Asset Pricing Review

Cochrane ch. 1.

- One of the workhorse models of modern asset pricing is the *consumption–based* asset pricing model.
- It falls out of the investor's consumption/investment problem.
- We will use it to introduce many issues in asset pricing that we will investigate empirically later on.

Risk Corrections

- It is well understood that more risky assets should pay more in expectation than less risky ones.
- What is less well understood is what the particular measure of risk should be.
  - For example, with the CAPM, the appropriate measure of risk is market beta, which is a (scaled) covariance of asset return with that of the market portfolio.
- If marginal utility is a measure of how you make tradeoffs, say, between consumption today and asset holdings, then one might expect that the covariance of asset returns with marginal utility governs risk premia.

Empirical Content of the Models

- Is due to additional restrictions.
- The restrictions relate marginal utility to observables, like aggregate consumption, or the return on the market portfolio.
- For example, the CAPM is a special case (i.e. returns normally distributed and/or quadratic utility) where in equilibrium, the market portfolio is MVE.

Deriving the basic pricing equation

We start with (almost) no assumptions.

- Consider a two period environment.
- We want to find the value at time t of a *payoff*[a] x_{t+1}.
- For example: buy a stock now and get the price appreciation plus dividend:

$$x_{t+1} = p_{t+1} + d_{t+1}$$

- The payoff is random because next period's stock price is.

[a] Not return, which we'll look at separately.

The Investor's Problem

- Start with (two period additively separable) utility in terms of consumption, c that obeys standard assumptions (e.g. concavity, etc.):

$$U(c_t, c_{t+1}) = u(c_t) + \beta E_t[u(c_{t+1}]$$

- Often, we'll have power utility:

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

or log utility:

$$\lim_{\gamma \to 1} \left[u(c_t) = \frac{1}{1-\gamma} c_t^{1-\gamma} \right] = ln(c_t).$$

Interpretation

- β captures impatience and is sometimes called the subjective discount factor.
- γ captures aversion to risk and also aversion to *intertemporal substitution.*
  - investors prefer steady consumption across states of nature and time.

The agent's problem, securities.

- Let the investor buy and sell the payoff, x_{t+1}, at price p_t.
- Let $e_\tau, \tau = \{t, t+1\}$ be his known endowments in each period.
- Let ξ be the amount of asset chosen.

Mathematical Representation

$$max_{(\xi)} u(c_t) + E_t[\beta u(c_{t+1})] s.t.$$

$$c_t = e_t - p_t \xi$$

$$c_{t+1} = e_{t+1} + x_{t+1} \xi$$

Now either substitute in the constraints or form a lagrangian and take $\frac{\partial}{\partial \xi}$.

The First-order Condition

Also called the Euler equation, this equates marginal utilities:

$$p_t u'(c_t) = E_t[\beta u'(c_{t+1}) x_{t+1}] \tag{1}$$

or

$$p_t = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} x_{t+1} \right] \tag{2}$$

This is *the* central asset pricing formula. All else follows from additional assumptions. We've not solved the model completely, but can still get a lot from this representation (after adding additional restrictions).

The MRS as a SDF

People have been defining the stochastic discount factor, $\equiv m_{t+1}$ as

$$m_{t+1} \equiv \beta \frac{u'(c_{t+1})}{u'(c_t)} \tag{3}$$

Then, the basic pricing formula is:

$$p_t = E_t \left[m_{t+1} x_{t+1} \right] \tag{4}$$

Sometimes, when it is clear, we'll write $p = E[mx]$.

Why SDF?

Because it generalizes standard ideas about discount factors. Without uncertainty, the PV formula gives:

$$p_t = \frac{1}{R^f} x_{t+1} \tag{5}$$

Then, R^f is the gross risk-free rate and $\frac{1}{R^f}$ is the discount factor. Riskier assets have lower prices to account for risk (usually)

$$p_t^i = \frac{1}{R^i} E_t[x_{t+1}^i]$$

Interpretation of $p_t = E_t\left[m_{t+1} x_{t+1}\right]$

- One can incorporate *all* risk corrections using a *single* discount factor - the same one for each asset.
- It is stochastic because at time t, m_{t+1} is not known with certainty.
- The correlation between the discount factor and asset return generates the risk corrections.

Gross Returns

- The gross payoff divided by today's price:

$$R_{t+1} = \frac{x_{t+1}}{p_t}$$

Therefore, returns are payoffs with price one:

$$1 = E[mR]$$

- This special case is something we will devote most of our time to.
- It's useful because returns are stationary over time, so that sample moments converge to their population counterparts.

Everything is not a return

- This is obvious, but importantly for us, consider borrowing a dollar today at interest rate R^f and investing it in an asset that pays expected return $E[R]$.
- Then, you have a *zero investment (zero cost) portfolio* that pays an *excess return*

$$R^e \equiv E[R] - R^f.$$

We will deal with excess returns often (CLM calls them Z). Remember, zero price does not imply zero payoff.

Conditioning Information

- Notice that we write

$$E_t[m_{t+1}R_{t+1}] = 1$$

- This is shorthand for

$$E[m_{t+1}R_{t+1}|I_t] = 1$$

where I_t is the agent's information set at time t.

- This implies that our asset pricing model is conditional on information the agent currently possesses.
- Later on, when we talk about overidentifying restrictions, we will use things in the agent's information set as *instruments*.

And Market Efficiency

- Instruments are a useful way of imposing market efficiency to estimate a model's parameters (and test a model).
- To see this, rewrite our asset pricing equation as:

$$E_t[m_{t+1}R_{t+1} - 1] = 0$$

We can think of the quantity in brackets as a "one step ahead forecast error."

- This is because in expectation, the discounted future payoff equals the price, one, but the actual realization will typically not. If the actual return is high (low) , then the forecast error will be positive (negative), and is zero on average.

- Note that returns can still be predictable if m is.
- This will be important later on — predictability is not *prima facie* evidence against market efficiency.

Empirically,

- It also must be true that things in the information set, i.e. past returns, should be unable to predict forecast errors — our asset pricing model is conditional on this information set.
- Hence, things in the information set must be orthogonal to the forecast errors:

$$E_t[(m_{t+1}R_{t+1} - 1)(z_t \in I_t)] = 0$$

which is a statement about market efficiency.

Classic Issues in Finance

We'll manipulate the basic asset pricing model to look at

- The economics of interest rates
- Risk adjustments
- Expected return - beta representations
- The mean variance frontier
- The slope of the frontier
- Time-varying expected returns

The Risk-free rate

The risk-free asset return pays off R^f in every state so that

$$1 = E[m_{t+1}R^f] = R^f E[m_{t+1}]$$

so that

$$R^f = \frac{1}{E[m]}$$

If the risk-free security is not traded, then we define $R^f = \frac{1}{E[m]}$ as a shadow risk-free rate, or (zero-beta) rate.

What affects R^f?

Use an example: $u'(c) = c^{-\gamma}$. Then

$$R^f = \frac{1}{\beta}\left(\frac{c_{t+1}}{c_t}\right)^{\gamma}$$

Note 3 effects:

3 effects

1. Real rates are high when β is low (people are impatient).
2. Real interest rates are high when consumption growth is high because it pays to invest more now to consume more later. High rates encourage saving.
3. Real rate are more sensitive to consumption growth when curvature (γ) is high. High γ makes the investor care more about maintaining a consumption profile that is smooth over time and less sensitive to interest rate incentives. it takes a larger interest rate change to induce consumption (un)smoothing.

Uncertainty

Refine our example: let consumption growth be lognormally distributed. This is a trick we will see again and again throughout the course. Then,

$$R_t^f = \frac{1}{E_t\left[\beta\left(\frac{c_{t+1}}{c_t}\right)^\gamma\right]}$$

Now, a normal z means that

$$E(e^z) = e^{E(z)+\frac{1}{2}\sigma^2(z)}$$

Then, with $\beta = exp(-\delta)$, $\Delta lnc_{t+1} = ln(c_{t+1}/c_t)$,

$$R_t^f = \left[e^{-\delta}e^{-\gamma E_t(\Delta lnc_{t+1})+(\gamma^2/2)\sigma_t^2(\Delta lnc_{t+1})}\right]^{-1}$$

Taking logs:

$$r_t^f = \delta + \gamma E_t(\Delta lnc_{t+1}) - \frac{\gamma^2}{2}\sigma_t^2(\Delta lnc_{t+1}). \quad (6)$$

$$r_t^f = \delta + \gamma E_t[\Delta lnc_{t+1}] - \frac{\gamma^2}{2}\sigma_t^2(\Delta lnc_{t+1}).$$

Interpretation:

1. Real rates are high when impatience (δ) is high.
2. Higher γ makes rates more sensitive to consumption growth.
3. The σ^2 term captures *precautionary savings.* When consumption is more volatile, people with this utility function worry more about low consumption states - they want to save, driving down interest rates.

   or: consumption growth is high when real rates are since people save now to spend later, and consumption is less sensitive to interest rates as the desire for smooth consumption, (captured by γ) rises.

What γ Does

With power utility, it controls

1. intertemporal substitution - aversion to consumption that varies over time
2. risk aversion - aversion to consumption that varies over states of nature
3. precautionary savings, which depends on the third derivative of the utility function

This link is particular to power utility.

Risk Corrections

Start with $p = E[mx]$ and recall that $cov(m, x) = E(mx) - E(m)E(x)$, so that

$$p = E(m)E(x) + cov(m, x).$$

Substituting in $R^f = 1/E(m)$,

$$p = \frac{E(x)}{R^f} + cov(m, x)$$

The first term is the standard present value formula. The second term is a risk adjustment.

- An asset whose payoff covaries positively with the sdf has its price raised (and return lowered) and vice versa.

- To understand this, substitute back in terms of consumption:

$$p = \frac{E(x)}{R^f} + \frac{cov(\beta u'(c_{t+1}), x_{t+1})}{u'(c_t)}$$

Marginal utility declines as c rises. So, an asset's price falls as its payoff covaries more positively with consumption and vice versa.

- Why? Investors don't like uncertainty about consumption.
- If an asset's payoff covarys positively with consumption, it pays off when you feel wealthy, i.e., when marginal utility is low. That's not worth so much.
- If an asset's payoff covarys negatively with consumption, it is like insurance, and you pay more for it.

Why Only Covariance Matters

- That is, why *doesn't* variance affect asset prices?
- This is a version of the diversification argument.
- If an agent can maintain steady consumption, individual asset volatility is irrelevant.
- To see why, consider an investor who buys a little more ξ of x. The variance of consumption is

$$\sigma^2(c + \xi x) = \sigma^2(c) + 2\xi cov(c, x) + \xi^2\sigma^2(x)$$

Now, ξ is small, so the third term is an order of magnitude smaller than the second.

With Returns

You get the same thing.

- Start with $1 = E(mR^i)$.
- Then, using the covariance decomposition:

$$1 = E[m]E[R^i] + cov(m, R^i)$$

- Using the risk-free rate:

$$E[R^i] - R^f = -R^f cov(m, R^i)$$

or

$$E[R^i] - R^f = -\frac{cov[u'(c_{t+1}), R^i_{t+1}]}{E[u'(c_{t+1})]}.$$

- All assets have an expected return equal to the risk-free rate plus a risk adjustment.
- Assets whose returns positively covary with consumption increase consumption volatility and must provide higher expected return.
- Assets that covary negatively with consumption, such as insurance can offer even *negative* net returns. Think insurance.

Expected Return and Beta

- Start with
$$E[R^i] - R^f = -R^f cov(m, R^i)$$
and write it as

$$\begin{aligned} E[R^i] &= R^f - R^f cov(m, R^i) && (7) \\ &= R^f - \frac{1}{E[m]} cov(m, R^i) && (8) \\ &= R^f \left(\frac{cov(m, R^i)}{Var[m]}\right)\left(-\frac{Var[m]}{E[m]}\right) && (9) \\ &\equiv R^f \beta_{i,m} \lambda_m && (10) \end{aligned}$$

where $\beta_{i,m}$ is the regression coefficient of return R^i on m. This is a *beta pricing model*, although it is not the CAPM.

- Expected return is proportional to the regression coefficient, or beta in a regression of that return on the discount factor, m.
- You can think of λ as the market price of risk — it depends critically on the volatility of the discount factor (which is important for thinking about the equity premium puzzle.)

Power Utility

$$m = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma}$$

We can take a Taylor Series approximation of equation (7) to express betas in terms of consumption growth, which you can measure:

$$E[R^i] = R^f + \beta_{i,\Delta c} \lambda_{\Delta c}$$

where

$$\lambda_{\Delta c} = \gamma var(\Delta c)$$

Interpretation

- Expected returns should increase linearly with their betas on consumption growth.
- The factor risk premium is determined by risk aversion and the volatility of consumption. The more risk averse people are, or the riskier their environment, the bigger the risk premium.

Mean-variance Frontier

Asset pricing theory has focused a lot on this. First we show (Hansen Jagannathan 1991) that all assets priced by m obey

$$|E[R^i] - R^f| \leq \frac{\sigma(m)}{E[m]}\sigma(R^i) \tag{11}$$

This follows from:

$$1 = E[mR^i] = E[m]E[R^i] + \rho_{m,R^i}\sigma(R^i)\sigma(m)$$

and hence,

$$E[R^i] = R^f - \rho_{m,R^i}\frac{\sigma(m)}{E[m]}\sigma(R^i)$$

Interpretations

1. Means and variances lie in the wedge from figure 1.1.
2. All returns on the frontier are perfectly correlated with the discount factor: $|\rho_{m,R^i}| = 1$. Returns on the upper (lower) part are perfectly negatively correlated with the discount factor, and positively (negatively) correlated with consumption.
3. All frontier returns are perfectly correlated with each other (obviously). So any frontier return can be synthesized from two others. For example, pick a frontier return R^m and then:

$$R^{mv} = R^f + a(R^m - R^f)$$

for some a.

4. Because each point on the frontier is also perfectly correlated with the discount factor, we must be able to pick a, b, c, d such that

$$m = a + bR^{mv}$$

and

$$R^{mv} = d + em$$

so that any MVE return carries all pricing information. If we have a MVE return, we can find a discount factor and vice versa.

In particular, if the CAPM holds, and R^{mv} is the market portfolio, then the discount factor is linear in the market portfolio. More on this later in the book.

5. Given a discount factor, we can also construct a single beta representation so expected returns can be described in a single

beta representation using any MVE return except R^f:

$$E[R^i] = R^f + \beta_{i,mv}(E[R^{mv}] - R^f).$$

That is, in mean/beta space, you should get a straight line — with the CAPM, it is the security market line.

Because the beta model applies to every return, if one plugs in R^{mv} and note it has a beta of 1, we see that the market risk premium is $R^{mv} - R^f$.

This means that there must be a relationship between discount factors, beta models, and mean variance frontiers (see chapter 6 of Cochrane's book).

6. Note that figure 1.1 can be used to break apart return into systematic, and idiosyncratic. The idiosyncratic part is not priced.

The Equity Premium Puzzle

The ratio of mean excess return to standard deviation is called the Sharpe ratio:

$$sr \equiv \frac{E[R^i] - R^f}{\sigma(R^i)}$$

This is an interesting number for MBA's. Why?

- The highest Sharpe Ratio available is the slope of the mean-standard deviation frontier.
- Let R^{mv} be on the frontier. recall equation (11)

$$|E[R^i] - R^f| \leq \frac{\sigma(m)}{E[m]}\sigma(R^i)$$

This implies that the slope of the frontier is

$$\left|\frac{E[R^i] - R^f}{\sigma[R^{mv}]}\right| = \frac{\sigma(m)}{E[m]} = \sigma(m)R^f$$

The slope of the frontier is governed by the volatility of the discount factor.

- Now, add power utility: $u'(c) = c^{-\gamma}$.

$$\left|\frac{E[R^i] - R^f}{\sigma[R^{mv}]}\right| = \frac{\sigma[(c_{t+1}/c_t)^{-\gamma}]}{E[(c_{t+1}/c_t)^{-\gamma}]}.$$

- The s.d. of the RHS is big when consumption is volatile or γ is large. Suppose further that consumption growth is lognormal. You will be asked to show that

$$\left|\frac{E[R^i] - R^f}{\sigma[R^{mv}]}\right| = \sqrt{exp(\gamma^2\sigma^2(\Delta lnc_{t+1})) - 1} \sim \gamma\sigma(\Delta lnc).$$

Interpretation

The slope of the frontier is higher if consumption is more volatile or investors are more risk averse.

Unfortunately, in postwar U.S. data, the slope of the historical frontier is "too big." Real returns have averaged 9% with a s.d. of 16% and the real T-bill return is about 1%.

Worse still, NDS consumption growth has a s.d. of about 1% so that to reconcile these facts with the above equation requires a γ of about 50. Either people are much more risk averse than we thought, or stock returns for the past 50 years were luck, or the model is wrong. We'll look at this again later on in the course (CLM chapter 8 and C chapter 21)

Time Varying Expected Returns

There is some confusion about what the efficient markets hypothesis means. It does not mean is that prices are unpredictable — quite.

This is important for tests of predictability in returns (CLM chapters 1,2). Predictability does not necessarily mean that markets are inefficient.

To see why, consider again, the basic FOC from our model

$$p_t u'(c_t) = E_t[\beta u'(c_{t+1})(P_{t+1} + d_{t+1})] \tag{12}$$

for a dividend paying stock.

martingales

- If (1) investors are risk-neutral, u' is linear, or (2) if there is no variation in consumption, or (3) there are no dividends, and further, we're dealing with short horizons so β is close to 1, this becomes
$$p_t \sim E_t[p_{t+1}]$$
- Prices then follow a time series:
$$p_{t+1} = p_t + \epsilon_{t+1}$$
- If $\sigma_t^2(\epsilon_{t+1} \forall t$ is constant, prices follow a random walk.
- More generally, prices follow a martingale.

- Put another way, expected returns, $E[p_{p+1}/p_t]$ should be unpredictable.
- Our general equation (12) says that prices should follow a martingale after adjusting for dividends and scaling by marginal utility.
- What people do is this scaling to take advantage of useful mathematical properties of a risk-neutral world — this is risk neutral pricing.
- Since consumption and risk aversion don't change much on a daily basis, short horizon returns should be pretty much unpredictable and technical analysis should not "work."
- For longer horizons, however, predictability is certainly possible — it could be that money is really to be made, or it could be that we simply have a conditional asset pricing model.

Example

Start with

$$
\begin{aligned}
E[R_{t+1} - R^f] &= -\frac{cov_t(m_{t+1}, R_{t+1})}{E_t[m_{t+1}]} \qquad (13) \\
&= \frac{\sigma_t(m_{t+1})}{E_t[m_{t+1}]}\sigma_t(R_{t+1})\rho_t(m_{t+1}, R_{t+1}) \\
&\sim \gamma_t \sigma_t(\Delta c_{t+1})\sigma_t(R_{t+1})\rho_t(m_{t+1}, R_{t+1})
\end{aligned}
$$

Interpretation

- The t emphasizes that these are conditional moments.
- The conditional mean and the unconditional mean of a r.v. need not be the same: knowing who is injured improves forecasting who will win a football game.
- Looking at equation 13 we see that returns can be predictable. If the conditional variance changes, we'd expect the conditional mean to as well.
- But letting the mean return move in and out along a line with constant Sharpe Ratio doesn't help much in the data — we have to look at conditional covariances as well or look at volatility of log consumption growth or changes in risk

aversion. Models that make these connections are a very active area of research in asset pricing.

2 Period Models

It's worth looking in the text to show that the solution to the 2 period problem is general in the sense that we get the same type of asset pricing formula with an infinite horizon.

A good place to look at a paper that solves everything out is Lucas (1978), which you saw in the first asset pricing course.

Assignment 1

Cochrane chapter 1: 1., 5., 6.

Asset Pricing Review II

Cochrane ch. 2: Using the Model.

So far, we've imposed no restrictions. To do any tests, we have to impose some. What we've not assumed is

- There is a representative investor — our equations apply individually to everyone. But if you want to use aggregate consumption data, you need a representative agent.
- Payoff and return distributions are anything. But to do econometrics, we're going too have to make assumptions about what's stationary, what's ergodic, ...
- Two periods.
- i.i.d. returns over time (see above)

- time– or state–separable utility (it just looks that way) from the exposition.
- investors have no other sources of income (e.g. labor, real estate, ...). People need this for the CAPM, but we've not assumed it here.

What People Did

Although in principle, the consumption–based asset pricing model answers all questions in the universe, once you make any assumptions and try to do empirical work, you find that it doesn't work.

This is what people did in the early 1980's:

- Assume power utility for a representative agent:

$$u'(c) = c^{-\gamma}$$

- Write down the equation governing excess returns in terms of aggregate consumption:

$$0 = E_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} R^e_{t+1} \right]$$

- Then, what people did was say that the pricing kernel was stationary over time, and that the only way conditioning information entered the estimation equation was to make the one-step-ahead forecast errors orthogonal to past information. Then, using current and past returns, they estimated γ and β and *blew the model out of the water.*
- What you could do instead is take unconditional expectations (getting rid of all the stuff in the agent's information set), rewrite our equation as

$$E[R^e_{t+1}] = -R^r cov \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\gamma}, R^e_{t+1} \right]$$

take some aggregate consumption data and some returns and again estimate γ and β and *blew the model out of the water.* See figure 2.4 in Cochrane's book. The best fit has a $\gamma = 241$! Experimental values of γ are about 2.

- It is an understatement to say that the model does not perform well.

So, what to do?

The first thing to do is recognize that aggregate consumption data is not very good.

The second thing to do is recognize that even if consumption data were any good, then the models are awful. Why, because aggregate consumption data (and hence the pricing operator in our simple model) is not that variable and returns are. So *of course* we either need better models or better data or a way of writing down models that don't need consumption data to directly explain anything.

Modern asset pricing theory tries every one of these strategies.

For Instance

- Different utility functions: nonseparabilities across time or goods; habit persistence.
- Individual consumption data (aggregation can lead the variance of aggregate wealth to appear in the pricing kernel (making it more variable).
- Models that allow consumption to be written in terms ofsomethingg better measurable (e.g. Campbell 1993).
- Ad-hoc factor models (for the discount factor, for example. We can show that the CAPM falls into this class

$$m_{t+1} = a + bR^W_{t+1}$$

m is linear in the return on the market portfolio. In turn, this has sparked a resurgence in the conditional CAPM (which has better measures of aggregate wealth than the NYSE value-weighted index. See Jagannathan and Wang 1996).

Asset Pricing review III

Cochrane chapter 5 (or CLM chapter 5). We will focus on sections 5.1 and 5.2: the Lagrangian Characterization. Skip the other stuff.