Portfolio Choice (3)

John Y. Campbell

Ec2723

September 2010

Outline

- The Capital Asset Pricing Model (CAPM)
  - The model in a nutshell
  - Asset pricing implications
  - The model without riskless borrowing or lending
- Econometrics of CAPM tests
  - Time-series approach
  - Cross-section approach
  - Fama-MacBeth approach
  - Common themes and empirical evidence
- Arbitrage pricing and factor models
  - Single-factor model
  - Multi-factor model
  - Conditional vs. unconditional CAPM

The CAPM: Assumptions

- All investors are price-takers.
- All investors care about returns measured over one period.
- There are no nontraded assets.
- Investors can borrow or lend at a given riskfree interest rate (Sharpe-Lintner version of the CAPM - this is relaxed in the Black version).
- Investors pay no taxes or transaction costs.
- All investors are mean-variance optimizers.
- All investors perceive the same means, variances, and covariances for returns.

The CAPM: Conclusion

- All investors work with the same mean-standard deviation diagram.
- All investors hold a mean-variance efficient portfolio.
- Since all mean-variance efficient portfolios combine the riskless asset with a fixed portfolio of risky assets, all investors hold risky assets in the same proportions to one another.
- These proportions must be those of the *market portfolio* or *value-weighted index* that contains all risky assets in proportion to their market value.
- Thus *the market portfolio is mean-variance efficient*.

The CAPM: Investment Implications

- A mean-variance investor need not actually perform the mean-variance analysis! The investor can "free-ride" on the analyses of other investors, and use the market portfolio (in practice, a broad index fund) as the optimal mutual fund of risky assets (tangency portfolio).
- The optimal capital allocation line (CAL) is just the *capital market line (CML)* connecting the riskfree asset to the market portfolio.

The CAPM: Asset Pricing Implications

- We look at properties of mean-variance efficient portfolios.
- The CAPM implies that these properties hold for the market index.

Covariance Properties of Efficient Portfolios

- An increase in portfolio weight w_i, financed by a decrease in the weight on the riskless asset, affects the mean and variance of the portfolio return as follows:

$$\frac{d\overline{R}_p}{dw_i} = \overline{R}_i - R_f \ .$$

$$\frac{d\mathrm{Var}(R_p)}{dw_i} = 2\mathrm{Cov}(R_i, R_p).$$

If portfolio p is efficient, the ratio of these two effects should be the same for all assets. Why?

Covariance Properties of Efficient Portfolios

Adjust two different portfolio weights, w_i and w_j:

$$d\overline{R}_p = (\overline{R}_i - R_f)dw_i + (\overline{R}_j - R_f)dw_j.$$

$$d\mathrm{Var}(R_p) = 2\mathrm{Cov}(R_i, R_p)dw_i + 2\mathrm{Cov}(R_j, R_p)dw_j.$$

Set dw_j so that the mean portfolio return is unchanged, $d\overline{R}_p = 0$:

$$dw_j = -\frac{(\overline{R}_i - R_f)}{(\overline{R}_j - R_f)}dw_i.$$

Then the portfolio variance must also be unchanged, because otherwise one could achieve a lower variance with the same mean, which would contradict the assumption that the portfolio is efficient. We have

$$d\mathrm{Var}(R_p) = \left[2\mathrm{Cov}(R_i, R_p) - 2\mathrm{Cov}(R_j, R_p)\frac{(\overline{R}_i - R_f)}{(\overline{R}_j - R_f)}\right] dw_i = 0.$$

Covariance Properties of Efficient Portfolios

This requires

$$\frac{\overline{R}_i - R_f}{2\mathrm{Cov}(R_i, R_p)} = \frac{\overline{R}_j - R_f}{2\mathrm{Cov}(R_j, R_p)}.$$

This equation must hold for all assets j, including the original portfolio itself. Setting $j = p$, we get

$$\frac{\overline{R}_i - R_f}{2\mathrm{Cov}(R_i, R_p)} = \frac{\overline{R}_p - R_f}{2\mathrm{Var}(R_p)},$$

$$\overline{R}_i - R_f = \frac{\mathrm{Cov}(R_i, R_p)}{\mathrm{Var}(R_p)}(\overline{R}_p - R_f) = \beta_{ip}(\overline{R}_p - R_f),$$

where $\beta_{ip} \equiv \mathrm{Cov}(R_i, R_p)/\mathrm{Var}(R_p)$ is the regression coefficient of asset i on portfolio p.

Back to the CAPM

Since the CAPM implies that the market portfolio m is efficient, this equation describes the market portfolio:

$$\overline{R}_i - R_f = \beta_{im}(\overline{R}_m - R_f).$$

If we consider the regression of excess returns on the market excess return,

$$R_{it} - R_{ft} = \alpha_i + \beta_{im}(R_{mt} - R_{ft}) + \varepsilon_{it},$$

$\alpha_i \equiv \overline{R}_i - R_f - \beta_{im}(\overline{R}_m - R_f)$ should be zero for all assets.
α_i is called *Jensen's alpha* and is used to try to find assets that are mispriced relative to the CAPM. The relationship

$$\overline{R}_i = R_f + \beta_{im}(\overline{R}_m - R_f)$$

is called the *security market line (SML)*, and α_i measures deviations from this line.

Harvard Policy Portfolio

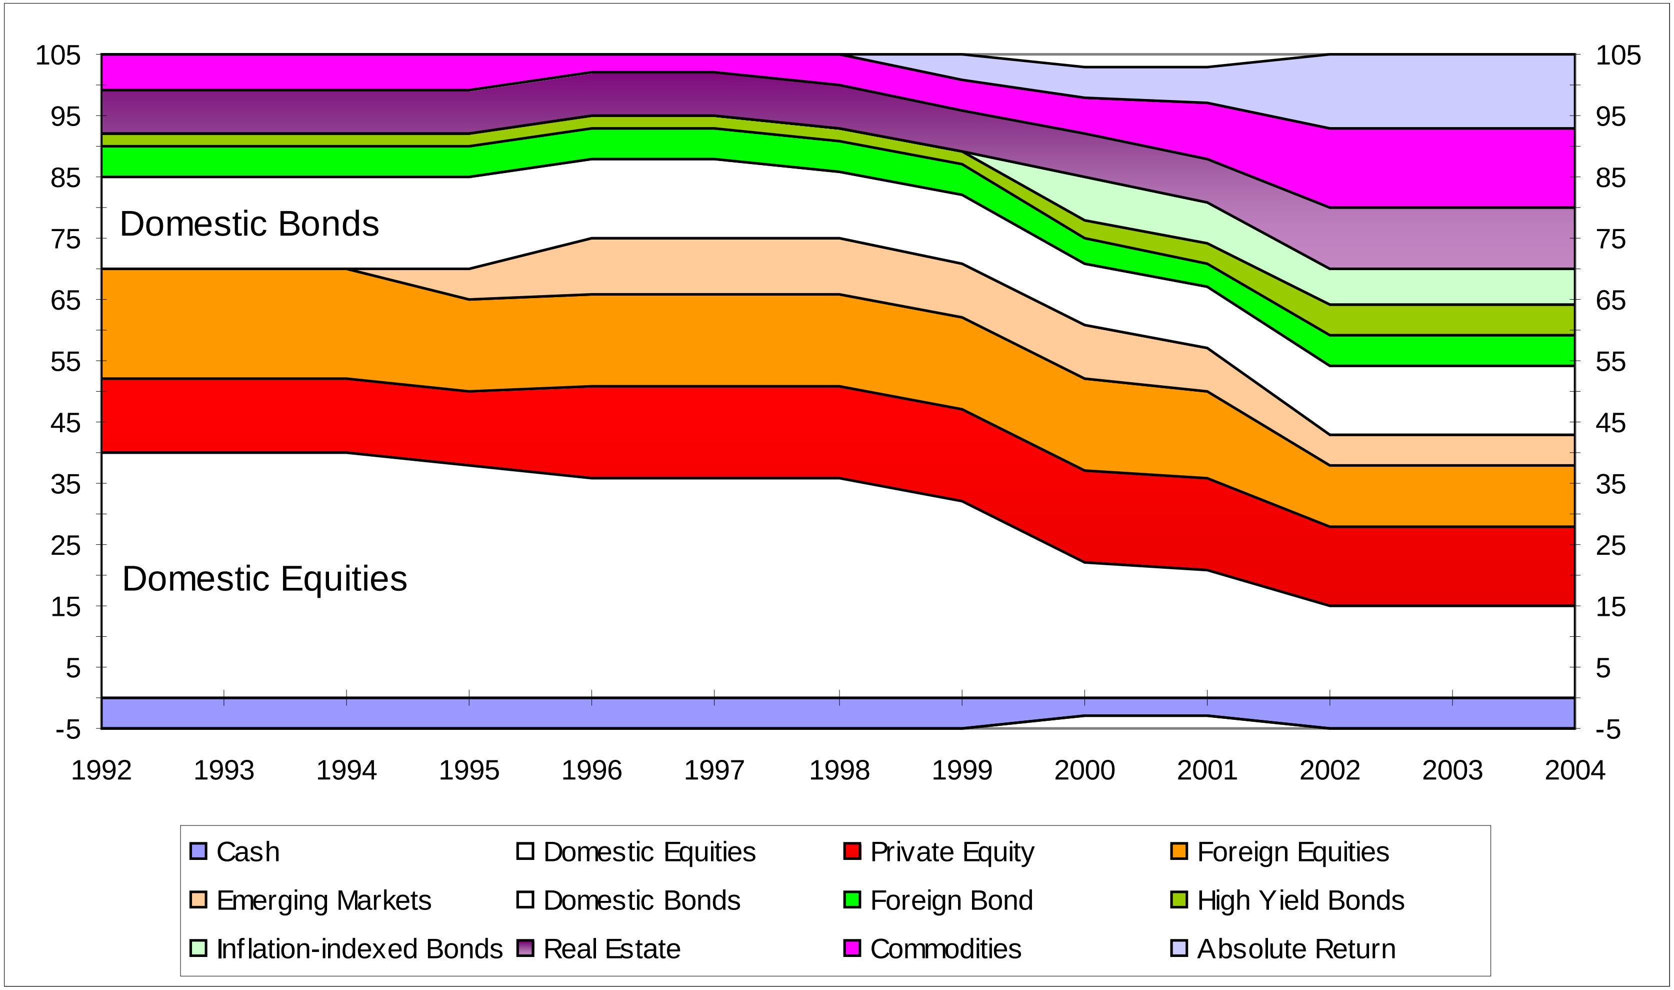

Harvard Investment Beliefs (1)

Source: HMC Capital Market Assumptions, 2004

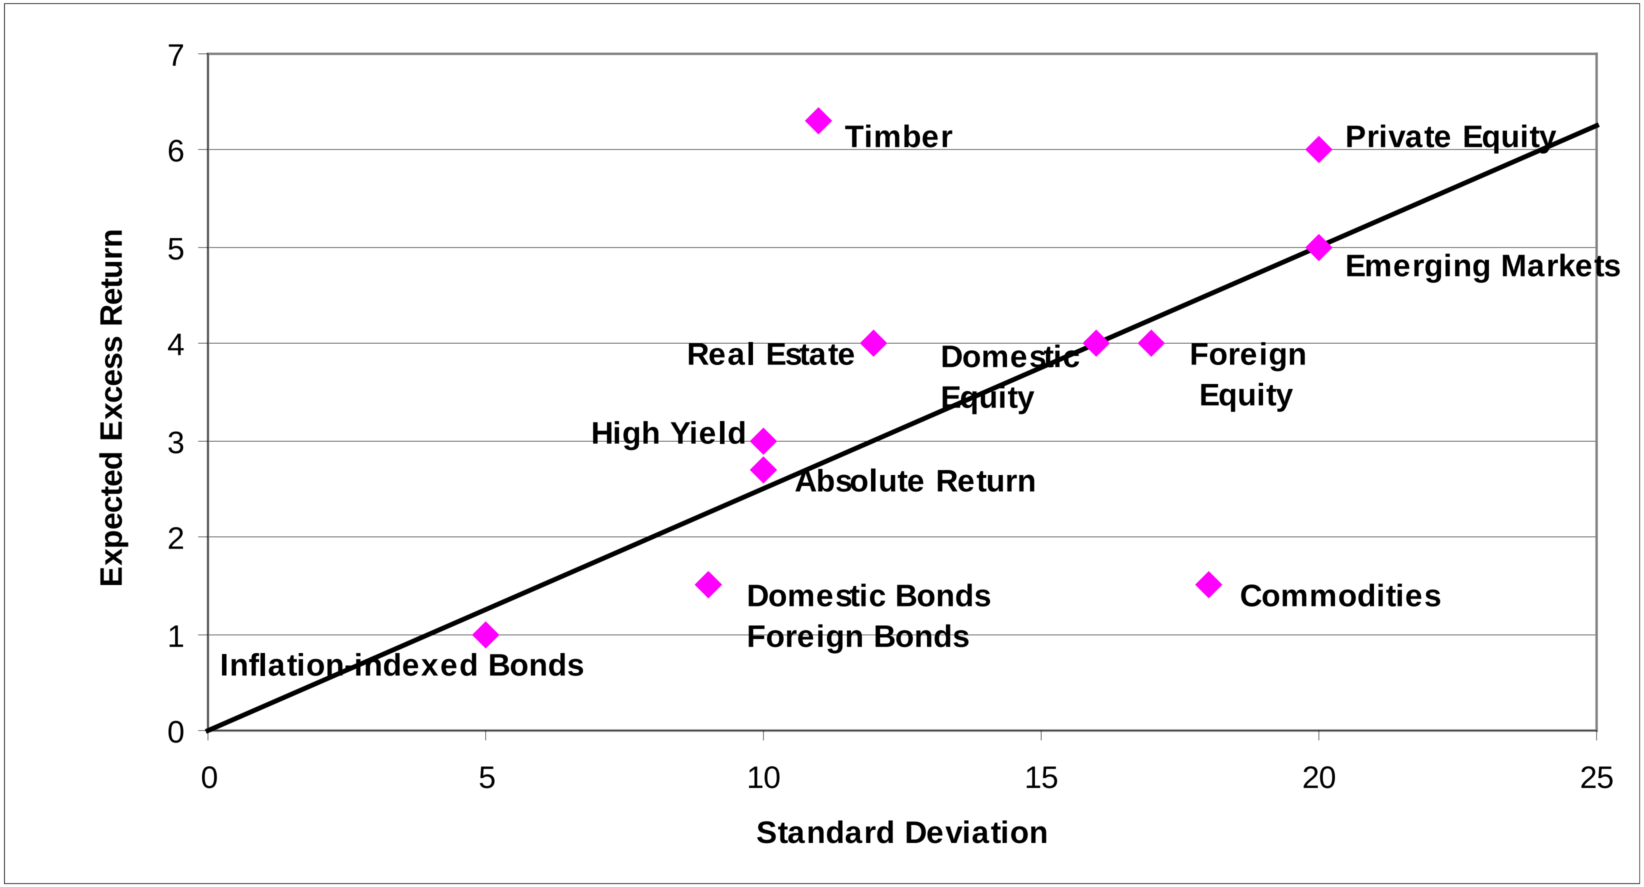

Harvard Investment Beliefs (2)

Source: HMC Capital Market Assumptions, 2004

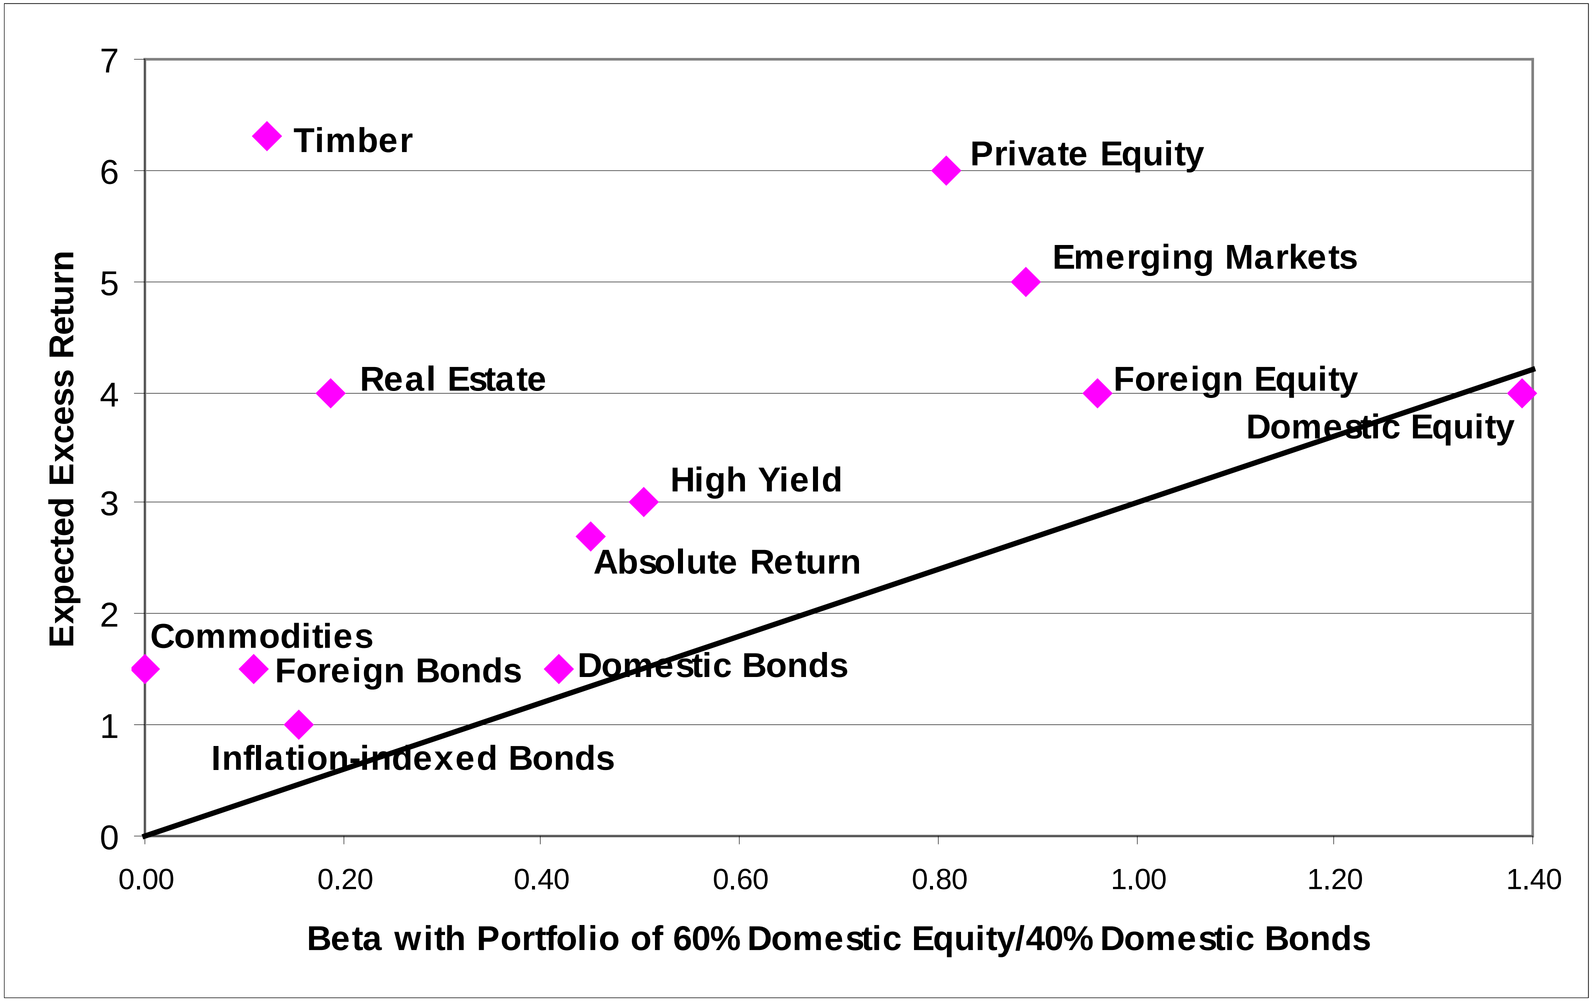

The CAPM Without Riskless Borrowing and Lending

- All investors choose combinations of the same two mutual funds.
- The market portfolio must be a combination of these mutual funds, and must therefore be efficient.
- Analysis of covariance properties goes through as before, except that we replace the riskless asset with an efficient portfolio z that is uncorrelated with the market portfolio. We get

$$\overline{R}_i - \overline{R}_z = \beta_{im}(\overline{R}_m - \overline{R}_z) \ ,$$

where β_{im} is defined as before.

- This version of the CAPM is due to Fischer Black.

Econometrics of CAPM Tests

Two approaches, time-series and cross-sectional, are related at a deep level.
Time-series approach:

$$R^e_{it} = \alpha_i + \beta_{im} R^e_{mt} + \varepsilon_{it},$$

Consider N assets jointly, create

- α as the N-vector of intercepts α_i
- Σ as the variance-covariance matrix of the **residuals** ε_{it}

Time-Series Approach: Asymptotics

$$T\left[1+\left(\frac{\overline{R}_m^e}{\sigma(R_{mt}^e)}\right)^2\right]^{-1}\widehat{\alpha}'\widehat{\Sigma}^{-1}\widehat{\alpha}\sim\chi_N^2$$

for large T. Intuition:

- Suppose there were no market return in the model. Then the vector α would be a vector of sample mean excess returns, with variance-covariance matrix $(1/T)\Sigma$.
- The quadratic form $\widehat{\alpha}'\widehat{\Sigma}^{-1}\widehat{\alpha}$ is a sum of squared intercepts, divided by its variance-covariance matrix, which has a χ_N^2 distribution.
- The term in square brackets is a correction for the presence of the market return in the model. Uncertainty about the betas affects the alphas, and more so when the market has a high expected return relative to its variance.

Time-Series Approach: Finite Sample

$$\left(\frac{T-N-1}{N}\right)\left[1+\left(\frac{\overline{R}_m^e}{\sigma(R_{mt}^e)}\right)^2\right]^{-1}\widehat{\alpha}'\widehat{\Sigma}^{-1}\widehat{\alpha} \sim F_{N,T-N-1}.$$

if ε_{it} are serially uncorrelated, homoskedastic, and normal.

- Degrees of freedom correction because Σ must be estimated.

Time-series approach: Geometry

Gibbons-Ross-Shanken (1989): Previous test statistic can be written as

$$\left(\frac{T-N-1}{N}\right)\left(\frac{\hat{S}_q^2-\hat{S}_m^2}{1+\hat{S}_m^2}\right).$$

where

- $\hat{S}_m^2$ is the estimated squared Sharpe ratio of the market
- $\hat{S}_q^2$ is the squared Sharpe ratio of the estimated tangency portfolio (the highest squared Sharpe ratio available from the set of test assets together with the market)

Cross-Section Approach

First estimate betas using time-series regression, then run a cross-sectional regression

$$\overline{R}_i^e = \lambda \beta_{im} + \alpha_i,$$

where

- There is no intercept in the regression (or it can be added to test whether it is zero, or allowed to be free in Black version of the CAPM)
- λ is the cross-sectional reward for bearing market risk
- The alphas are now regression residuals
- We want to test whether the residuals are zero. How is this possible?

Cross-Section Approach: Using Correlation

As before, we know that $\mathrm{E}(\alpha\alpha') = (1/T)\Sigma$. Thus the residuals of the cross-section regression are correlated with one another. To correct for this, we can run Generalized Least Squares and get

$$\widehat{\lambda}_{GLS} = (\beta'\Sigma^{-1}\beta)^{-1}\beta'\Sigma^{-1}\overline{R}^{e},$$

$$\widehat{\alpha}_{GLS} = \overline{R}^{e} - \widehat{\lambda}_{GLS}\beta.$$

Cross-Section Approach: Asymptotics

An asymptotic test statistic based on the GLS cross-sectional regression, and correcting for the fact that the betas are not known but estimated from prior time-series regressions, is

$$T\left[1+\left(\frac{\widehat{\lambda}_{GLS}}{\sigma(R^e_{mt})}\right)^2\right]^{-1}\widehat{\alpha}'_{GLS}\widehat{\Sigma}^{-1}\widehat{\alpha}_{GLS}\sim\chi^2_{N-1}.$$

This is very similar to the time-series test statistic

$$T\left[1+\left(\frac{\overline{R}^e_m}{\sigma(R^e_{mt})}\right)^2\right]^{-1}\widehat{\alpha}'\widehat{\Sigma}^{-1}\widehat{\alpha}\sim\chi^2_N$$

Can we make the two identical?

Cross-Section and Time-Series Approach

- We lost one degree of freedom by estimating the reward for beta in the cross-section rather than from the average excess market return.
- We can earn that degree of freedom back by adding the market to the set of assets in the cross-sectional GLS regression. Then the regression puts all the weight on the market in estimating the reward for market risk (since other assets are just the market plus noise, so the GLS regression knows they are less informative about this parameter).
- The result is a test statistic that is exactly the same as the time-series test statistic
- The cross-sectional approach has the advantage that it can be implemented even when the factor is not the return on a traded portfolio. In that case we need to use the cross-section to estimate the reward for bearing factor risk.

Fama-MacBeth Approach

First estimate betas, then run a series of period-by-period cross-sectional regressions.

$$R_{it}^e = \lambda_t \beta_{im} + \alpha_{it}.$$

- The observations in each regression run from $i = 1...N$.
- The regressions are run separately for each $t = 1...T$.
- The coefficients and residuals are then averaged over time to estimate the average reward for beta exposure $\widehat{\lambda}$ and the average alphas $\widehat{\alpha}_i$.
- We use the variability of the coefficients and residuals over time to estimate the standard errors of these averages and construct test statistics for the model.

Properties of Fama-MacBeth Approach

- When the explanatory variables in the regression do not vary over time, the Fama-MacBeth approach is equivalent to cross-sectional OLS using the entire sample average, or to a pooled time-series cross-sectional OLS regression, with standard errors corrected for cross-sectional correlation of residuals.
- When the explanatory variables do vary over time, Fama-MacBeth is different because it gives equal weight to each time period, even if the explanatory variables are more dispersed in one period than another.
- The basic Fama-MacBeth approach does not adjust for the fact that betas are not known but must be estimated from time-series regressions. However it does easily allow for changing betas over time. Thus it is more appropriate as a method to estimate the rewards to observable characteristics of firms (which could include their lagged historical betas), than as a method to test the CAPM.

Common Themes in CAPM Tests

- Looking at more assets (increasing N) will tend to find larger deviations from the model.
- But increasing N also increases the size of the deviations you *need* to find to reject the model statistically.
- Thus, to get a powerful test you need to group assets into a few portfolios that summarize the behavior of a larger set of assets.
- However it is cheating to do this after looking at the average returns on the full set of assets and picking portfolios based on this information. This *data-snooping* leads to spurious rejections of the model.
- Much of the debate about the empirical validity of the CAPM centers on this issue.

Figure 2

Average Annualized Monthly Return versus Beta for Value Weight Portfolios Formed on Prior Beta, 1928–2003

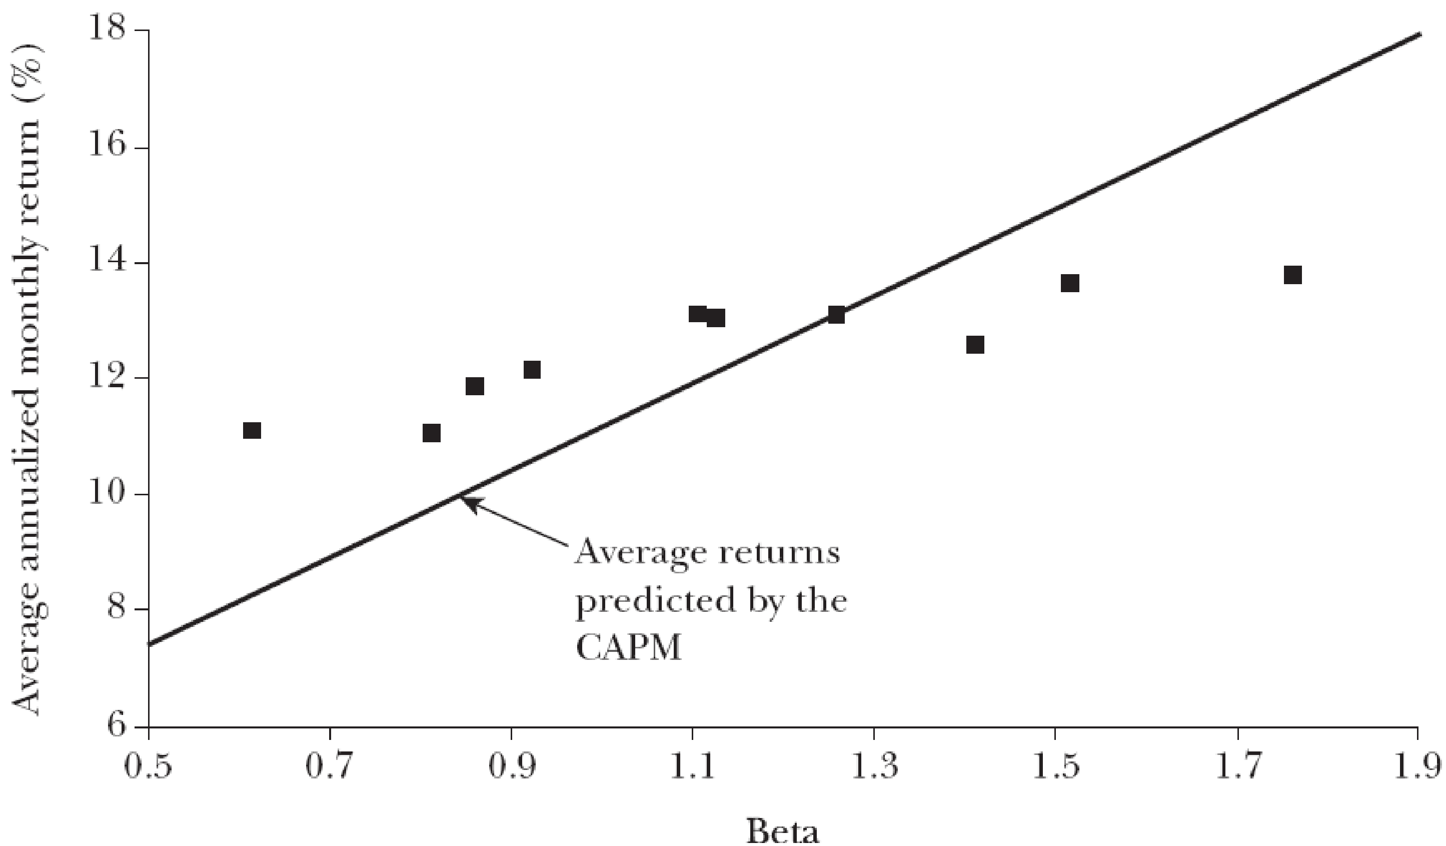

Figure 3

Average Annualized Monthly Return versus Beta for Value Weight Portfolios Formed on B/M, 1963–2003

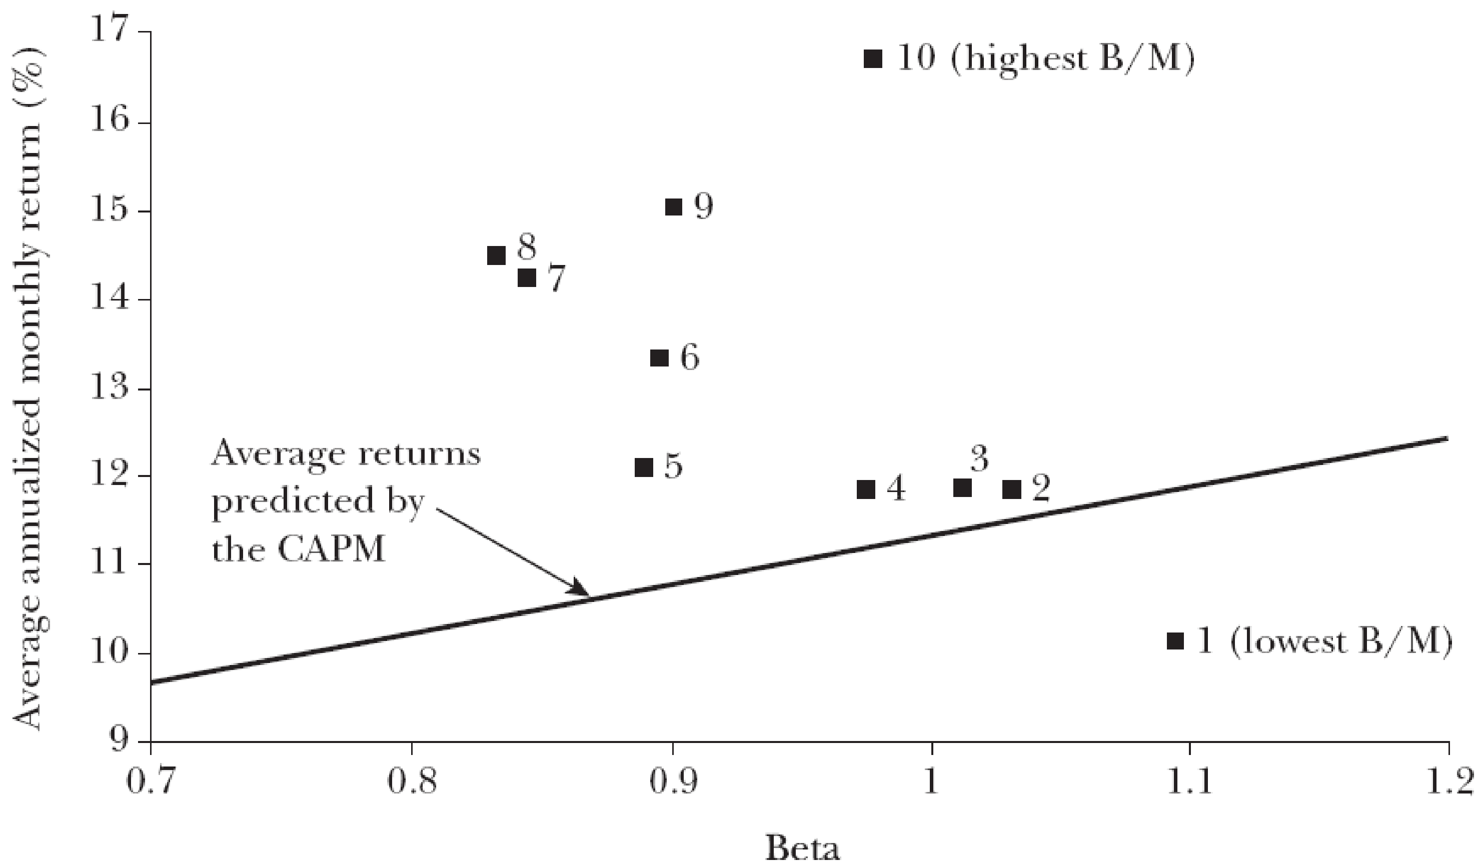

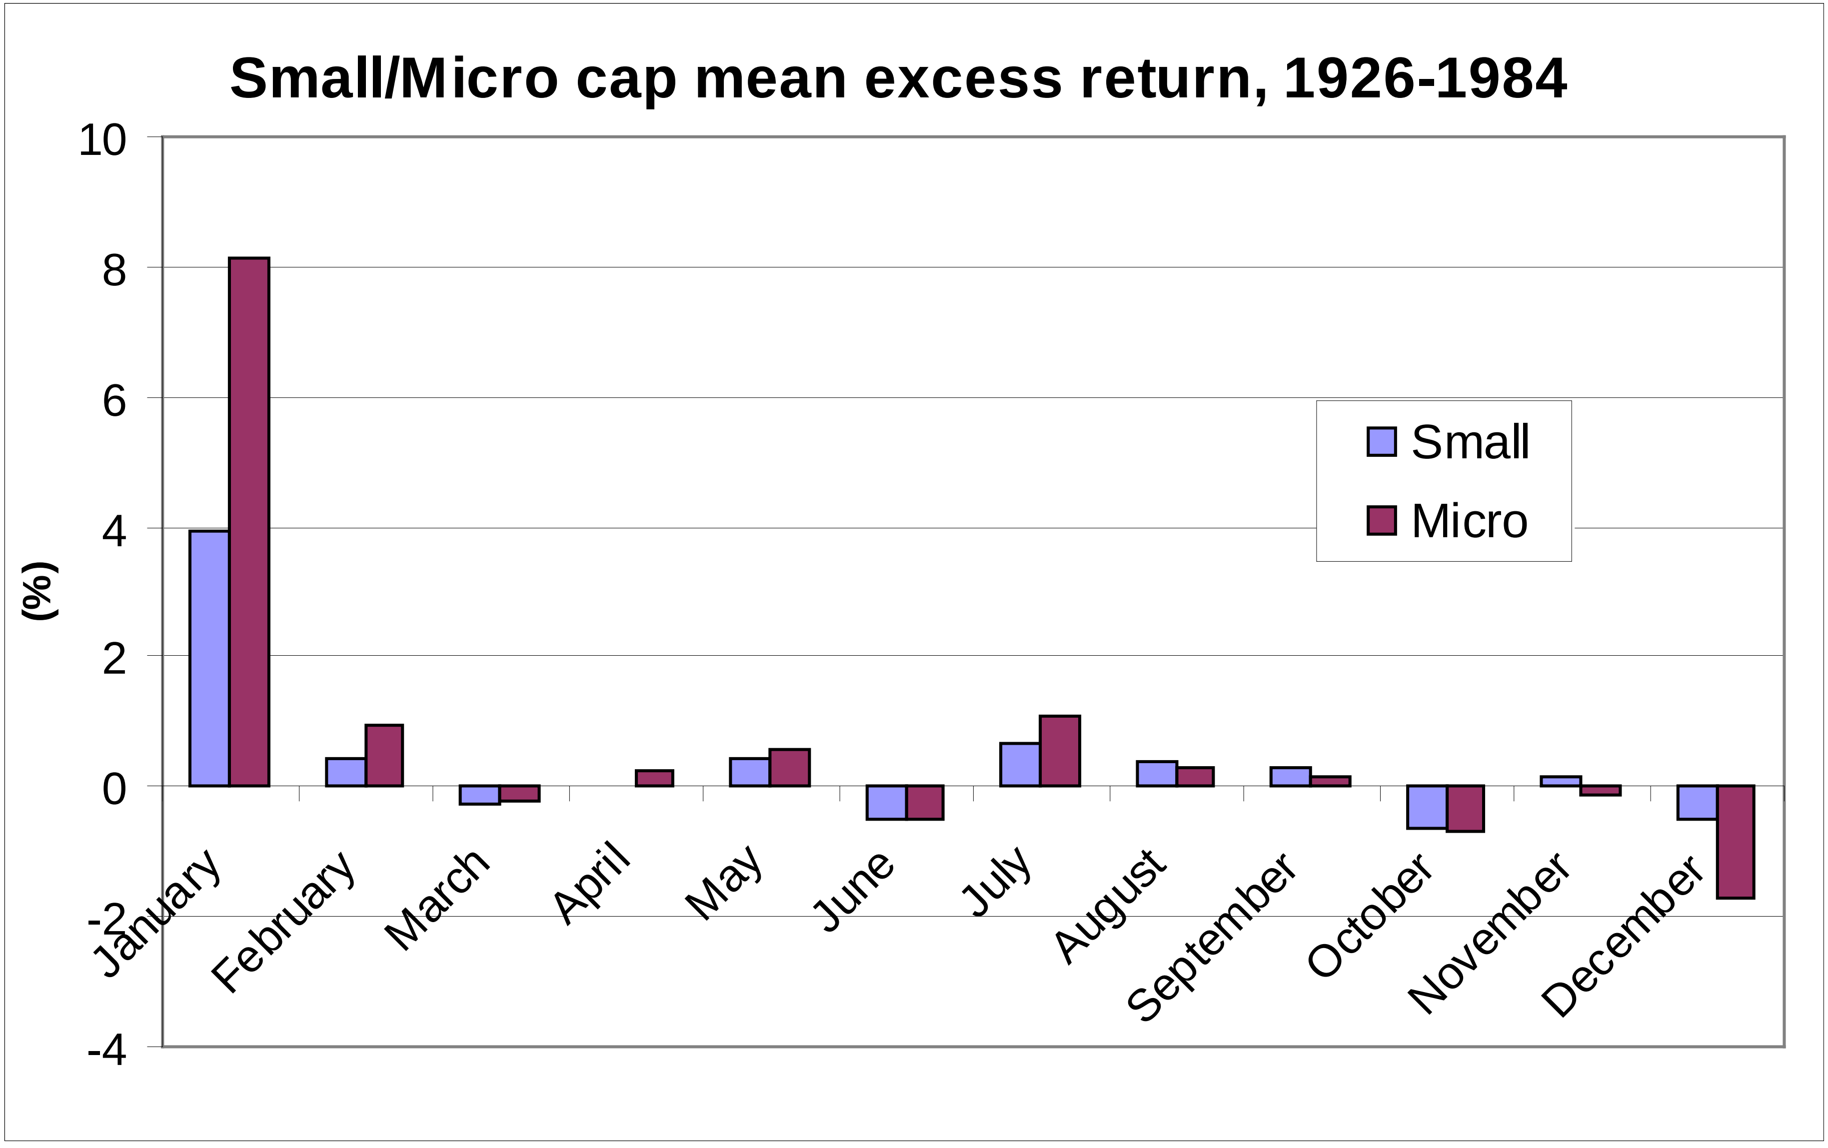

Small/Micro cap mean excess return, 1926-1984
10
8
6
4
2
0
-2
-4
(%)
Small
Micro
January
February
March
April
May
June
July
August
September
October
November
December

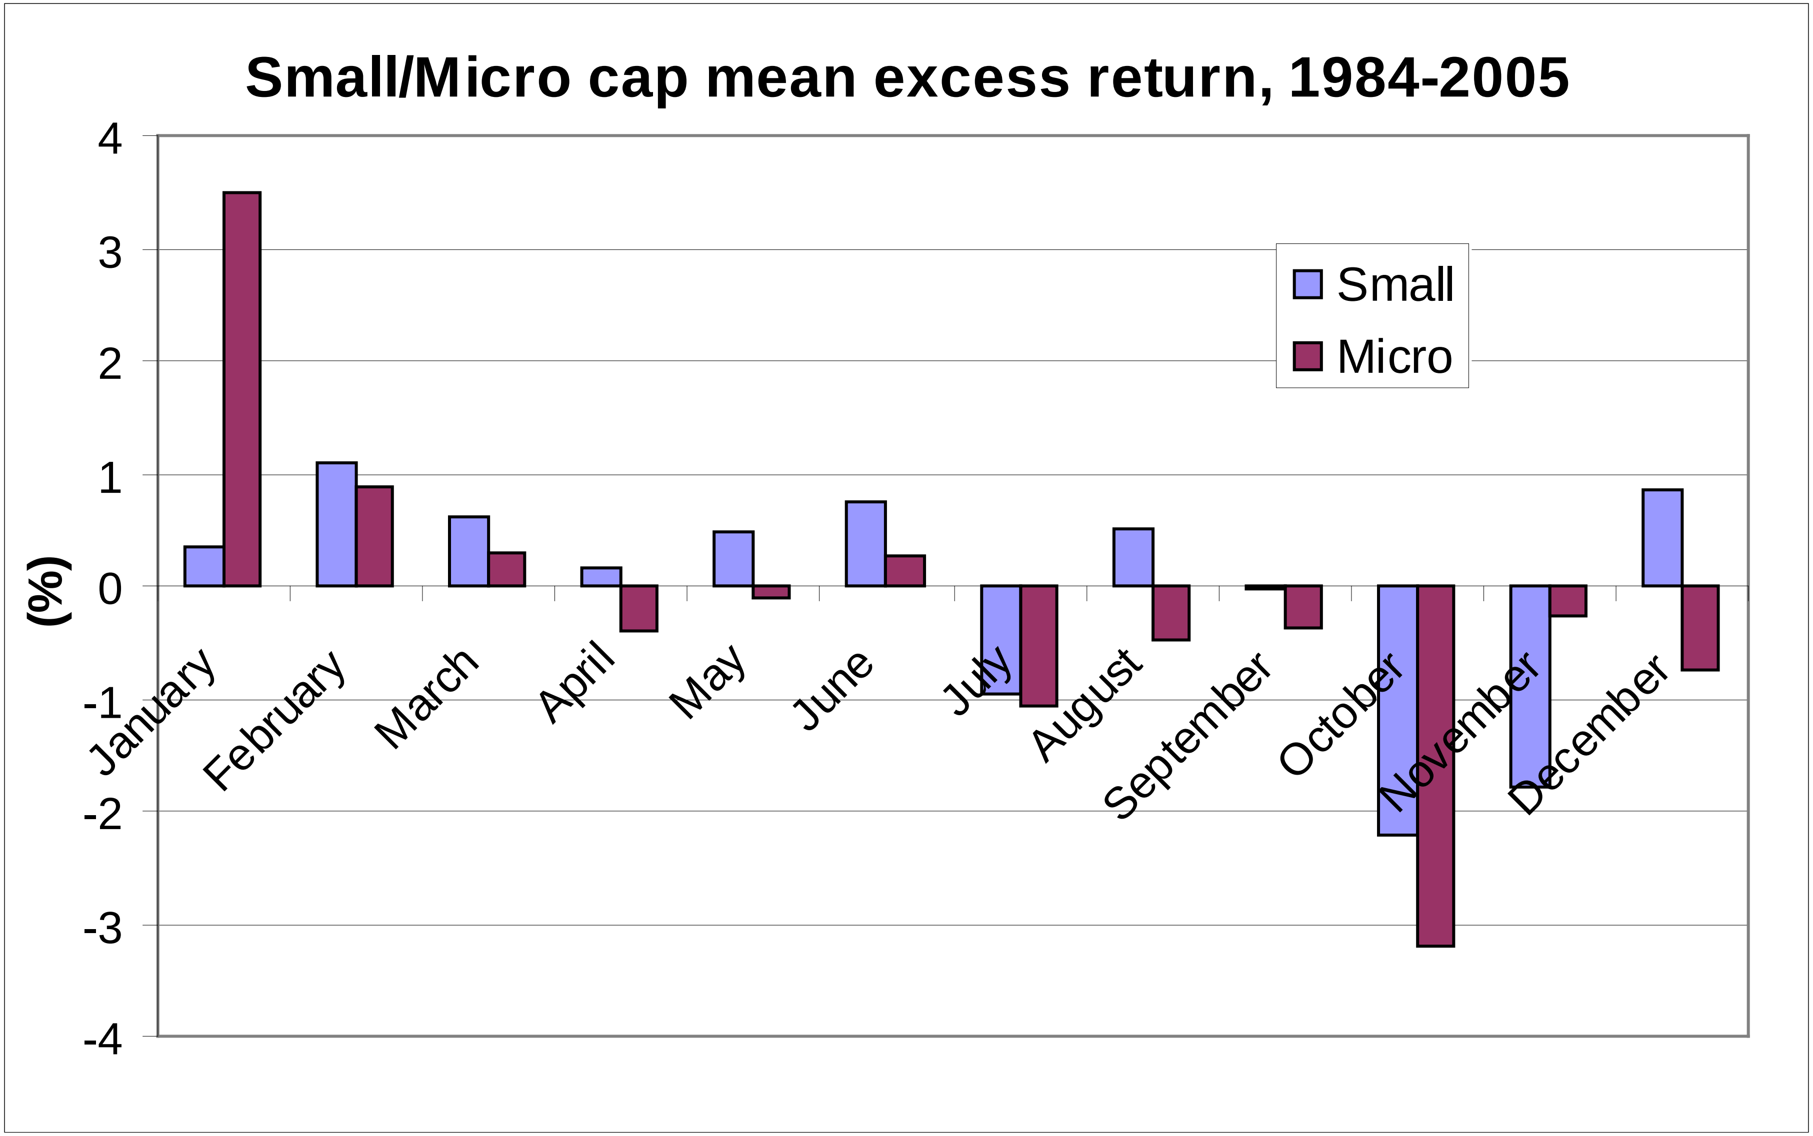

Small/Micro cap mean excess return, 1984-2005
4
3
2
1
0
-1
-2
-3
-4
(%)
Small
Micro
January
February
March
April
May
June
July
August
September
October
November
December

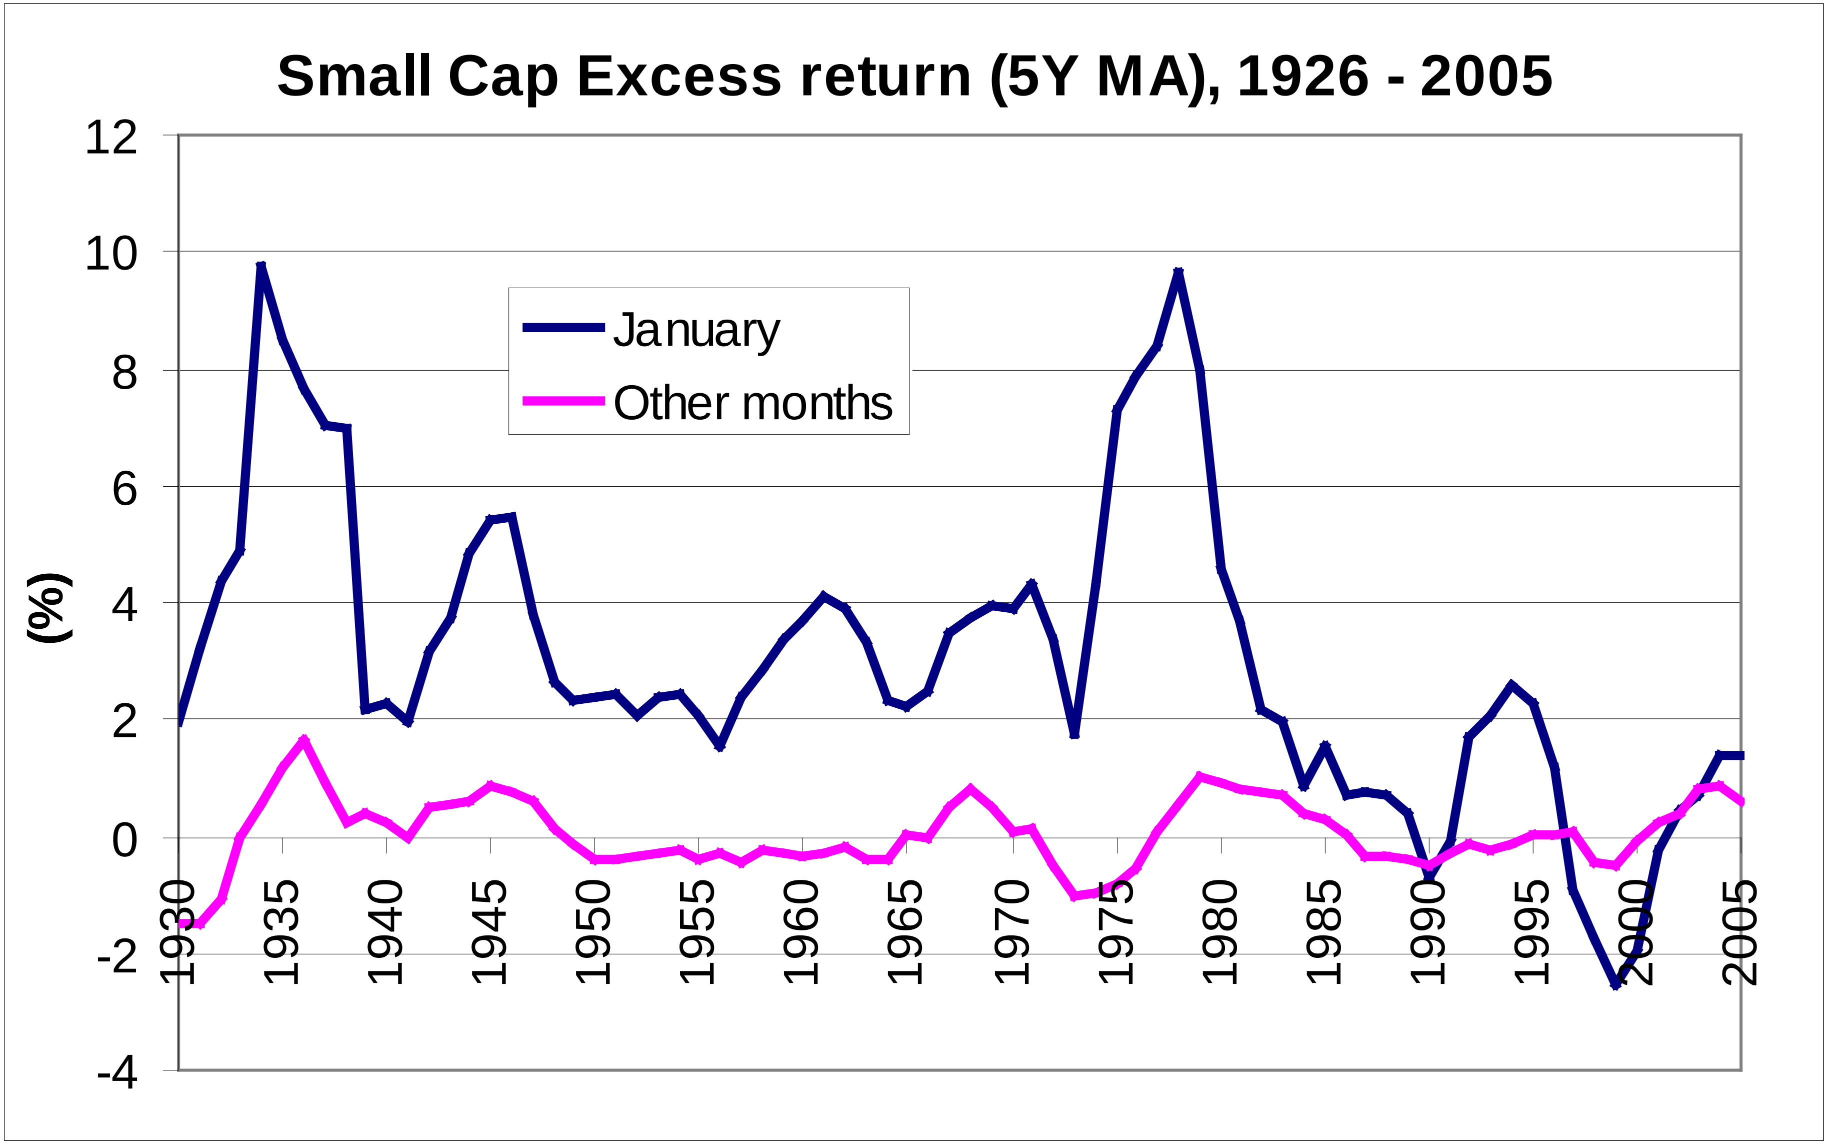
Small Cap Excess return (5Y MA), 1926 - 2005
(%)
12
10
8
6
4
2
0
-2
-4
1930
1935
1940
1945
1950
1955
1960
1965
1970
1975
1980
1985
1990
1995
2000
2005
January
Other months

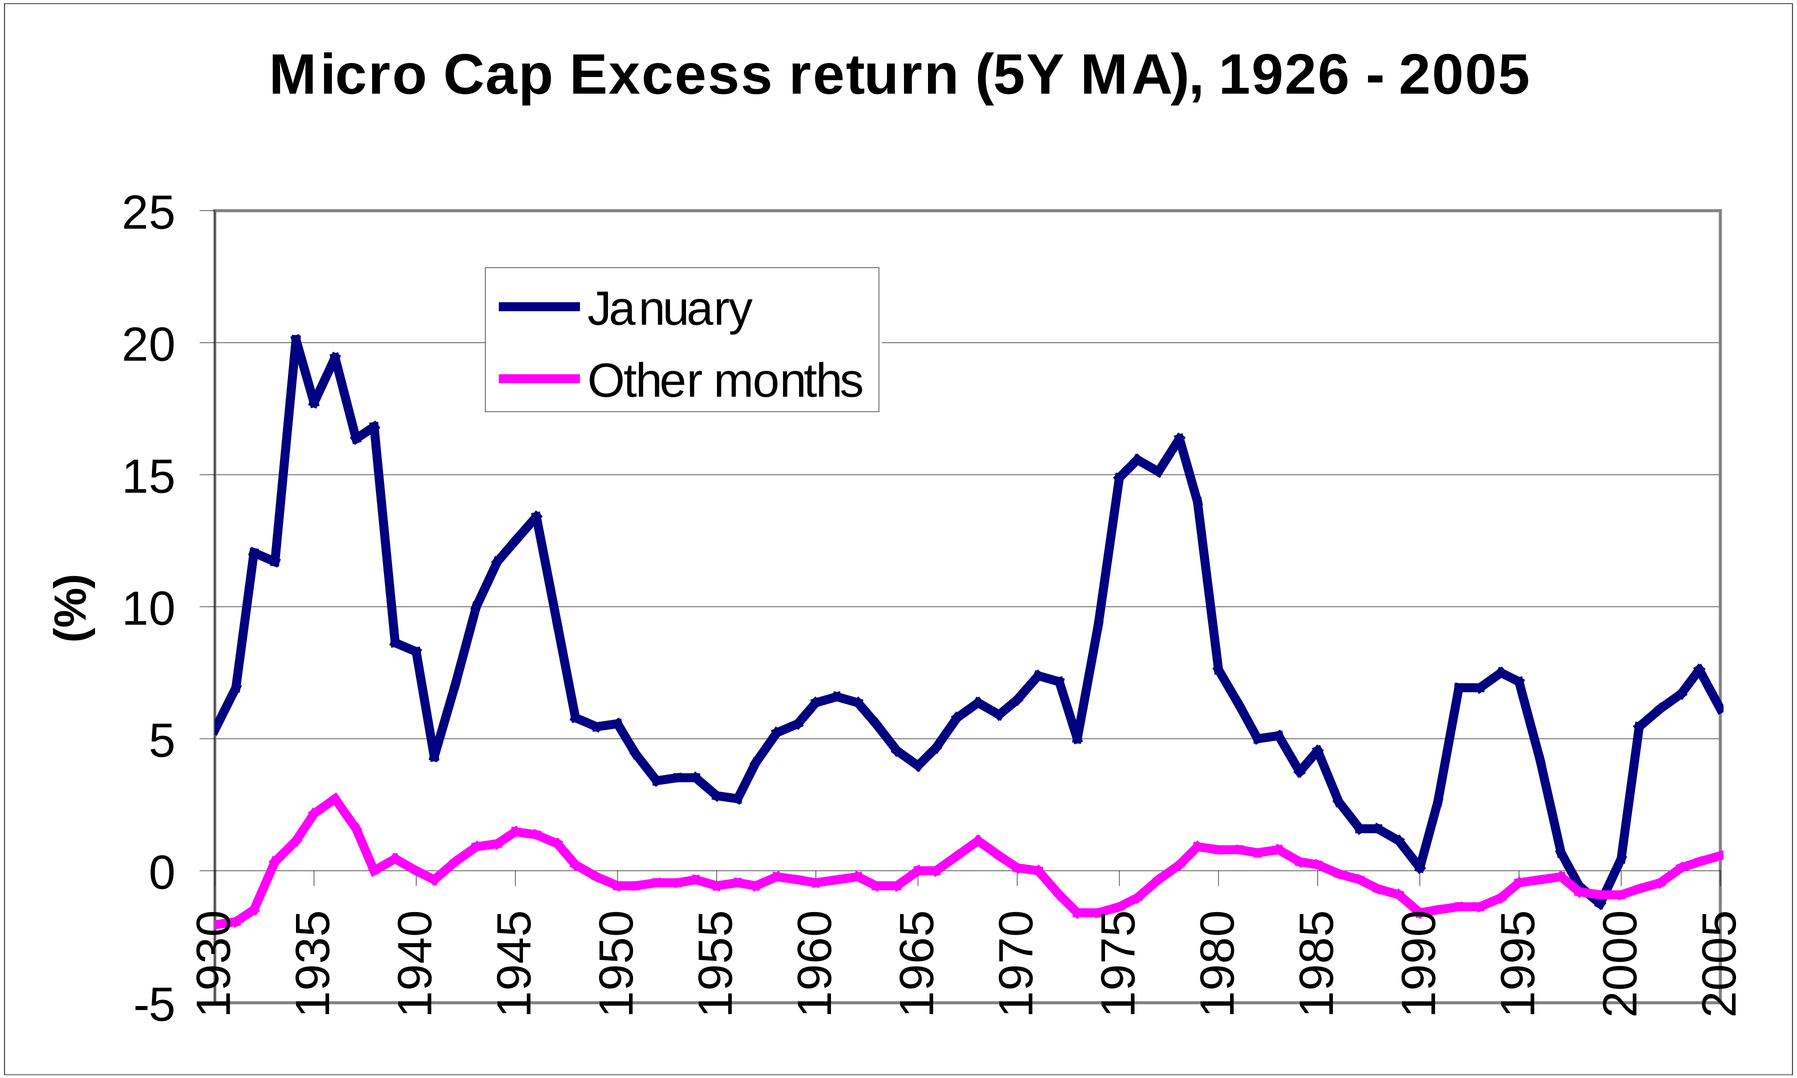

Micro Cap Excess return (5Y MA), 1926 - 2005
(%)
25
20
15
10
5
0
-5
1930
1935
1940
1945
1950
1955
1960
1965
1970
1975
1980
1985
1990
1995
2000
2005
January
Other months

A Single-Factor Model

$$R_{it} - R_f = \alpha_i + \beta_{im}(R_{mt} - R_f) + \epsilon_{it}.$$

This relationship is called the *market model*. It is the leading example of a *single-factor model* with a single common factor moving stock returns. Suppose that the errors in this equation are uncorrelated across stocks:

$$\mathrm{E}[\epsilon_{it}\epsilon_{jt}] = 0$$

for $i \neq j$. Then the residual risk in any stock is *idiosyncratic*, unrelated to the residual risk in any other stock.
Implications:

- Covariances are easy to estimate for mean-variance analysis because

$$\mathrm{Cov}(R_{it}, R_{jt}) = \beta_{im}\beta_{jm}\sigma_m^2.$$

- If many assets are available, we should expect α_i typically to be very small in absolute value. Why?

Arbitrage Pricing in a Single-Factor Model

Form a portfolio of N assets i. The portfolio return will be

$$R_{pt} - R_f = \alpha_p + \beta_{pm}(R_{mt} - R_f) + \epsilon_{pt},$$

where $\alpha_p = \sum_{j=1}^{N} w_j \alpha_j$, $\beta_{pm} = \sum_{j=1}^{N} w_j \beta_{jm}$, and $\epsilon_{pt} = \sum_{j=1}^{N} w_j \epsilon_{jt}$.
The variance of ϵ_{pt} will be

$$\mathrm{Var}(\epsilon_{pt}) = \sum_{j=1}^{N} w_j^2 \mathrm{Var}(\epsilon_{jt}),$$

which will shrink rapidly with N provided that no single weight w_j is too large.

Arbitrage Pricing in a Single-Factor Model

Now suppose we have enough stocks, with a small enough weight in each one, that the residual risk $\mathrm{Var}(\epsilon_{pt})$ is negligible. We say that the portfolio is *well diversified*. For such a portfolio, we can neglect ϵ_{pt} and write the return as

$$R_{pt} - R_f = \alpha_p + \beta_{pm}(R_{mt} - R_f).$$

But then we must have $\alpha_p = 0$. If not, there is an arbitrage opportunity: short β_{pm} units of the market, long one unit of the portfolio. This gives a riskless excess return of α_p.

Arbitrage Pricing in a Single-Factor Model

Ross APT (1976) builds on this to show that $\alpha_p = 0$ for all well diversified portfolios implies that

$$\lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} \alpha_i^2 = 0.$$

Almost all individual assets have α_i very close to zero. Beta pricing without the CAPM assumptions.

Arbitrage Pricing in a Multi-Factor Model

$$R_{it} - R_f = \alpha_i + \sum_{k=1}^{K} \beta_{ik}(R_{kt} - R_f) + \epsilon_{it}.$$

We assume that the residual is uncorrelated across stocks. The prediction of the model is that $\alpha_i = 0$ for almost all stocks. This is restrictive if $K << N$.

Alternatively, if we measure the factors directly as mean-zero shocks (for example, innovations to macroeconomic variables), then we have

$$R_{it} - R_f = \mu_i + \sum_{k=1}^{K} \beta_{ik} F_{kt} + \epsilon_{it},$$

and the prediction of the model is that

$$\mu_i = \sum_{k=1}^{K} \beta_{ik} \lambda_k,$$

where λ_k is the *price of risk* of the k'th factor.

Interpretation of Multi-Factor Models

- Mean-variance analysis: The full frontier can be constructed from the K factor portfolios, so mean-variance investors always hold some combination of these.
- Some portfolio is always ex post mean-variance efficient. Thus we know we can always get a 1-factor model to fit the data. A fortiori, we can always get a K-factor model to fit the data. What does this tell us about the world?
- How to pick the factors?
- The theory does not determine the risk prices.

Conditional vs. Unconditional CAPM

Suppose that the CAPM holds conditionally:

$$\mathrm{E}_t R^e_{i,t+1} = \beta_{imt} \mathrm{E}_t R^e_{m,t+1}.$$

Taking unconditional expectations,

$$\mathrm{E} R^e_{i,t+1} = \overline{\beta}_{im} \mathrm{E} R^e_{m,t+1} + \mathrm{Cov}(\beta_{imt}, E_t R^e_{m,t+1}).$$

An asset can have a higher unconditional average return than predicted by the unconditional CAPM, if its beta moves with the market risk premium.

Conditional vs. Unconditional CAPM

One way to test a conditional model is to parameterize the variables that shift betas over time. For example, we might write

$$\beta_{imt} = \beta_{i0} + \beta_{i1} z_t.$$

The conditional model can then be written as

$$\mathrm{E}_t R^e_{i,t+1} = \beta_{i0} \mathrm{E}_t R^e_{m,t+1} + \beta_{i1} \mathrm{E}_t z_t R^e_{m,t+1}$$

and now we can take unconditional expectations to get

$$\mathrm{E} R^e_{i,t+1} = \beta_{i0} \mathrm{E} R^e_{m,t+1} + \beta_{i1} \mathrm{E} z_t R^e_{m,t+1}.$$

- Multifactor model with market and z-scaled market as factors
- Test using time-series or cross-sectional approach
- If cross-sectional regression is used, must include the market and scaled market in the set of test assets.