infrequent high-price states.

The deeper lesson is that one should pair distributional assumptions and preferences. Portfolio theory makes the most sense when the assumed return distribution is characteristic of what a market of investors with the given preferences will produce, allowing for defined dimensions of investor heterogeneity.

The strange behavior of the long-run mean-variance frontier in the i.i.d. lognormal environment really has nothing to do with long-run mean-variance frontiers per se; it is mirrored in the traditional discrete-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 (2005) look at the mean-variance properties of long-run *log* returns, which are well-behaved, but not the solution to a portfolio problem. Martin (2012) is a nice paper which expands on the long-run pathologies of lognormal models.)

To understand this claim, let us examine the mean-variance properties of $T$ period returns, in the standard lognormal i.i.d. 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

$$E\left(\frac{S_T}{S_0}\right) = e^{\mu T}$$

$$\sigma\left(\frac{S_t}{S_0}\right) = e^{\mu T}\sqrt{e^{\sigma^2 T}-1}$$

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 or mean-variance 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:

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

The max sharpe ratio increases without bound!

To find the Sharpe-ratio maximizing payoff/portfolio, we can exploit the fact that $\Lambda_T/\Lambda_0$ is on the mean-variance frontier, and then find the self-financing strategy that yields the final payoff $\Lambda_T/\Lambda_0$. 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

$$W_t = E_t\left(\frac{\Lambda_T}{\Lambda_t}W_T\right) = W^b e^{-r(T-t)} - \lambda E_t\left(\frac{\Lambda_T}{\Lambda_t}\frac{\Lambda_T}{\Lambda_0}\right) = W^b e^{-r(T-t)} - \lambda V_t \quad (109)$$

Using time-0 wealth to eliminate $\lambda$,

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

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 (109),

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

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.

## 1.9 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 7 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 7 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.

Together with the fact in Figure 7 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.

## 1.10 Time-varying returns with unspanned state variables

Finding the payoff spaces and discount factors based on return models with time-varying investment opportunities and incomplete markets is trickier than it may appear.

For example, suppose the risk free rate and risky returns are controlled by a state

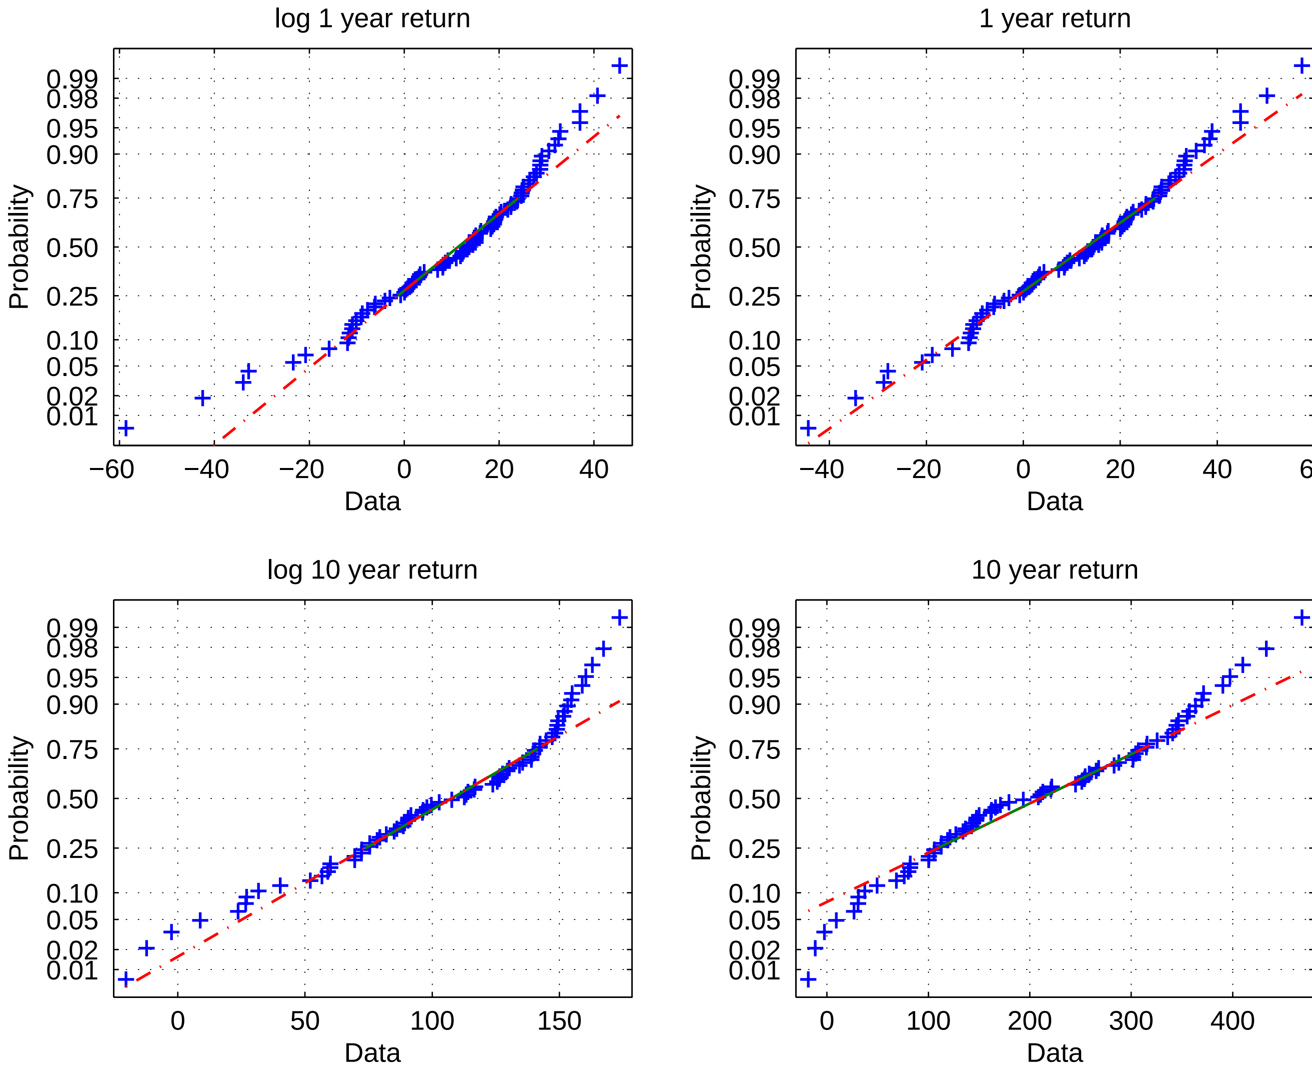


Figure 7: Distribution of CRSP value-weighted return, 1927-2006. Each plot presents the cumulative distrbibution of the data (blue cross), with the y axis scaled so that normally distributed data will fall on a line. The red dashed line is fit to the first and third quartiles to help assess linearity (matlab function normplot). The 10 year returns are overlapping annual observations. The units are percent in all cases.

variable $z_t$,

$$\begin{aligned} dr_t &= \left[ r^f(z_t) + \mu(z_t) \right] dt + \sigma(z_t) dB_{1t} \\ dz_t &= \mu_z(z_t)dt + \sigma_{z1}(z_t)dB_{1t} + \sigma_{z2}(z_t)dB_{2t}. \end{aligned}$$

In order to price $r_t^f$ and $dr_t$, the discount factor must be of the form

$$\frac{dx_t^*}{x_t^*} = \left[ \rho - r^f(z_t) \right] dt - \frac{\mu(z_t)}{\sigma(z_t)} dB_{1t} + \sigma_{2t}^* dB_{2t}. \tag{110}$$

$(E_t(dx_t^*/x_t^*) = \left(\rho - r^f\right) dt$ and $E_t\left(dr_t - r^f dt\right) = -E_t(dr_t dx_t^*/x_t^*)$.) For portfolio theory, we have to choose correctly $\sigma_{2t}^*$ so that $u'(\hat{x}_t) = \lambda x_t^*$ produces a tradeable $\hat{x}_t$. With quadratic utility that requirement simply means producing a traded $x_t^*$.

One might think that the *traded* discount factor payoff is simply generated by $\sigma_{2t}^* = 0$. After all, the resulting discount factor is the only one whose shocks are spanned by the shock $dB_{1t}$ to the traded assets. Alas, this is a subtle mistake. "Traded" here means an achievable *payoff* from a valid dynamic trading strategy, not an achievable *portfolio return.* Having shocks spanned by the asset return shocks $\sigma_{2t}^* = 0$ is neither necessary or sufficient for $x_t^*$ to be achievable as the dividend process of a valid trading strategy, i.e. to be a traded payoff. Typically, in fact, traded discount factors will have $\sigma_{2t}^* \neq 0$.

As an example, consider a finite-horizon economy with payoffs from 0 to $T$, a varying interest rate

$$dr_t^f = \sigma dB_t,$$

and no risky assets. The shock $dB_t$ is not in the span of traded returns. A "traded payoff" is an $\{x_t\}$ that results from the trading strategy

$$dV_t = \left(r_t^f V_t - x_t\right) dt \tag{111}$$

with $V_0 < \infty$, $V_T = 0$, and hence

$$V_0 = \int_{t=0}^{T} e^{-\int_{s=0}^{t} r_s ds} x_t dt. \tag{112}$$

The right hand side holds ex-post, so cannot depend on the realization of $\{dB_t\}$. As in the permanent-income example, a traded payoff $x_t$ can load on unspanned shocks such as $dB_t$, but if $x_t$ draws down wealth based on some non-traded shock, a future value of $x_{t+s}$ must pay back the debt so the strategy always ends up at $V_T = 0$.

In this example, the discount factor from (110) must be of the form

$$\frac{dx_t^*}{x_t^*} = \left(\rho - r_t^f\right) dt + \sigma_t^* dB_t.$$

The obvious candidate $\sigma_t^* = 0$, though the unique traded *portfolio* return, does not give a traded *payoff.* With $\sigma_t^* = 0$, we have

$$\frac{x_t^*}{x_0^*} = e^{\int_{s=0}^{t} \left(\rho - r_s^f\right) ds}. \tag{113}$$

Then (112) becomes

$$V_0 = x_0^* \int_{t=0}^{T} e^{\int_{s=0}^{t} (\rho - 2r_s) ds} dt.$$

The right hand side depends on the realization of $\{dB_t, 0 < t < T\}$ through $r_s$. Intuitively, a low realization of $r_t^f$ means that $x_t^*$ will rise through (113). Money will be paid out from the account generating $x_t^*$ in (111), lowering $V_t$. But with this $x^*$ process, money will not be paid back in the future, and since the interest rate is now lower, $V_t$ cannot recover by a rising rate of return either. That's why we see $2r^f$ in the formula. We need $\sigma_t^* \neq 0$ to generate a traded *payoff.*

# 2 Growing bliss points

A second pathology of global applications of the very simple model presented so far is that wealth never grows higher than the amount necessary to pay for bliss point consumption forever, $W_t < c^b/r^f$, and consumption never grows past the bliss point, $c_t < c^b$. This is easy to see in (94) and (96). Even though the underlying technologies allow growth, the investor chooses to eat rather than invest, and to invest more conservatively, as wealth grows. This is certainly not a good feature to bring to data that allow long-run growth. Even if the right tail of long-run wealth is thinner than a lognormal, it does not stop at zero.

One natural approach, recommended by Hansen and Sargent (2004), is to specify a deterministically or stochastically growing bliss point to accommodate growth. Here are three examples. In each example, the previous analysis applies to the difference between consumption and growing bliss points, which thus can accommodate growth.

## 2.1 Geometrically growing bliss points

Consider a geometrically growing bliss point $c_t^b = c^b e^{gt}$.

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

is unaffected by this change. Thus, the new consumption process is simply

$$c^b e^{gt} - c_t = \left(c^b - c_0\right) x_t^* \tag{114}$$

or, explicitly

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

where, taking the time-zero price of (114) and rearranging,

$$c^b - c_0 = \left(\frac{c^b}{g - r^f} - W_0\right)\left(2r^f - \rho - \mu'\Sigma^{-1}\mu\right)$$

Similarly, wealth follows

$$\left(\frac{c^b e^{gt}}{r - g} - W_t\right) = \left(\frac{c^b}{g - r^f} - W_0\right) x_t^*$$

or, explicitly

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

This modification rather transparently allows growth, though deterministic growth. All the previous characterizations apply to the difference between consumption and the growing bliss point $c^b e^{gt}$ rather than between consumption and the fixed bliss point $c^b$. Yes, consumption never exceeds the growing bliss point, and wealth never exceeds its growing perpetuity value. But at least consumption and wealth now grow over time.

## 2.2 Stochastic bliss point

By allowing a stochastic bliss point, we can allow the upper limit to grow stochastically. This structure also allows us to approximate any utility function. As an extreme example, we can construct a stochastic bliss point following which, the quadratic utility investor will choose exactly the power utility consumption and wealth process.

Denote by $c_t^p$ the consumption process for the power utility investor, starting with wealth $W_0$, e.g. from (98),

$$c_t^p \equiv c_0^p e^{\frac{1}{\gamma}\left[\left(r^f+\frac{1}{2}\mu'\Sigma^{-1}\mu-\rho\right)t+\mu'\Sigma^{-1}\sigma\int_0^t dB_s\right]} \tag{115}$$

$$c_0^p = \frac{1}{\gamma}\left[\rho + (\gamma-1)\left(r^f + \frac{1}{2}\frac{1}{\gamma}\mu'\Sigma^{-1}\mu\right)\right] W_0$$

and denote by $c_t^q$ the consumption process for the quadratic utility investor. $c_t^p$ is derived to follow the power utility investor's first order conditions,

$$c_t^{p-\gamma} = c_0^{p-\gamma} x_t^*$$

and $p(c_t^p) = W_0$ .

Now, regard (115) as an exogenous stochastic process, and suppose the quadratic utility investor has bliss point

$$c_t^b = c_t^{p-\gamma} + c_t^p. \tag{116}$$

The quadratic utility investor's first order conditions are then

$$c_t^b - c_t^q = \left(c_0^b - c_0^q\right) x_t^*$$

$$c_t^{p-\gamma} + c_t^p - c_t^q = \left(c_0^{p-\gamma} + c_0^p - c_0^q\right) x_t^*$$

$$c_0^{p-\gamma} x_t^* + c_t^p - c_t^q = \left(c_0^{p-\gamma} + c_0^p - c_0^q\right) x_t^*$$

$$c_t^p - c_t^q = \left(c_0^p - c_0^q\right) x_t^*$$

Together with $p(c_t^p) = p(c_t^q) = W_0$, we conclude that

$$\hat{c}_t^q = c_t^p$$

where I have added a ˆ to emphasize that this is the optimal choice.

In sum, if we specify the bliss point as in (116), the quadratic utility investor will choose *exactly* the same wealth and consumption process as the power utility investor.

This result could be useful as the basis of an approximation. Choose this bliss point so that quadratic and power match in this simple i.i.d. environment. Then tweak the environment to a more interesting specification.

## 2.3 Bliss points tied to past consumption

While putting in a bliss point that is a function of technology shocks is defensible as a way to produce an approximation to another model – find the point at which we approximate power by quadratic utility – it doesn't make much economic sense.

The point is to allow the bliss point to grow so that the quadratic utility investor stays below the bliss point, yet wealth can grow. Introducing a temporal nonseparability is a more appetizing route.

Suppose we write the period utility function as

$$u(c) = -\frac{1}{2}\left(c^b - c_t - \int_{s=0}^{\infty} b(\tau)c_{t-\tau}d\tau\right)^2$$

In the conventional interpretation, $b(\tau) > 0$ represents habit persistence, in which past consumption raises marginal utility, and $b(\tau) < 0$ represents durability. In this context, the $b(\tau)$ function amounts to a stochastic bliss point, which moves in reaction to past consumption decisions. Shifting the bliss point up lowers risk aversion. Thus, $b(\tau) < 0$ means that after consumption rises, the bliss point will rise, and offset the usual rise in risk aversion in quadratic models. Cochrane (2012b) works out asset pricing formulas for this specification.

# 3 References

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

Cochrane, John H., 2012a, A mean-variance benchmark for intertemporal portfolio theory, Manuscript, University of Chicago.

Cochrane, John H., 2012b, A continuous-time asset pricing model with habits and durability, Manuscript, University of Chicago.

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

Hansen, Lars Peter and Thomas J. Sargent, 2004, *Recursive Models of Dynamic Linear Economies*, Manuscript, University of Chicago and New York University.

Kim, Tong Suk and Edward Omberg, 1996, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141-61.

Martin, Ian, 2012, On the Valuation of Long-Dated Assets, *Journal of Political Economy* 120, 346-358