Stochastic Calculus in Finance
MATH 6910 - Salisbury

Introduction to Portfolio Optimization

(1) We work in a simple market with a stock – worth S_t – and a money market fund – worth R_t – whose evolution is (as usual)

$$dS_t = \mu S_t\,dt + \sigma S_t\,dB_t, \qquad dR_t = rR_t\,dt.$$

It is worth emphasizing that we are working under the physical measure P, under with B_t is a Brownian motion. Our wealth X_t consists of a self-financing portfolio invested in the stock and the bond:

$$dX_t = \phi_t S_t + \psi_t R_t.$$

The problem we want to answer is to find the *optimal* choice of ϕ_t and ψ_t. But "optimal" will mean different things to different people - risk averse individuals will interpret it differently from more aggressive investors. To account for risk preferences we specify a *utility function* $u(x)$, which is concave and increasing. "Increasing" means the impact on your utility because of gaining \$1 is always positive, but "concave" means that the magnitude of this impact decreases the wealthier you get.

We'll work with the special utilities

$$u(x) = \frac{x^{1-\gamma}}{1-\gamma}$$

where $\gamma > 0$ is a risk-aversion parameter.
[In fact we also need $\gamma \neq 1$ here. We could have got rid of that condition by using $\tilde{u}(x) = (x^{1-\gamma} - 1)/(1-\gamma)$ instead, which approaches $\log x$ as $\gamma \to 1$. Note also that $\gamma = 0$ would correspond to a linear or *risk-neutral* utility. Our utility goes by the name of CRRA, for Constant Relative Risk Aversion (I won't explain why). Increasing γ corresponds to becoming more risk averse.]

Our basic problem is to choose the portfolio to maximize $E[u(X_T)]$, where T is some fixed horizon (eg retirement).

(2) We'll reformulate the problem in terms of π_t, which will denote the proportion of total wealth invested in the stock. We normally think of having $0 < \pi_t < 1$, but in fact $\pi_t < 0$ (short selling) or $\pi_t > 1$ (leverage, or borrowing in order to purchase stock) are allowed.

The total value of our stock holdings are $\phi_t S_t$ which therefore equals $\pi_t X_t$. Likewise the total value of our money market holdings are $\psi_t R_t = (1-\pi_t)X_t$.

So the self-financing condition is that

$$\begin{aligned} dX_t &= \phi_t\, dS_t + \psi_t\, dR_t \\ &= \phi_t \mu S_t\, dt + \phi_t \sigma S_t\, dB_t + \psi_t r R_t\, dt \\ &= \mu \pi_t X_t\, dt + \sigma \pi_t X_t\, dB_t + r(1-\pi_t) X_t\, dt \\ &= r X_t\, dt + (\mu - r)\pi_t X_t\, dt + \sigma \pi_t X_t\, dB_t. \end{aligned}$$

There is a given initial value x_0 for wealth, so $X_0 = x_0$. And when necessary we'll write X_t^π for X_t to remind us that wealth depends on the chosen portfolio.

What conditions must π_t satisfy? First it must be adapted to the filtration $\mathcal{F}_t$ of B_t (so we can't peek into the future). Second, it turns out we have to rule out what are called *suicide strategies* (like the gambling trick of "doubling up"). There are various ways of doing this, but the most elementary is to choose a large value M and insist that $|\pi_t| \le M$ (eg that there is a limit to how leveraged the portfolio can be). Strategies π_t satisfying these two conditions are said to be *admissable*. And our problem becomes to find

$$\boxed{V_0 = \sup_{\text{admissable } \pi_t} E[u(X_T^\pi)]} \tag{A}$$

More precisely, we want to find the choice of π_t that realizes this supremum.

(3) Merton was the first to solve this problem. There are many variants, and what I want to emphasize is really the method. But we start with the solution itself, which turns out to be

$$\boxed{\pi_t = \frac{\mu - r}{\gamma \sigma^2}}$$

Note that
- The solution does not depend on t, T, or wealth x. In this sense it is *myopic* - it doesn't change over time, as the horizon T gets closer.
- $\pi_t > 0$ so there is no short-selling. But we may have $\pi_t > 1$ (leverage), for example if γ is small, or if μ is big.
- A static asset allocation doesn't mean you don't trade. In fact, the portfolio must be continuously rebalanced to maintain the allocation. But the effect is to force you to sell stocks when they go up in price, and to buy them when they are go down. In practice, financial advisors do often tell people to maintain a roughly constant asset allocation, with rebalancing taking place a few times per year.

(4) There are several ways of obtaining this solution. One, which I won't give details for, is to use the risk neutral measure in two steps:
- Focus on the terminal wealth $Y = X_T$, and choose it to maximize $E[u(Y)]$ subject to a constraint. Namely that the risk neutral expectation $\tilde{E}[e^{-rT}Y] =$

x_0, the initial wealth. This follows from the fact that the discounted portfolio value is a risk-neutral martingale. This optimization is usually done using *duality methods* (basically Lagrange multipliers).

- Then extract the asset allocation π_t that realizes the martingale $\tilde{E}[e^{-rT}Y|\mathcal{F}_t]$ as a discounted portfolio $e^{-rt}X_t$.

(5) Instead the approach I'll follow is that of *Stochastic Control Theory* and dynamic programming. This involves moving backwards from the horizon T using the process

$$V_t = \sup_{\text{admissable } \pi} E[u(X_T^\pi)|\mathcal{F}_t]$$

in which we start with a given wealth X_t at time t and then manage the portfolio optimally between times t and T. The variable you *control* is π_t. The *Hamilton-Jacobi-Bellman equations* for the *value function* $v(t,x)$ are that $\sup_\pi L_\pi v(t,x) = 0$ for every t, x, where

$$L_\pi v = v_t + v_x[r + (\mu - r)\pi]x + \frac{1}{2} v_{xx}\sigma^2\pi^2x^2. \tag{B}$$

We then have the following

HJB verification theorem:

Suppose that we can find

- a smooth function $v(t,x)$ with $v(T,x) = u(x)$ such that $v_t + v_x[r + (\mu - r)\pi]x + \frac{1}{2}v_{xx}\sigma^2\pi^2x^2 \le 0$ for every t, x, π;
- a function $\pi^*(t,x)$ such that $v_t + v_x[r + (\mu - r)\pi^*]x + \frac{1}{2}v_{xx}\sigma^2(\pi^*)^2x^2 = 0$ for every t, x;
- a wealth process X_t^* corresponding to a portfolio strategy π_t such that $\pi_t = \pi^*(t, X_t)$.

Then the solution to (A) is $V_0 = v(0, x_0)$. More generally, $V_t = v(t, X_t)$ and the optimal control is the above π_t.

(6) **Proof:**

First let π_t be any admissable strategy, and X_t the corresponding wealth process. By Ito's lemma,

$$dv(s, X_s) = [v_s + v_x[r + (\mu - r)\pi]X_s + \frac{1}{2}v_{xx}\sigma^2\pi_s^2X_s^2]\,ds + v_x\sigma\pi_sX_s, dB_s.$$

The ds term is negative, so when we integrate we get

$$u(X_T) = v(T, X_T) \le v(t, X_t) + \int_0^t v_x\sigma\pi_sX_s\,dB_s.$$

Using the boundedness of π_s we can show that the integrand is square integrable, so that the stochastic integral has mean zero (this is the only place we use the bound M that came up when we defined admissability). Thus $E[u(X_T)|\mathcal{F}_t] \le v(t, X_t)$ and so $V_t \le v(t, X_t)$.

Going through the same argument using $\pi_s = \pi^*(s, X_s^*)$ gives that $E[u(X_T^*)|\mathcal{F}_t] = v(t, X_t)$. Thus V_t actually $= v(t, X_t)$ and $\pi_s = \pi^*(s, X_s)$ is the optimal allocation.

(7) **Merton's solution to HJB.**

We'll show that there is a solution to HJB of the form $v(t, x) = f(t)x^{1-\gamma}/(1-\gamma)$ with $f(t) > 0$ and $f(T) = 1$. Then $v_{xx} = -\gamma f(t)x^{-\gamma-1} < 0$. So as a function of π, (B) is a downward pointing parabola, which does indeed have a maximum at

$$\pi = -\frac{(\mu - r)v_x}{\sigma^2 x v_{xx}}. \tag{C}$$

Substituting back into (B) gives that

$$v_t + rxv_x - \frac{(\mu - r)^2 v_x^2}{2\sigma^2 v_{xx}} = 0$$

(which also goes by the name of the HJB equation). In some other circumstances we have to solve this kind of nonlinear PDE numerically. But in this case we can substitute in our conjectured form for the solution and get

$$\frac{x^{1-\gamma}}{1-\gamma}\Big[f'(t) + f(t)r(1-\gamma) + f(t)\frac{(1-\gamma)(\mu - r)^2}{2\gamma\sigma^2}\Big] = 0.$$

This is an ODE for f, whose solution is $f(t) = e^{(T-t)(1-\gamma)[r + \frac{(\mu-r)^2}{2\gamma\sigma^2}]}$. Not only does this give us an explicit formula for $v(t, x)$, but we can substitute back into (C) and get the earlier formula $\pi^* = (\mu - r)/\gamma\sigma^2$. And now the verification theorem tells us that this solves our problem.

(8) **Derivation of HJB.**

The verification theorem is nice and tidy, but it doesn't explain how we come up with HJB in the first place. In practice one goes through the following heuristic to figure out what PDE to try and solve. And once you solve the PDE you go back and use the verification theorem to show that your (tentative) solution actually works.

Pretend we know that there is an optimal control π_t^*. Then $V_t = E[u(X_T^{\pi^*})|\mathcal{F}_t]$. Note that now we have conditional expectations of a single random variable (ie no more "$\sup_\pi$" term). So V_s is a martingale, when we use the optimal control.

What if we use some other sub-optimal control π_s? Define a new control $\tilde{\pi}_s$ which consists of following π_s for $s \le t$ and then switching to π^* for $t < s \le T$.

Then for $q < t$,

$$\begin{aligned} E[v(t, X_t^{\pi})|\mathcal{F}_q] &= E[v(t, X_t^{\tilde{\pi}})|\mathcal{F}_q] \\ &= E\Big[E[u(X_T^{\tilde{\pi}})|\mathcal{F}_t]|\mathcal{F}_q\Big] \\ &= E[u(X_T^{\tilde{\pi}})|\mathcal{F}_q] \le v(q, X_q^{\pi}). \end{aligned}$$

In other words $v(s, X_s^{\pi})$ is a *supermartingale* (that is, a process Y_t with the property that $E[Y_t|\mathcal{F}_q] \le Y_q$ whenever $q < t$).

Supermartingales drift downwards, so if we compute $dv(t, X_t)$ using Ito's lemma, we get that the dt term is ≤ 0 for general controls, and $= 0$ for the optimal control. Those two statements are just what HJB says.