

Empirical Asset Pricing

Francesco Franzoni

Swiss Finance Institute - University of Lugano

Testing Asset Pricing Models: Overview

1. Testing CAPM: historical development
2. Statistical description of the asset pricing tests
3. Stochastic Discount Factor Representation and Tests
4. Conditional Asset Pricing Models
5. Statistical discipline

Relevant readings:

- Cochrane: chapters 6, 7, 8, 12, 14, 15, 16
- Huang and Litzenberger: chapter 10

1. Testing CAPM: historical development

Quick Review: The CAPM

- The Sharpe-Lintner-Mossin (SLM) version of the model assumes that the risk-free rate is available. For any asset i :

$$E(R_i) = R_f + \beta_i (E(R_m) - R_f) \quad (1)$$

where R_m is the total wealth (i.e., market) portfolio and

$$\beta_i = \frac{\text{Cov}(R_i, R_m)}{\text{Var}(R_m)} \quad (2)$$

- The model states that the **market return is mean-variance efficient**
- That is: the expected return on any asset can be spanned by the expected return on the market and the risk free rate
- The Black (1972) version of the model is derived without assuming the existence of the risk-free rate:

$$E(R_i) = E(R_0) + \beta_i (E(R_m) - E(R_0)) \quad (3)$$

where R_0 is the return on a mean-variance efficient portfolio with zero covariance with R_m . It turns out that $E(R_m) - E(R_0) > 0$

SLM version:

1. **Time-series test.** Consider the time-series regression

$$R_{it} - R_f = \alpha_i + \beta_i (R_{mt} - R_f) + \varepsilon_{it}.$$

Take expectations of both sides. The model predicts

$$\alpha_i = 0 \quad \forall i = 1 \dots N$$

2. **Cross-sectional test.** Consider the cross-sectional regression

$$E_T(R_i^e) = a + \beta_i \lambda + d_i \gamma + u_i \quad i = 1 \dots N \quad (4)$$

$E_T(R_i^e)$: time-series average excess return (it is an estimate of the expected excess return)

d_i : vector of **stock characteristics** different from beta (e.g. size, dividend yield, volatility, etc.):

- (a) $\lambda = E(R_m) - R_f > 0$
- (b) $\alpha_i \equiv a + u_i = 0 \quad \forall i$ (zero pricing errors)
- (c) $\gamma = 0$

Conceptual Issues in the Tests

Even before worrying about econometrics, testing CAPM raises three fundamental conceptual hurdles

1. **Ex ante vs. sample moments.** The model predictions are stated in terms of *ex ante* expected returns and betas but tested on sample moments

- The solution is to assume Rational Expectations: investors' beliefs coincide with the true distribution of returns. This means forecast errors $R_{i,t+1} - E_t(R_{i,t+1})$ are mean zero, so sample averages of realized returns converge to the *ex ante* expected returns the model predicts

2. **Non-stationarity.** CAPM holds period by period: it is a conditional model. How to handle non-stationarity in the moments of the *ex ante* distribution?

- One solution is to assume that CAPM holds unconditionally too.
Implications of this assumption:
The conditional CAPM:

$$E_t(R_{i,t+1}^e) = \beta_{it} E_t(R_{m,t+1}^e)$$
$$\beta_{it} = \frac{Cov_t(R_{i,t+1}, R_{m,t+1})}{Var_t(R_{m,t+1})}$$

Take unconditional expectations of both sides

$$\begin{aligned} E \left(E_t \left(R_{i,t+1}^e \right) \right) &= E \left(\beta_{it} E_t \left(R_{m,t+1}^e \right) \right) \\ E \left(R_{i,t+1}^e \right) &= Cov \left(\beta_{it}, E_t \left(R_{m,t+1}^e \right) \right) \\ &\quad + E \left(\beta_{it} \right) E \left(R_{m,t+1}^e \right) \end{aligned}$$

For the model to hold unconditionally, one needs to assume that

(a) $Cov \left(\beta_{it}, E_t \left(R_{m,t+1}^e \right) \right) = 0$

(b) $E \left(\beta_{it} \right) = \beta_i = \frac{Cov(R_i, R_m)}{Var(R_m)}$

- This is not necessarily true. For example:
 - Beta can change along the business cycle
 - Beta can change along the firm's life-cycle
- The other solution is to model the conditional moments. We will discuss this strategy later — and see that even conditional models face a fundamental testability problem (Hansen and Richard critique). For now, the focus is on unconditional tests

3. **Roll's (1977) critique:** the total wealth portfolio includes non-marketable assets (e.g. human capital, private businesses, etc.). Instead, the tests typically focus on a proxy of the market that is traded (e.g. the S&P 500, the CRSP index, etc.)
- Consequence: measurement error in the market return.
Rejection of CAPM may depend on use of incorrect market portfolio
 - Most tests ignore the unobservability and assume proxy is mean-variance efficient
 - Also: if the true market portfolio is sufficiently correlated with the proxy (above 70%), a rejection of the proxy implies a rejection of the true portfolio (Stambaugh (1982), Kandal and Stambaugh (1987), Shanken (1987))
 - Some authors have tried to compute a broader proxy by including the return on human capital (Jagannathan and Wang (1996))

Econometric Issues in the Tests

1. **Heteroskedastic and correlated errors.** The errors in equation (4) are likely to be heteroskedastic and correlated across assets
 - OLS coefficients are unbiased but inefficient (OLS is not BLUE, GLS is)
 - OLS standard errors are biased
 - **Possible Solutions:**
 - (a) Use feasible GLS...I have never seen this in practice
 - (b) Use OLS to estimate the slopes and adjust the standard errors with an estimate of the correct variance-covariance matrix
 - (c) Fama and MacBeth (1973) procedure


2. **Errors-in-variables.** The betas in equation (4) are measured with error because they are estimates of true betas

- Measurement error causes $\hat{\gamma}$ to be biased towards zero and CAPM rejected
- **Main Solution:** Group stocks into portfolios. Averaging over L stocks reduces measurement error variance proportionally to $1/L$, so estimates become consistent as $L \rightarrow \infty$
- However, grouping also compresses beta dispersion (it reduces between-group variance), reducing the efficiency of cross-sectional estimates
 - See Gagliardini, Ossola, and Scaillet (2016) for a stock-level approach to estimating risk premia that maximizes efficiency
- **Best grouping practice:** sort by pre-ranking betas estimated from non-overlapping data. This maximizes between-group beta dispersion while keeping measurement errors uncorrelated over time

3. **Time-varying betas.** The stock-level betas are likely to change over time

- Betas change over the life-cycle in a non-stationary way
- Historical solution: form portfolios according to a stationary characteristic

- Does it work? See Franzoni (2002): Where's beta going?



4. **High idiosyncratic volatility.** High volatility of individual stock returns

- For individual stocks, cannot reject hypothesis that average returns are all the same
- St. error of mean is $\frac{\sigma_i}{T^{1/2}}$. If σ_i between 40% and 80% and $T = 60$ months, very large confidence intervals
- **Solution:** form portfolios

5. **Non-synchronous trading.**

- Small stocks react slowly to information
- Solution: Scholes and Williams (1977)

$$\hat{\beta} = \frac{\hat{\beta}^+ + \hat{\beta} + \hat{\beta}^-}{1 + 2\rho}$$

where $\hat{\beta}^+$, $\hat{\beta}^-$, and $\hat{\beta}$ come from the simple regression of stock returns on leads, lags, and the contemporaneous market return:

$$R_{it}^e = \alpha_i^+ + \beta_i^+ R_{m,t+1} + \varepsilon_t^+$$

$$R_{it}^e = \alpha_i + \beta_i R_{m,t} + \varepsilon_t$$

$$R_{it}^e = \alpha_i^- + \beta_i^- R_{m,t-1} + \varepsilon_t^-$$

ρ is the 1st order auto-correlation of the market return.

Note that if $\rho = 1$, it is as if you were running three times the same regression. Then, you have to divide the sum of the betas by $1 + 2\rho = 3$

If $\rho = 0$, each of the three regressions provides independent information on the reaction of stock returns to market returns. Then, the sum of the three betas does not need to be normalized.

- Another solution is due to Dimson (1979)

$$\hat{\beta} = \hat{\beta}^+ + \hat{\beta} + \hat{\beta}^-$$

where the estimates come from

$$R_{it}^e = \alpha_i + \beta_i^+ R_{m,t+1} + \beta_i R_{m,t} + \beta_i^- R_{m,t-1} + \varepsilon_t$$

2. Statistical description of AP tests

- General focus on multi-factor asset pricing models

$$E(R_i^e) = \beta_i' \lambda \quad (5)$$

where β_i is a K -dimensional vector of multiple-regression slopes and λ is K -vector of factor risk premia

- For example: APT, ICAPM
- CAPM is the case with $K = 1$
- The methodology extends unambiguously from the CAPM tests to tests of multifactor models
- Questions we address:
 - **How to estimate parameters?**
 - **Standard errors?**
 - **How to test the model predictions?**

Estimation Approach: Time-Series Regressions

- You can apply this approach **only if factors** $f_1, \dots, f_K$ **are returns**

- Then, the model (5) applies to factors as well:

$$E(f_k) = \lambda_k \quad k = 1 \dots K \quad (6)$$

- So, one can re-write the model in (5) as:

$$E(R_i^e) = \beta_i' E(f) \quad (7)$$

where $E(f)$ is the vector of **expected excess returns** on the k factors

- Assume $K = 1$ for simplicity, consider the time-series regression:

$$R_{it}^e = \alpha_i + \beta_i f_t + \varepsilon_t \quad (8)$$


- Take expectation of each side. Then, the implication of (7) is

$$\alpha_i = 0 \quad i = 1 \dots N$$

- You can estimate α_i by running regression (8) for each asset

Estimating the Model Predictions

- Then, use t -tests to test $\alpha_i = 0$
- In an $E(R)$ - β space, you are pricing the factor correctly by fitting a line through the factor – i.e., the factor has zero alpha – and allowing pricing errors on the other assets



- You want to test the hypothesis that all alphas are jointly zero

- The ε_{it} are correlated across assets with variance-covariance matrix $\Sigma = E(\varepsilon_t \varepsilon_t')$
- The asymptotically valid test for

$$H_o : \alpha = 0$$

is

$$\hat{\alpha}' \left[\widehat{Var}(\hat{\alpha}) \right]^{-1} \hat{\alpha} \sim \chi_N^2$$

where α is a N -dimensional vector of pricing errors

- Intuition: reject the model if the weighted sum of the squared errors is far from zero
- In the case $K = 1$, assuming no autocorrelation and stationarity, and noting that $\hat{\alpha}_i$ contains the average ε_{it} , the test becomes

$$T \left[1 + \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2 \right]^{-1} \hat{\alpha}' \hat{\Sigma}^{-1} \hat{\alpha} \sim \chi_N^2$$

where $E_T(f)$ is the sample mean of the factor (an estimate of the risk premium) and $\hat{\Sigma}$ is the sample vcov matrix of the residuals from the N regressions

Proof. (Sketch): $\widehat{Var}(\hat{\alpha})$ is $N \times N$ matrix

$$\widehat{Var}(\hat{\alpha}) = \begin{bmatrix} \hat{\sigma}_{\hat{\alpha}_1}^2 & \hat{\sigma}_{\hat{\alpha}_1\hat{\alpha}_2} & \cdots & \hat{\sigma}_{\hat{\alpha}_1\hat{\alpha}_N} \\ \hat{\sigma}_{\hat{\alpha}_2\hat{\alpha}_1} & \hat{\sigma}_{\hat{\alpha}_2}^2 & & \\ \vdots & & \ddots & \\ \hat{\sigma}_{\hat{\alpha}_N\hat{\alpha}_1} & & & \hat{\sigma}_{\hat{\alpha}_N}^2 \end{bmatrix}$$

Focus on:

$$\begin{aligned} \hat{\alpha}_i &= \bar{R}_i^e - \hat{\beta}_i E_T(f) \\ &= \alpha_i + \beta_i E_T(f) + \bar{\varepsilon}_i - \hat{\beta}_i E_T(f) \\ &= \alpha_i + (\beta_i - \hat{\beta}_i) E_T(f) + \bar{\varepsilon}_i \end{aligned}$$

So:

$$\begin{aligned}
Var(\hat{\alpha}_i) &= Var(\hat{\beta}_i) E_T^2(f) + Var(\bar{\varepsilon}_i) \\
&= \frac{\sigma_{\varepsilon_i}^2}{\hat{\sigma}^2(f) T} E_T^2(f) + \frac{1}{T} \sigma_{\varepsilon_i}^2 \\
&= \frac{1}{T} \sigma_{\varepsilon_i}^2 \left[1 + \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2 \right]
\end{aligned}$$

Similarly:

$$Cov(\hat{\alpha}_i, \hat{\alpha}_j) = \frac{1}{T} \sigma_{\varepsilon_i \varepsilon_j} \left[1 + \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2 \right]$$

Hence, it follows that:

$$Var(\hat{\alpha}) = \frac{1}{T} \left[1 + \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2 \right] \Sigma$$

So:

$$\widehat{Var}(\hat{\alpha}) = \frac{1}{T} \left[1 + \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2 \right] \hat{\Sigma}$$

Then:

$$\hat{\alpha}' \left[\widehat{Var}(\hat{\alpha}) \right]^{-1} \hat{\alpha} = T \left[1 + \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2 \right]^{-1} \hat{\alpha}' \hat{\Sigma}^{-1} \hat{\alpha}$$

You only need to prove that the distribution is a chi-squared. For that you can use standard theorems on the limit distribution of squared residuals. ■

The Gibbons, Ross, and Shanken (1989) Test

- Assuming joint normality of the ε_{it} , Gibbons, Ross, and Shanken (GRS, 1989) provide a small sample test that has an intuitive geometric interpretation

$$\frac{T - N - 1}{N} \left[1 + \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2 \right]^{-1} \hat{\alpha}' \hat{\Sigma}^{-1} \hat{\alpha} \sim F_{N, T-N-1}$$

- Using efficient set algebra, one can prove that


$$\left(\frac{\mu_q}{\sigma_q} \right)^2 = \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2 + \hat{\alpha}' \hat{\Sigma}^{-1} \hat{\alpha} \quad (9)$$

where $\frac{\mu_q}{\sigma_q}$ is the Sharpe ratio of the tangency portfolio and $\frac{E_T(f)}{\hat{\sigma}(f)}$ is the Sharpe ratio of the factor

- Then, the GRS statistic can be rewritten as

$$\frac{T - N - 1}{N} \frac{\left(\mu_q / \sigma_q \right)^2 - \left(E_T(f) / \hat{\sigma}(f) \right)^2}{1 + \left(\frac{E_T(f)}{\hat{\sigma}(f)} \right)^2} \sim F_{N, T-N-1}$$

- In this form, the interpretation is: **Reject the null if the factor is far from the tangency portfolio on the ex-post efficient frontier**
- In other words: reject CAPM if the market is far from the ex-post efficient frontier



- In case of $K > 1$ factors, the GRS statistic is

$$\frac{T - N - K}{N} \frac{\left(\frac{\mu_q}{\sigma_q}\right)^2 - E_T(f)' \hat{\Omega}^{-1} E_T(f)}{1 + E_T(f)' \hat{\Omega}^{-1} E_T(f)} \sim F_{N, T-N-K}$$

where

$$\hat{\Omega} = \frac{1}{T} \sum_{t=1}^T [f_t - E_T(f)] [f_t - E_T(f)]'$$

is the sample vcov matrix of the factors and

$$E_T(f)' \hat{\Omega}^{-1} E_T(f)$$

is the multifactor equivalent of the squared Sharpe ratio of the factor

- **Empirical application:** GRS rejects CAPM and the Fama-French three-factor model on the 25 size/BM portfolios, though with a much smaller statistic for the latter

Proof of Equation (9)

- Let q be a portfolio on the mean variance efficient frontier. Let μ be the vector of expected excess returns on the $N + 1$ assets in the market (i.e. N test assets and 1 factor). Hence, q has minimum variance for given expected return e

$$\begin{aligned} \text{Min} \quad & q'Vq \\ \text{s.t. } & q'\mu = e \end{aligned}$$

- The Lagrangean for this problem is

$$L = q'Vq + 2\lambda(e - q'\mu)$$

which gives first order conditions with respect to q and λ (the Lagrange multiplier)

$$\begin{aligned} 2Vq - 2\lambda\mu &= 0 \\ e - q'\mu &= 0 \end{aligned}$$

- Hence, the frontier portfolio q has to satisfy

$$q = \lambda V^{-1}\mu \tag{10}$$

- Using Equation (10), the variance of q is therefore

$$\sigma_q^2 = q'Vq = \lambda^2 \mu'V^{-1}\mu$$

- The squared expected excess return of q is

$$\mu_q^2 = (q'\mu)^2 = \lambda^2 (\mu'V^{-1}\mu)^2$$

- We can then compute the squared Sharpe ratio of q as

$$\frac{\mu_q^2}{\sigma_q^2} = \mu' V^{-1} \mu \quad (11)$$

- Now, let us remember that this market is composed of $N + 1$ assets, where the first asset is the factor f
- The vector of returns of the N assets can be written as

$$R_t = \alpha + \beta f_t + \varepsilon_t,$$

where α and β are N -dimensional vectors capturing the alphas and betas of the N assets with respect to the factor f . Also, ε is the idiosyncratic component of returns, which has variance-covariance matrix equal to Σ

- Hence, the vector μ can be written as

$$\mu = \begin{bmatrix} \mu_f \\ \alpha + \beta \mu_f \end{bmatrix} \quad (12)$$

- For the same reason, the variance-covariance matrix V can be written as

$$V = \begin{bmatrix} \sigma_f^2 & \beta' \sigma_f^2 \\ \beta \sigma_f^2 & \beta \beta' \sigma_f^2 + \Sigma \end{bmatrix} \quad (13)$$

- Using the formula for the inverse of a partitioned matrix (see, e.g., Greene, Econometric Analysis), we can obtain the inverse of V as

$$V^{-1} = \begin{bmatrix} \frac{1}{\sigma_f^2} + \beta' \Sigma^{-1} \beta & -\beta' \Sigma^{-1} \\ -\Sigma^{-1} \beta & \Sigma^{-1} \end{bmatrix} \quad (14)$$

- Using Equations (12) and (14), after some tedious algebra, we can re-write Equation (11) as

$$\frac{\mu_q^2}{\sigma_q^2} = \frac{\mu_f^2}{\sigma_f^2} + \alpha' \Sigma^{-1} \alpha,$$

which completes the proof

Estimation Approach: Cross-Sectional Regressions

- **This is the only approach available when the factors are not returns.** E.g., oil shocks, interest rate shocks, etc.

- Two-pass methodology:

1. Estimate β_i from time-series regressions with multiple factors if $K > 1$


$$R_{it}^e = a_i + \beta_i' f_t + \varepsilon_{it} \quad t = 1 \dots T \quad (15)$$

2. Regress average returns on estimated β_i

$$E_T(R_{it}^e) = \beta_i' \lambda + \alpha_i \quad i = 1 \dots N \quad (16)$$

where E_T denotes a time-series average over the T observations in the sample.

- In this case, you are fitting a line through the expected returns, trying to minimize pricing errors, estimating a risk premium for the factor, and allowing a non-zero pricing error for the factor



Standard Errors

- To test the significance of the $\hat{\lambda}$, you could use the estimate of the variance that takes into account error correlation

$$\sigma^2 (\hat{\lambda}_{ols}) = \frac{1}{T} \left[(\beta' \beta)^{-1} \beta' \Sigma \beta (\beta' \beta)^{-1} \right]$$

- Using standard OLS results, we can compute the vcov of $\hat{\alpha}$ as

$$Var(\hat{\alpha}) = \frac{1}{T} \left(I - \beta (\beta' \beta)^{-1} \beta' \right) \Sigma \left(I - \beta (\beta' \beta)^{-1} \beta' \right)$$

- Specifically, to derive the formula use the fact that $Var(\hat{u}) = MVar(u)M$, where $\hat{u}$ are the fitted residuals and $M = I - X(X'X)^{-1}X'$ is the residual making matrix; in our case $u = \alpha$ and $X = \beta$
- Notice that these formulas assume that β **are known**. See next for corrections that take into account that β is estimated

Estimating the Model Predictions

- You can test the AP model with

$$\hat{\alpha}' [Var(\hat{\alpha})]^{-1} \hat{\alpha} \sim \chi^2_{N-K}$$

Shanken Correction

- In fact, β_i are estimates, not the true parameters
- Hence, need to account for sampling error in β_i , when computing standard errors
- **Shanken's (1992) correction:**

$$\sigma^2(\hat{\lambda}_{ols}) = \frac{1}{T} \left[(\beta' \beta)^{-1} \beta' \Sigma \beta (\beta' \beta)^{-1} (1 + \lambda' \Omega^{-1} \lambda) \right] + \frac{1}{T} \Omega \quad (17)$$

where Ω is the vcov of the factors

- You have a **multiplicative** $(1 + \lambda' \Omega^{-1} \lambda)$ correction and an **additive** correction $\frac{1}{T} \Omega$
- Is the multiplicative correction important?
- Consider CAPM. In annual data: $(\lambda_m / \sigma_m)^2 = (0.08 / 0.16)^2 = 0.25$
The correction is important.
In monthly data $(\lambda_m / \sigma_m)^2 = 0.25 / 12 \approx 0.02$. The multiplicative correction is not important

- The additive correction is likely to be important
 - $\frac{1}{T}\Omega$ is the variance of the mean of the factor, which is non-negligible in case f is R_m
 - **Numerical assessment.** Example with $\beta = \mathbf{1}_N$, $\Sigma = \sigma_\varepsilon^2 I_N$, the base term (ignoring the multiplicative correction) simplifies to $\sigma_\varepsilon^2/(NT)$ and the additive term is σ_m^2/T . Their ratio is $N\sigma_m^2/\sigma_\varepsilon^2$. For the 25 Fama and French (1993) portfolios in monthly data ($\sigma_m \approx 4.5\%$, $\sigma_\varepsilon \approx 2.7\%$):

$$\frac{N\sigma_m^2}{\sigma_\varepsilon^2} = \frac{25 \times (0.045)^2}{(0.027)^2} \approx 69$$

The additive term is ≈ 70 times larger than the base term and **dominates the total variance of $\hat{\lambda}$**

- If the whole system, including betas, risk premia, and pricing errors is **estimated via GMM**, the correction of the standard errors is made “automatically” (see Cochrane, ch. 12). No need for Shanken correction

Time-Series vs. Cross-Section

- How are the two approaches different?
1. **TS can be applied only if factors are returns**
 - In CAPM, it fits a line through the pricing errors by forcing a zero pricing error on R_m
 - Test: $\alpha_i = 0$ for all i
 2. **CS only alternative when factors are not returns**
 - In CAPM, it minimizes all pricing errors, by allowing some error on R_m
 - Historically CS has been **used to test for characteristics in the cross-section of returns** (especially using Fama-MacBeth approach)

Fama-MacBeth (1973) Methodology

- The FM procedure addresses both time-varying betas and cross-sectionally correlated errors by running the cross-sectional regression separately at each date t , then averaging the resulting estimates
- It is a **three-pass procedure**:
 1. **First pass**: Obtain β_{it} from time-series regressions, using data up to $t - 1$
 2. **Second pass**: At each date t , run a cross sectional regression

$$R_{it}^e = \beta_{it}' \lambda_t + d_{it}' \gamma_t + \alpha_{it} \quad i = 1 \dots N$$

and obtain a time series of $\hat{\lambda}_t$, $\hat{\gamma}_t$, and $\hat{\alpha}_{it}$ $t = 1 \dots T$. Note that d_{it} are cross-sectional characteristics such as profitability, size, B/M, etc.

3. **Third pass:** Finally, obtain full-sample estimates as time series means

$$\hat{\lambda} = \frac{1}{T} \sum_{t=1}^T \hat{\lambda}_t$$

$$\hat{\gamma} = \frac{1}{T} \sum_{t=1}^T \hat{\gamma}_t$$

$$\hat{\alpha}_i = \frac{1}{T} \sum_{t=1}^T \hat{\alpha}_{it}$$

and use the standard error of the mean

$$\sigma(\hat{\lambda}) = \frac{1}{T^{1/2}} \left(\underbrace{\frac{\sum_{t=1}^T (\hat{\lambda}_t - \hat{\lambda})^2}{T-1}}_{\widehat{Var}(\hat{\lambda}_t)} \right)^{1/2}$$
$$\sigma(\hat{\alpha}_i) = \frac{1}{T^{1/2}} \left(\underbrace{\frac{\sum_{t=1}^T (\hat{\alpha}_{it} - \hat{\alpha}_i)^2}{T-1}}_{\widehat{Var}(\hat{\alpha}_{it})} \right)^{1/2}$$

and similarly you can obtain the standard error for $\hat{\gamma}_t$

- **Two main advantages** of this approach:

1. It allows for **time-varying betas**
2. It computes **standard errors** by getting around the problem of heteroskedastic and correlated errors: it exploits sample variation in $\hat{\lambda}_t$, $\hat{\gamma}_t$, and $\hat{\alpha}_{it}$

3. SDF Representation and Tests

Equivalence of Representations

- In general one can always go from a **Stochastic Discount Factor (SDF)** representation

$$\begin{aligned} E(R_{it}m_t) &= 1 \\ m_t &= a - bf_t \end{aligned} \tag{18}$$

where m_t is the SDF expressed as a linear function of K factors f_t , **to an expected return-beta representation**

$$\begin{aligned} E(R_{it}) &= \lambda_0 + \lambda'\beta \\ \beta &= Cov(R_{it}, f_t) E(f_t f_t)^{-1} \end{aligned}$$

and vice versa

- Here, we show the direction SDF $\longrightarrow$ E(R)- β , see Cochrane ch. 6 for the proof of the other direction

- Start from Equation (18), let $K = 1$ for simplicity, and replace m_t into the first equation

$$\begin{aligned}
E(R_{it}(a - bf_t)) &= 1 \\
aE(R_{it}) - bE(R_{it}f_t) &= 1 \\
aE(R_{it}) &= 1 + bE(R_{it}f_t) \\
aE(R_{it}) &= 1 + bE(R_{it})E(f_t) + bCov(R_{it}, f_t) \\
E(R_{it})(a - bE(f_t)) &= 1 + bCov(R_{it}, f_t)
\end{aligned} \tag{19}$$

Note that by pricing the risk-free rate you get

$$\begin{aligned}
R_f E(m) &= 1 \\
R_f &= \frac{1}{a - bE(f_t)}
\end{aligned}$$

So, by replacing R_f into equation (19), you get

$$E(R_{it}) = R_f + bR_fCov(R_{it}, f_t)$$

which can be expressed in E(R)- β form

$$\begin{aligned} E(R_{it}) &= \lambda_0 + \beta_{if}\lambda_1 \\ \beta_{if} &= \frac{Cov(R_{it}, f_t)}{Var(f_t)} \\ \lambda_1 &= bR_f Var(f_t) \\ \lambda_0 &= R_f \end{aligned} \tag{20}$$

Testing the SDF Representation

- The SDF representation can be tested using GMM on the moment conditions given by the pricing errors

$$E [R_{it}m_t(\delta)] = 1$$

where

$$m_t(\delta) = a - bf_t$$

with $\delta = \begin{bmatrix} a & b \end{bmatrix}$

- Given that we have N assets, define the pricing errors

$$w_t(\delta) = R_tm_t(\delta) - \mathbf{1}_N$$

where R_t is an N -vector of asset returns and $\mathbf{1}_N$ is an N -vector of ones

- The N moment conditions for the vector of pricing errors $w_t(\delta)$ are

$$E [w_t(\delta)] = 0$$

- The GMM estimates the parameters in δ so as to minimize a quadratic form in the N moment conditions

$$E[w_t(\delta)]' A E[w_t(\delta)] \quad (21)$$

where A is a weighting matrix

- Assume that $w_t(\delta)$ is i.i.d. over time
- In this case, the optimal weighting matrix by Hansen and Singleton (1982) reduces to

$$A = [Var(w_t(\delta))]^{-1}$$

which is also called **second-stage** weighting matrix, because you estimate it by computing pricing errors from a first stage in which the weighting matrix is typically the identity

- However, the optimal quadratic form in (21) cannot be used as a metric to make comparisons across different asset pricing models
 - If a model contains more noise, $Var(w_t(\delta))$ is larger

- Then the quadratic form is smaller just because of the noise, and not because of smaller pricing errors. This argument echos Fama and French's (1993) defense against the rejection of their model by the GRS test

- The solution proposed by Hansen and Jagannathan (1994) is to use

$$A = [E(R_t R_t)]^{-1}$$

which is the matrix of second moments of returns

- The advantage is that this matrix remains the same across different specifications of the AP model and allows comparisons of different models
- Also, the authors show that the square root of the resulting quadratic form, called the **Hansen-Jagannathan (H-J) distance**, is the pricing error of the most mispriced portfolio among the N assets by a given AP model
- They derive the distribution of this statistic, which is in general non-standard
- This approach is used in the Conditional CAPM application of Jagannathan and Wang (1996)

4. Conditional Asset Pricing

Conditional Asset Pricing (Cochrane ch. 8)

- In the case of the expected return-beta representation, we have seen above that a **conditional model does not imply an unconditional model with the same factor in general**
- **The same conclusion holds for the SDF representation**
- Consider, for example, the SDF for CAPM

$$m_{t+1} = a_t + b_t R_{t+1}^m$$

- By applying the pricing equation to R_{t+1}^m and $R_{f,t}$, it can be shown that

$$b_t = -\frac{E_t(R_{t+1}^m) - R_{f,t}}{R_{f,t} \text{Var}_t(R_{t+1}^m)}$$

- Thus, the parameters of the SDF should be considered as **time-varying** because they are conditional moments

- The pricing statement is also conditional on time t information

$$1 = E_t \left[\underbrace{\left(a_t + b_t R_{t+1}^m \right)}_{m_{t+1}} R_{it+1} \right]$$

- In general, this statement does not imply an unconditional pricing statement. Let's try. Take unconditional expectations of each side:

$$\begin{aligned} 1 &= E \left[a_t R_{it+1} + b_t R_{t+1}^m R_{it+1} \right] \\ &= E[a_t] E[R_{it+1}] + \text{Cov}(a_t, R_{it+1}) \\ &\quad + E[b_t] E[R_{t+1}^m R_{it+1}] + \text{Cov}(b_t, R_{t+1}^m R_{it+1}) \end{aligned}$$

- Thus, **you can get an unconditional model only if the covariance terms are zero:**

$$1 = E \left[\left(\underbrace{E[a_t]}_a + \underbrace{E[b_t]}_b R_{t+1}^m \right) R_{it+1} \right]$$

which is **not** in general the case

Hansen and Richard (1987) Critique

- Testing conditional models may not be enough to do things correctly
- To be sure to test the correct model, one needs to account for all the relevant conditioning information that investors use
- **One needs to know investors' information set**
- This is arguably **impossible**
- Hence, according to Hansen and Richard **conditional factor models are not testable!**
- It resonates with the Roll Critique

A partial solution: scaled factors

- Because investors' information changes over time, the parameters a_t and b_t in the SDF change
- One can try to model this time-variation by using a set of L instruments z_t , which are variables that plausibly enter investors' information set

$$\begin{aligned}a_t &= a'z_t \\ b_t &= b'z_t\end{aligned}$$

- Suppose, for example, 1 factor and 1 instrument

$$\begin{aligned}m_{t+1} &= a_t + b_t f_{t+1} \\ &= a_0 + a_1 z_t + (b_0 + b_1 z_t) f_{t+1} \\ &= a_0 + a_1 z_t + b_0 f_{t+1} + b_1 z_t f_{t+1}\end{aligned}$$

- In place of a conditional 1-factor model, one obtains an unconditional 3-factor model, with factors $(z_t, f_{t+1}, z_t f_{t+1})$ and fixed coefficients

- $z_t f_{t+1}$ are scaled factors
- In general, from K factors and L instruments one gets $(K + 1) \times (L + 1)$ unconditional factors (including a constant)
- This is a partial solution, because one could be omitting some relevant variable from z_t that is instead in investors' information set (Hansen and Richard Critique)

A (very brief) history of Conditional Asset Pricing

- For a period of time, conditional factor models seemed to provide a solution to explain the failure of CAPM
 - Jagannathan and Wang (1996, JF) used a conditional CAPM with human capital to explain the size anomaly
 - Lettau and Ludvigson (2001, JF) used a conditional CAPM to explain the book-to-market anomaly (see later for the anomalies)
- However, Lewellen and Nagel (2006) put an end to this literature by showing that the empirical success of conditional asset pricing models was based on false premises

- **Key insight:** conditional alphas estimated at high frequency are significantly different from zero, and the covariation between β_{it} and the equity premium γ_t is too small to explain unconditional anomalies. Here is the formal argument.
- In $E(R)$ - β representation, the conditional CAPM says

$$\begin{aligned} E_t(R_{it+1}^e) &= \beta_{it}\gamma_t \\ \beta_{it} &= \frac{Cov_t(R_{it+1}, R_{t+1}^m)}{Var_t(R_{t+1}^m)} \\ \gamma_t &= E_t(R_{t+1}^m) - R_{f,t} \end{aligned}$$

- Take unconditional expectations

$$\begin{aligned} E(R_{it+1}^e) &= E(\beta_{it}\gamma_t) \\ &= E(\beta_{it})E(\gamma_t) + Cov(\beta_{it}, \gamma_t) \\ &= \bar{\beta}_i\gamma + Cov(\beta_{it}, \gamma_t) \end{aligned} \tag{22}$$

- By definition, the unconditional alpha in a time-series regression has the following Plim

$$\text{Plim } \alpha_i^u = E(R_{it+1}^e) - \beta^u \gamma \quad (23)$$

where $\beta^u = \frac{\text{Cov}(R_{it+1}, R_{t+1}^m)}{\text{Var}(R_{t+1}^m)}$ is the unconditional beta and $\gamma = E(\gamma_t)$

- Replacing (22) into (23) gets a clear expression for the unconditional alpha

$$\alpha^u = \gamma(\bar{\beta} - \beta^u) + \text{Cov}(\beta_t, \gamma_t)$$

- Under mild assumptions, you have that $\bar{\beta} = \beta^u$. So, that

$$\alpha^u = \text{Cov}(\beta_t, \gamma_t)$$

- Lettau and Ludvigson's (2001) point is that β_t varies along the business cycle with the equity premium γ_t and this variation is enough to explain unconditional alphas
 - So, firms that have higher betas during recessions (when the equity premium is high)—i.e., value stocks—have higher unconditional alphas: the so-called value premium
- **Lewellen and Nagel (2006)** show that this covariation is just not large enough
- They compute conditional alphas directly using higher frequency data and show that they are **not zero**

What were previous papers missing?

- L&N argue that these papers only test the *qualitative* implications of the conditional models. They **do not test the constraints imposed by the theory**
- In other words, they treat some coefficients as **free parameters**, which allows them to get more explanatory power in the cross-sectional regressions
- Concretely: previous papers parametrize $\beta_{it} = \beta_i + \delta_i z_t$, which implies the cross-sectional restriction $c_1 = Cov(z_t, \gamma_t)$. But they treat c_1 as a free parameter, gaining an extra degree of freedom that inflates R^2 and reduces pricing errors artificially
- **Lewellen and Nagel show that, when it is imposed, the restriction is actually rejected in the data and the R-squared drops substantially**

5. Statistical Discipline

Statistical Discipline According to Cochrane

- From efficient set mathematics: Mean variance efficiency of a factor implies a $E(R) = \beta' E(R_q)$ representation, where β are betas on the factor and $E(R_q)$ is the mean return on the factor q , where q is a portfolio on the mean-variance frontier constructed from the (sample or population) moments of the N assets

Proof. Take a portfolio x of the N assets. You have that

$$\beta_{x,q} = \frac{\text{Cov}(R_x, R_q)}{\text{Var}(R_q)} = \frac{\text{Cov}(x'R, q'R)}{\text{Var}(q'R)} = \frac{x' \text{Var}(R) q}{q' \text{Var}(R) q} = \frac{x' V q}{q' V q} \quad (24)$$

where x is the vector of weights for the portfolio, q is the vector of weight for a mean-variance efficient portfolio, and V is the variance-covariance matrix of the N assets

To obtain a mean-variance efficient portfolio, it has to be the case that q solves the problem

$$\begin{aligned} & \text{Min } q' V q \\ \text{s.t. } & q' \mu = e \end{aligned}$$

where e is a fixed level of expected excess return and μ is the vector of expected excess returns for the N assets. The Lagrangean for the problem is

$$\mathcal{L} = q'Vq + 2\lambda (e - q'\mu)$$

and the first order conditions are

$$\begin{aligned} Vq - \lambda\mu &= 0 \\ q'\mu &= e \end{aligned}$$

from the first order conditions we obtain

$$Vq = \lambda\mu$$

which we can use in equation (24):

$$\begin{aligned} \beta_{x,q} &= \frac{x'Vq}{q'Vq} \\ &= \frac{\lambda x'\mu}{\lambda q'\mu} \end{aligned}$$

Rearranging and simplifying λ :

$$x'\mu = \beta_{x,q}q'\mu$$

or:

$$E(R_x) = \beta_{x,q} E(R_q)$$

■

- This theorem holds whatever distribution one uses to compute moments: objective, subjective, or ex-post (sample) distribution
- Hence: in the sample, there exists an $E(R) = \beta' E(R_q)$ representation based on sample moments
- That is: there is an ex-post mean-variance efficient portfolio that prices all assets
- **Danger of “Fishing for Factors”!** Finding the portfolio that in-sample works
- Limit to the ability of out-of-sample tests: international data dirty; need to wait for 30 years for a new U.S. sample
- **Discipline:** a factor needs to be economically motivated

- Since any set of returns has an in-sample mean-variance efficient portfolio, a significant factor proves nothing on its own. Economic theory is the only reliable guide to which factors deserve attention.
- The Stochastic Discount Factor derives from the first order conditions for intertemporal consumption

$$E_t \left(m_{t+1} R_{i,t+1}^e \right) = 0$$

where

$$m_{t+1} = \frac{u'(c_{t+1})}{u'(c_t)} = 1 - b' f_{t+1}$$

- Different AP models can be derived from different specifications of SDF
- In any case, sensible factors must be somehow related to consumption growth. E.g.: return on wealth portfolio (CAPM), state variables of hedging concern (ICAPM)
- It does not mean, however, that covariance of factors with consumption growth implies fully rational investors (see later: Kozak, Nagel, and Santosh, 2018)

Data Snooping/Data Mining

- Related to “fishing for factors”, but different problem. **Here the issue is statistical**
- If you search the data for significant relationships over and over again, you will find something (with 5% chance)
- In principle, you should test model on a **different data set, i.e. out-of-sample**
- Objectively difficult
- **Need to go to the data with economically founded models**
- The multiple testing problem — how to set significance thresholds when hundreds of factors have been tried — is covered in a later lecture