ECON 2723, Asset Pricing, Section 1

Jennifer Walsh[1]

September 15, 2023

[1] Thanks to Alex Wu, Marcela Mello, Roman Sigalov, and Argyris Tsiaras for their help.

Admin

- Problem set is due on Thursday, October 5. Let me know if you have any questions!
- I will hold office hours on Fridays, 9:30-10:30am (weeks when we don't have section).
- Always feel free to email/contact me!

ECON 2723 Topics

1. Chapter 1: Risk
2. Chapter 2: Mean-variance analysis
3. Chapter 3: CAPM, APT
4. Chapter 3: Factor models
5. Chapter 4: SDF
6. Chapter 5: Interpreting the level of returns (cash flows or discount rates)
7. Chapter 6: Consumption based Asset Pricing
8. Chapter 7: Production based Asset Pricing
9. Chapter 8: Fixed Income (bonds)
10. Chapter 9: Intertemporal models
11. Chapter 10: Household finance
12. Chapter 11, 12: Incomplete markets, asymmetric information, market microstructure

Contents

CARA-Normal

Suppose we invest *dollar amount* θ in an asset whose returns satisfy

$$\tilde{x} \sim \mathcal{N}(\mu, \sigma^2).$$

- Investor solves

$$\max_{\theta} \ \mathbb{E}\big[- \exp\big(- A(W_0 + \theta \tilde{x})\big)\big].$$

- Consider the moment generating function for a normal random variable,

$$\mathbb{E}\big\{ \exp\big[\mathcal{N}(\mu, \sigma^2)\big]\big\} = \exp\left(\mu + \frac{1}{2}\sigma^2\right),$$

- or the properties of the log-normal distribution,

$$\log\big[\mathbb{E}(\tilde{z})\big] = \mathbb{E}\big[\log(\tilde{z})\big] + \frac{1}{2} Var\big[\log(\tilde{z})\big].$$

CARA-Normal (2)

- Rewrite the objective function as

$$\max_{\theta} \; -\exp\left\{\mathbb{E}\big[-A(W_0+\theta\tilde{x})\big]+\frac{1}{2}Var\Big[-A(W_0+\theta\tilde{x})\Big]\right\} \iff$$

$$\max_{\theta} \; -\exp\left\{-A\big[W_0+\theta\mathbb{E}(\tilde{x})\big]+\frac{1}{2}A^2\theta^2 Var(\tilde{x})\right\} \iff$$

$$\max_{\theta} \; -\exp\left\{-A\Big[\theta\mathbb{E}(\tilde{x})-\frac{1}{2}A\theta^2 Var(\tilde{x})\Big]\right\} \iff$$

$$\max_{\theta} \left\{\theta\mathbb{E}(\tilde{x})-\frac{1}{2}A\theta^2 Var(\tilde{x})\right\} \iff$$

$$\max_{\theta} \left\{\theta\mu-\frac{1}{2}A\theta^2\sigma^2\right\}.$$

- The first order condition yields

$$\mu - A\theta^*\sigma^2 = 0 \implies \theta^* = \frac{\mu}{A\sigma^2},$$

where θ is in dollar terms.

CARA-Normal (3)

Its main advantage is to be very tractable...

- **Easy to aggregate**. For example, in Chapter 12 going to cover Rational Expectations Equilibrium. With CARA

$$\text{Total supply} = \sum_{i \in \text{Agents}} \theta_i^* = \sum_{i \in \text{Agents}} \frac{\mu_i - P}{A\sigma_i^2}$$

  where it is easy to solve for price

- Can incorporate idiosyncratic labor risk as in Chapter 10

$$\max_{\theta} \mathbb{E}[-\exp(-A(W_0 + \tilde{Y} + \theta\tilde{x}))] \text{ where } \tilde{Y} \sim \mathcal{N}(\mu_Y, \sigma_Y^2)$$

- Can analyze belief disagreement, market design and speculation as in Simsek (2013) in a very tractable form.

- Very easy to extend to multiple assets. The maximization problem reduces to mean variance

$$\max_{\theta} \left\{ \theta'\mu - \frac{A}{2}\theta'\Sigma\theta \right\} \overset{FOC}{\Longrightarrow} \mu = A\Sigma\theta \implies \theta = \frac{1}{A}\Sigma^{-1}\mu$$

CARA-Normal (4)

But there are important disadvantages.

- Wealth doesn't affect the amount invested in the risky asset
- Utility is bounded above but returns are unbounded below
- Growth in consumption and wealth with multiplicative risks implies increasing absolute risks (so there should be an upward trend in risk premia)
- The normality assumption cannot hold over more than one time interval, because the compounding of returns over many periods converts a normal distribution into a right-skewed, non-normal distribution.

CAPM with CARA-Normal (1)

We can use the CARA-Normal model to derive the CAPM in a 2-period setting.

- Suppose an investor invests dollar amount θ_i in each risky asset i and θ^f into the risk free asset.
- Then, from the budget constraint

$$\theta^f + \theta'\mathbf{1} = W_0 \implies \theta^f = W_0 - \theta'\mathbf{1}.$$

- Consumption in period 1 is

$$C_1 = W_1 = W_0 + \theta^f R^F + \theta' R.$$

- Returns are normally distributed

$$R \sim \mathcal{N}\Big(\mathbb{E}(R), \Sigma\Big).$$

- The objective function is

$$\max_{\theta^f, \theta} \; -\mathbb{E}\big[e^{-AC_1}\big] \iff \max_{\theta} \left\{ (W_0 - \theta'\mathbf{1})R^F + \theta'\mathbb{E}(R) - \frac{1}{2}A\theta'\Sigma\theta \right\}.$$

CAPM with CARA-Normal (2)

- The FOC is

$$\mathbb{E}(R) - R^f = A \cdot \Sigma \cdot \theta$$

- Note that $R^W = \theta^F R^F + \theta' R$, so that

$$Cov(R, R^W) = Cov(R, \theta' R) = \Sigma \cdot \theta.$$

- Therefore, FOC can be rewritten as

$$\mathbb{E}(R) - R^F = A \cdot Cov(R, R^W) = \frac{Cov(R, R^W)}{Var(R^W)} \cdot A \cdot Var(R^W).$$

- Evaluate this expression at the wealth portfolio:

$$\mathbb{E}(R^W) - R^F = A \cdot Var(R^W).$$

- If all investors are identical, then $R^W = R^M \implies$ CAPM:

$$\mathbb{E}(R) - R^F = \frac{Cov(R, R^M)}{Var(R^M)} \cdot \left[\mathbb{E}(R^M) - R^F\right] = \beta' \left[\mathbb{E}(R^M) - R^F\right].$$

CRRA-Lognormal (1)

We are going to work with the CRRA-Lognormal A LOT. So let's cover it in detail.

- Optimization problem

$$\max \ \mathbb{E}_t \left[\frac{W_{t+1}^{1-\gamma}}{1-\gamma} \right], \text{ where } W_{t+1} = W_t(1 + R_{p,t+1}),$$

- and $\log\left(R_{p,t+1}\right) = r_{p,t+1} \sim \mathcal{N}\left(\mathbb{E}_t r_{p,t+1}, \sigma^2_{p,t}\right)$.

- Assume (w.l.g.) that $\gamma < 1$, drop $1-\gamma$, and take the log of the objective function:

$$\max \ \log\left(\mathbb{E}_t\left[W_{t+1}^{1-\gamma}\right]\right).$$

CRRA-Lognormal (2)

- Let $w_{t+1} \equiv \log(W_{t+1})$. Lognormal properties imply the objective function

$$\max \left\{ \mathbb{E}_t \left[(1-\gamma) w_{t+1} \right] + \frac{1}{2}(1-\gamma)^2 Var(w_{t+1}) \right\}.$$

- Take logs of the budget constraint

$$W_{t+1} = W_t(1 + R_{p,t+1}) \implies w_{t+1} = w_t + r_{p,t+1}.$$

- Plug it into the maximization problem, cancel $1-\gamma$, and drop w_t since it is a constant

$$\max \left\{ \mathbb{E}_t \left[r_{p,t+1} \right] + \frac{1}{2}(1-\gamma) Var(r_{p,t+1}) \right\}$$

CRRA-Lognormal (3)

- What's up with the variance?

$$\max \left\{ \mathbb{E}_t \left[r_{p,t+1} \right] + \frac{1}{2}(1-\gamma) Var(r_{p,t+1}) \right\}$$

  Split the expression

$$\max \underbrace{\mathbb{E}_t \left[r_{p,t+1} \right] + \frac{1}{2} Var(r_{p,t+1})}_{\log \mathbb{E}[1+R_{p,t+1}]} - \frac{\gamma}{2} Var(r_{p,t+1})$$

- What is the effect of increased variance? When we fix geometric average, increase in variance raises mean arithmetic return. For an aggressive investor, with $\gamma < 1$, this is good news.

- This is going to reappear in Chapter 6: increase in variance for a conservative investor, $\gamma > 1$, is perceived as a deterioration in investment opportunities $\Longrightarrow$ asset prices go down.

CRRA-Lognormal (4)

- Before solving the problem, we need to express the *excess portfolio return* in terms of the *return* and *portfolio weight* on the *risky asset*, and the *risk-free rate*.

$$
\begin{aligned}
r_{p,t+1} - r_{f,t+1} &= \log(1 + R_{p,t+1}) - \log(1 + R_{f,t+1}) \\
&= \log[\alpha_t(1 + R_{t+1}) + (1 - \alpha_t)(1 + R_{f,t+1})] - \log(1 + R_{f,t+1}) \\
&= \log\left[1 + \alpha_t\left(\frac{1 + R_{t+1}}{1 + R_{f,t+1}} - 1\right)\right] \\
&= \log\left[1 + \alpha_t\left(\exp(r_{t+1} - r_{f,t+1}) - 1\right)\right]
\end{aligned}
$$

CRRA-Lognormal (5)

- A 2^{nd}-order Taylor expansion of function $f(x) = \log\left[1+\alpha\left(\exp(x)-1\right)\right]$ around $x=0$ gives

$$\begin{aligned}
\log\left[1+\alpha\left(\exp(x)-1\right)\right] =& 0 + \left.\frac{\alpha\exp(x)}{1+\alpha\left(\exp(x)-1\right)}\right|_{x=0} x \\
&+ \left.\frac{\alpha\exp(x)\left[1+\alpha\left(\exp(x)-1\right)\right]-\alpha^2\exp(2x)}{\left[1+\alpha\left(\exp(x)-1\right)\right]^2}\right|_{x=0}\frac{x^2}{2} \\
&+ \overline{o}(x^2) \\
=& \alpha x + \frac{1}{2}(\alpha-\alpha^2)x^2 + \overline{o}(x^2)
\end{aligned}$$

- Applying this to $r_{p,t+1} - r_{f,t+1}$ we get

$$r_{p,t+1} - r_{f,t+1} \approx \alpha(r_{t+1} - r_{f,t+1}) + \frac{1}{2}\alpha(1-\alpha)\underbrace{(r_{t+1} - r_{f,t+1})^2}_{\to\sigma^2 \text{ as time interval shrinks}}$$

A side note on continuous time approximation (credits to Argyris Tsiaras)

- Non-trivial result from stochastic calculus: *the quadratic variation of a Brownian motion is equal to its variance*. Here goes a heuristic argument...

- A Brownian motion is defined by
$$z_{t+\Delta t} - z_t \sim \mathcal{N}(0, \Delta t),$$
which implies the (heuristic) formulas
$$\mathbb{E}_t[dz_t] = 0,$$
$$Var_t[dz_t] = \mathbb{E}_t\left[(dz_t)^2\right] = dt.$$

- A key property of the Brownian motion is that the latter equality also holds **ex post**: $(dz_t)^2 = dt$.

A side note on continuous time approximation (credits to Argyris Tsiaras)

- A heuristic argument is the following:

$$\frac{(z_{t+\Delta t} - z_t)}{\sqrt{\Delta t}} \sim \mathcal{N}(0,1) \implies \frac{(z_{t+\Delta t} - z_t)^2}{\Delta t} \sim \chi_1^2$$

- Therefore,

$$\mathbb{E}_t\left[(z_{t+\Delta t} - z_t)^2\right] = \Delta t \cdot \mathbb{E}\left[\chi_1^2\right] = \Delta t,$$
$$Var_t\left[(z_{t+\Delta t} - z_t)^2\right] = (\Delta t)^2 \cdot Var\left[\chi_1^2\right] = 2(\Delta t)^2.$$

- So $Var_t\left[(z_{t+\Delta t} - z_t)^2\right]$ vanishes at the order of dt^2 and, hence, can be **ignored** as the time interval shrinks.

- It follows that $(z_{t+\Delta t} - z_t)^2$ is **instantaneously deterministic**: $(z_{t+\Delta t} - z_t)^2 = dt$.

A side note on continuous time approximation (credits to Argyris Tsiaras)

- Going back to our problem, $r_{t+\Delta t} - r_{f,t+\Delta t}$ is an arithmetic Browninan motion *conditional on information at time* t, so, by the same argument,

$$
\begin{aligned}
(r_{t+dt} - r_{f,t+dt})^2 =& \mathbb{E}_t(r_{t+dt} - r_{f,t+dt})^2 \\
=& \mathbb{E}_t(r_{t+dt} - \mathbb{E}_t r_{t+dt} + \mathbb{E}_t r_{t+dt} - r_{f,t+dt})^2 \\
\overset{[r_{f,t+dt} \text{ known at time } t]}{=}& \mathbb{E}_t(r_{t+dt} - \mathbb{E}_t r_{t+dt})^2 + (\underbrace{\mathbb{E}_t r_{t+dt} - r_{f,t+dt}}_{\text{Risk Premium}})^2 \\
\overset{[\text{risk premium instantaneously deterministic}]}{=}& \sigma_t^2 dt.
\end{aligned}
$$

CRRA-Lognormal (6)

- Investor's (approximate) problem is then

$$\begin{cases} \max_{\alpha_t} \left\{ \mathbb{E}_t \left[r_{p,t+1} \right] + \frac{1}{2}(1-\gamma) Var(r_{p,t+1}) \right\} \\ r_{p,t+1} \approx r_{f,t+1} + \alpha_t (r_{t+1} - r_{f,t+1}) + \frac{1}{2}\alpha_t (1-\alpha_t)\sigma_t^2 \end{cases}$$

- Plug the approximation to get

$$\max_{\alpha_t} \left\{ r_{f,t+1} + \alpha_t (\mathbb{E}_t r_{t+1} - r_{f,t+1}) + \frac{1}{2}\alpha_t \sigma_t^2 - \frac{1}{2}\gamma \alpha_t^2 \sigma_t^2 \right\}$$

- The first order condition is

$$\mathbb{E}_t r_{t+1} - r_{f,t+1} + \frac{1}{2}\sigma_t^2 = \gamma \alpha_t \sigma_t^2$$

$$\Longleftrightarrow$$

$$\alpha_t = \frac{\mathbb{E}_t r_{t+1} - r_{f,t+1} + \frac{1}{2}\sigma_t^2}{\gamma \sigma_t^2} \approx \frac{\mathbb{E}_t R_{t+1} - R_{f,t+1}}{\gamma \sigma_t^2}.$$

Growth Optimal Portfolio (1)

- When $\gamma = 1$, the CRRA utility reduces to *log-utility*. The investor, thus, chooses the portfolio that maximizes log-returns, called the **Growth Optimal Portfolio**:

$$\max \ \mathbb{E}_t[r_{p,t+1}]$$

- If returns are *iid*, this portfolio *outperforms every other portfolio* with increasing probability as the investment horizon increases.
  - To see this, note that

    $$r^{GO}_{t+1} - r_{p,t+1} \sim \mathcal{N}(\Delta, \sigma^2),$$

    where $\Delta > 0$ because the growth optimal portfolio maximizes log-returns.
  - Given *iid* returns, compounding implies that

    $$r^{GO}_{t+k} - r_{p,t+k} \sim \mathcal{N}(k\Delta, k\sigma^2).$$

  - Let $\Phi(\cdot)$ be the CDF of the standard normal distribution, then

    $$\mathbb{P}(r^{GO}_{t+k} - r_{p,t+k} < 0) = \Phi\left(\frac{0 - k\Delta}{\sqrt{k}\sigma}\right) = \Phi\left(-\sqrt{k}\frac{\Delta}{\sigma}\right) \to_{k\to\infty} 0.$$

Growth Optimal Portfolio (2)

- Opposing forces balance each other when $\gamma = 1$:
  1. When $\gamma > 1$, the investor seeks a **safer** portfolio by *penalizing the variance of log-returns*.
  2. When $\gamma < 1$, the investor seeks a **riskier** portfolio since a *higher variance*, conditional on a given mean log-return, *corresponds to higher mean simple return*.

- We are going to see GO portfolio several times in Chapter 4:
  1. The return on GO portfolio will be related to inverse of SDF.
  2. It will provide lower bound on the entropy of SDF.

Contents

Arbitrage Pricing Theory (1)

- The main idea of the APT can be summarized in the following points
  - There is a big common component to stock returns.
  - Idiosyncratic movements in asset returns should not carry any risk prices, since investors can diversify away by holding portfolios.
  - Thus, risk prices or expected returns should depend only on the asset's exposure to the common risk sources.
  - In particular, if many assets are available ($N \to \infty$), we should expect α's to be very small in absolute value.

- This is a statistical argument which contrasts it to CAPM. Let's formalize this...

- Suppose that asset i's excess return is given by the **market model**

$$R_{it}^e = \alpha_i + \beta_{im} R_{mt}^e + \varepsilon_{it}, \text{ with } \mathbb{E}_t[\varepsilon_{it}\varepsilon_{jt}] = 0.$$

- Consider forming a *well diversified*[2] portfolio of such assets with weights w. Portfolio's return is then

$$R_{pt}^e = \underbrace{\sum_{i=1}^{N} w_i \alpha_i}_{\alpha_p} + \underbrace{\left(\sum_{i=1}^{N} w_i \beta_{im} \right)}_{\beta_{pm}} R_{mt}^e + \underbrace{\sum_{i=1}^{N} w_i \varepsilon_{it}}_{\varepsilon_p}$$

[2] Small enough weight in each asset, so that the variance of the portfolio's idiosyncratic component is negligible.

Arbitrage Pricing Theory (2)

- Suppose that $\alpha_p > 0$.

- Now suppose we do the following trade:
  - Go long \$1 worth of this portfolio $\implies$ "get" $\alpha_p + \beta_{pm} R^e_{mt} + \varepsilon_p$
  - Go short \$$\beta_{pm}$ of market $\implies$ "lose" $\beta_{pm} R^e_{mt}$
  - This means that our net average gain is $\alpha_p + \varepsilon_p$

- What is the Sharpe ratio on our portfolio? In a simple case when $w_i = \frac{1}{N}$ (but in general when each w_i is "not too large") the variance of portfolio residual is

$$Var_t(\varepsilon_{pt}) = \frac{\sigma^2}{N} \to 0$$

  Hence the Sharpe ratio is

$$S = \frac{\mathbb{E}[\alpha_p + \varepsilon_p]}{\sqrt{Var_t(\varepsilon_{pt})}} = \frac{\alpha_p}{\sigma/\sqrt{N}} \to \infty$$

- This is too attractive. Investors will come in to close this asymptotic arbitrage and drive α_p to zero.

APT vs CAPM

- APT's result is weaker than CAPM's: we only have a limiting result that nonzero alphas are rare in the first, while all alphas should exact zero in the second.
- Recall the CARA-Normal CAPM derivation from earlier. There we got

$$\mathbb{E}(R^i) - R^F = \beta_{iW}\left[\mathbb{E}(R^W) - R^f\right]$$

  In principle can price any asset (will return to this in Chapter 4). In a sense it is *absolute pricing* (pricing securities by reference to fundamental sources of risk). But it is highly model dependent.
- APT, in contrast, gives the *price of the asset relative to other prices*. But it doesn't specify, for example, what is the return on the market.
  - If we assume that the market is priced correctly, APT allows us to price other securities relative to the market.
  - If we assume that the market model holds, APT asymptotically yields the CAPM formula without requiring the restrictive and unrealistic assumptions of the latter.
  - The asymptotic approximation of an APT model holds well for (i) portfolios that naturally have high R^2,[3] or (ii) well-diversified portfolios in large markets.

[3] Factor model is defined as a regression: $Var(R_i^e) = Var(\beta_i' f) + Var(\varepsilon_i) \implies Var(\varepsilon_i)/Var(R_i^e) = 1 - R^2$.

Conditional CAPM

- Our CAPM so far has had a constant beta over time

$$\mathrm{E}R^e_{i,t+1} = \beta_{im}\mathrm{E}R^e_{m,t+1} \tag{1}$$

- What if stock betas vary over time? Can this explain asset price movements?
- Suppose that CAPM holds conditionally and take the unconditional expectation

$$\mathbb{E}_t[R^e_{i,t+1}] = \beta_{imt}\mathbb{E}_t[R^e_{m,t+1}] \stackrel{\text{uncond.}}{\Longrightarrow}$$

$$\mathbb{E}[R^e_{i,t+1}] = \mathbb{E}[\beta_{imt}]\mathbb{E}[R_{m,t+1}] + Cov\left(\beta_{imt}, \mathbb{E}_t[R_{m,t+1}]\right) \iff$$

$$\mathbb{E}[R^e_{i,t+1}] = \underbrace{\beta_{im}\mathbb{E}[R^e_{m,t+1}]}_{\text{Uncond. CAPM}} + \underbrace{\left(\mathbb{E}[\beta_{imt}] - \beta_{im}\right)}_{\text{Volatility Timing}}\mathbb{E}[R^e_{m,t+1}] + \underbrace{Cov\left(\beta_{imt}, \mathbb{E}_t[R^e_{m,t+1}]\right)}_{\text{Market Timing}}$$

- In the second line, we add and subtract the unconditional beta times the unconditional equity premium (that is, the unconditional CAPM's prediction of the asset's unconditional expected excess return).
- The first term is the unconditional CAPM term, and the second and third terms represent deviations from those.

Conditional CAPM

$$\mathbb{E}[R^e_{i,t+1}] = \underbrace{\beta_{im}\mathbb{E}[R^e_{m,t+1}]}_{\text{Uncond. CAPM}} + \underbrace{\left(\mathbb{E}[\beta_{imt}] - \beta_{im}\right)}_{\text{Volatility Timing}}\mathbb{E}[R^e_{m,t+1}] + \underbrace{Cov\left(\beta_{imt}, \mathbb{E}_t[R^e_{m,t+1}]\right)}_{\text{Market Timing}}$$

- The second term is positive (i.e. an asset can have higher returns than what is predicted by the conditional CAPM) if the time-series average of the asset's conditional beta is higher than it's unconditional beta. This can happen if the market return is heteroskedastic.
- Lewellen and Nagel (2006) show that one can approximate the difference $\mathrm{E}\beta_{imt} - \beta_{im}$ as:

$$\mathrm{E}\beta_{imt} - \beta_{im} \approx -\frac{\mathrm{Cov}\left(\beta_{imt}, \sigma^2_{mt}\right)}{\sigma^2_m} \tag{2}$$

- The unconditional beta (averaged across time periods) is a time series regression coefficient which places high weight on periods with large variance of market returns (of either sign).
- If an assets tends to have a high conditional beta when market volatility is low, then the time-series average of its conditional beta will be higher than its unconditional beta. This is a high-beta stock when markets are calm, and a low-beta stock when markets are volatile.
- The third term is positive if the beta covaries positively with the equity premium. If beta is high when the market risk premium is high, the asset delivers market exposure at times when market risk is highly rewarded, which increases the unconditional average return.

Example of Market Timing

- Suppose $R_{m,t+1} = \mu_t + \varepsilon_{t+1}$, and that $\mu_t = R_f - \Delta$ or $\mu_t = R_f + \Delta$ with equal probability.
- Suppose there is a strategy that invests 100% in the risk free asset when $\mu_t = R_f - \Delta$ and 100% in the market when $\mu_t = R_f + \Delta$.
- Unconditional return on the market is $\mathbb{E}[R_{mt}] = 0.5(R_f + \Delta) + 0.5(R_f - \Delta) = R_f$. This means that the expected market excess return is zero. Therefore, unconditional CAPM will predict that every expected excess return is zero.
- However, the unconditional expected return of this strategy is

$$E[R_t] = \underbrace{0.5R_f}_{\text{when } \mu_t = R_f - \Delta} + \underbrace{0.5(R_f + \Delta)}_{\text{when } \mu_t = R_f + \Delta} = R_f + 0.5 \cdot \Delta > R_f$$

- The discrepancy comes from the covariance term: when expected market return is low the beta of this strategy is 0 and when expected market return is high the beta of this strategy is 1.

Can Market Timing Explain Excess Alpha?

- Lewellen and Nagel (2006) note that:

$$\mathrm{Cov}\left(\beta_{imt}, E_t R^e_{m,t+1}\right) = \sigma\left(\beta_{imt}\right)\sigma\left(E_t R^e_{m,t+1}\right)\rho_{\beta_{imt}, E_t R^e_{m,t+1}} \tag{3}$$

- We can estimate the standard deviation of beta and the risk premium, and set the correlation to 1 to get an upper bound on how much the stock's unconditional alpha can be above the conditional CAPM prediction.
- Estimate $\sigma_\beta = 0.3$ and $\sigma\left(E_t R^e_{m,t+1}\right) = 0.5\%$.
- Then, unconditional α is at most 0.15% if β and γ are perfectly correlated.
- Empirically, the B/M strategy has an alpha of 0.59% monthly (std. error, 0.14%), and a momentum strategy has an alpha of 1.01% monthly (std. error, 0.28%).

Conditional CAPM vs APT (inspiration for factors)

- Where should our inspiration for factors come from in the conditional CAPM framework?
  - In a one period model we care only about consumption next period. Hence, the relevant covariances are with return on wealth portfolio.
  - In multiperiod models this will change. Suppose there are some state variable that predict return, i.e. high X_t indicates that future returns are high. This increases our value function today. Hence, such variables can also be factors.
  - This is just to make you think and we will go back to state variables later in the class.
- In contrast, the APT framework suggests that one start with a statistical analysis of the covariance matrix of returns and find portfolios that characterize common comovement.

Testing Multifactor Pricing Models: Step 1, Time Series Regression

- The β_{ik}'s are defined to be the slope coefficients in the *time-series regression* of excess returns on all K factors:

$$R_{it}^e = a_i + \beta_{i1} f_{1t} + ... + \beta_{iK} f_{Kt} + \varepsilon_{it}, \; t = 1, ..., T. \tag{4}$$

- Your asset pricing model constrains the values of β_{ik}. Consequently, the time series regression gives the model falsifiability.

- Hence, estimating the β_{ik}'s (running time-series regressions) is a **necessary** step in *any* test of a factor pricing model.

- The question is whether running the time-series regression for each asset (or portfolio) i is **sufficient** to test a factor pricing model.

Testing Multifactor Pricing Models: Step 1, Time Series Regression

- The β_{ik}'s are defined to be the slope coefficients in the *time-series regression* of excess returns on all K factors:

$$R^e_{it} = a_i + \beta_{i1} f_{1t} + ... + \beta_{iK} f_{Kt} + \varepsilon_{it},\ t = 1, ..., T. \tag{4}$$

- Your asset pricing model constrains the values of β_{ik}. Consequently, the time series regression gives the model falsifiability.

- Hence, estimating the β_{ik}'s (running time-series regressions) is a **necessary** step in *any* test of a factor pricing model.

- The question is whether running the time-series regression for each asset (or portfolio) i is **sufficient** to test a factor pricing model.

- **Question:** What is/are the factor(s) in the (not conditional) CAPM from class yesterday?

Testing Multifactor Pricing Models: Step 2

- A factor pricing model with factors f_1, f_2, ..., f_K posits that

$$\mathbb{E}[R_i^e] = \sum_{k=1}^{K} \beta_{ik} \lambda_k, \tag{5}$$

for all assets $i = 1, ..., N$.

- The expected return is the price of risk times the quantity of risk.
  - R_i^e is the excess return of asset i over the the risk-free rate or zero-beta rate.[4]
  - The λ_k's are the factor risk **prices**.
    - Interpretation: λ_k is the price of (compensation for) the exposure to the risk associated with factor k.
  - The β_{ik}'s are **NOT** free parameters in the model.
    - They are determined by the time series of an assets' return (last slide).
    - Interpretation: β_{ik} measures the (**quantity** of) exposure of an asset i to the risk associated with factor k.

[4] Defined as the expected return on a portfolio that is uncorrelated with all the factor portfolios.

When factors are also excess returns...

- Note:

$$E[R_{it}^e|f_t] = \beta_i' f_t = Cov(R_i^e, f')Var^{-1}(f_t)f_t$$
$$E[E[R_{it}^e|f_t]] = E[R_{it}^e] = \beta_i E[f_t]$$

- Assume $E[f_k] = E[R_k^e]$ for all periods t.
- Apply the definition in (4) to factor k's excess return. It follows that:
  - $\beta_{kk} = 1$: factor k has a beta of 1 with itself.
  - $\beta_{kj} = 0$ for all factors $j \in \{1, \ldots, K\}$, $j \neq k$: factor k has a beta of zero with all other factors.
  - $a_k = 0$.
- Take the expectation of (4) for $i = k$ and compare it with the defining equation of the model, in (5):

$$\beta_{kk} = 1,\ \beta_{kj} = 0,\ a_k = 0 \implies \lambda_k = \mathbb{E}[R_k^e].$$

  - The price of (compensation for) factor k's risk *equals* the expected excess return on factor k.This pins down the free parameters λ_k, $k = 1, \ldots, K$.
- Now take the expectation of (4) for any asset i that is not a factor. The finding above implies:

$$\mathbb{E}[R_i^e] = a_i + \sum_{k=1}^{K} \beta_{ik}\lambda_k.$$

  - The model in (5) then implies $a_i = 0$ for all i.
  - Therefore, we can test that model by testing the null that all intercepts of the times-series regressions are zero.

When factors are not returns...

- When factors are not excess returns (for example, inflation or GDP growth), we **cannot** pin down the free parameters λ_k **as the expected values** of factors.

- To see this, take the expectation of the time-series in (4) and combine it with the model in (5):

$$\mathbb{E}[R_{it}^e] = a_i + \sum_{k=1}^{K} \beta_{ik}\mathbb{E}[f_{kt}] \overset{(5)}{\Longrightarrow} a_i = \sum_{k=1}^{K} \beta_{ik}\left(\lambda_k - \mathbb{E}[f_k]\right).$$

- Clearly, however, this intercept restriction cannot be tested without an estimate for λ_k.
  - This is what the cross-sectional approach solves: it estimates the λ_k's from a cross-sectional regression of average excess returns on estimated betas (next slide).

Cross-Sectional Regression

- Two-step procedure:
  1. Estimate β_{ik}'s from the time-series regressions in (4).
  2. Regress average excess returns on estimated betas in the cross-section of assets:

$$\mathbb{E}[R_{it}^e] = \sum_{k=1}^{K} \beta_{ik}\lambda_k + \alpha_i, \text{ for } i = 1, \ldots, N,$$

  and where the α_i's are the pricing errors (cross-sectional residuals) we are interested in testing.[5]

- The model in (5) predicts that $\alpha_i = 0$ for all i, and we can test this prediction.

[5] This regression can be ran either with or without an intercept. If we want to imposed the null as far as possible, we should run a no-intercept regression.

Time-Series vs Cross-Sectional Approach (when factors are excess returns)

- When factors are excess returns, we have the option of testing the model either directly through the time-series approach or through the cross-sectional regression, since $\alpha_i = a_i = 0$ for all i under the null.
- But the time-series approach and the OLS cross-sectional regressions are **not** equivalent.
  - The time-series regressions sets
    $$\hat{\lambda}_k = \frac{1}{T}\sum_{t=1}^{T} f_{kt} \text{ (sample mean)},$$
    which forces the estimate for the pricing error, $\hat{a}_k$, to be *always* zero.
  - The OLS cross-sectional estimates, however, the λ_k's so as to yield the best fit for all assets in an equally weighted regression. So it is possible that $\hat{\alpha}_k \neq 0$.
- The time-series approach, however, **is equivalent** to a GLS cross-sectional regression that includes the factor portfolios as test assets.

Other Approaches

- Fama MacBeth:
  1. Run time series regression for every asset i to estimate β_i
  2. Then, run cross-sectional regression for every period to estimate λ_t, α_t.

     $$R^e_{it} = \boldsymbol{\beta}_i' \lambda_t + \alpha_{it}$$

     for $t = 1, \ldots, T$.
  3. Define

     $$\widehat{\lambda} = \frac{1}{T} \sum_{t=1}^{T} \widehat{\lambda}_t \text{ and } \widehat{\boldsymbol{\alpha}} = \frac{1}{T} \sum_{t=1}^{T} \widehat{\boldsymbol{\alpha}}_t$$

     and do inference on the alphas.
- Equivalent to pooled time-series cross-sectional OLS with standard errors clustered by time in special case where explanatory variable doesn't vary over time
- Less efficient than panel model, gives equal weight to each time period