

Asset and Risk Management I

Giorgia Simion and Otto Randl

WU Vienna University of Economics and Business

Stocks

1. Time series evidence

- Motivation
- Time series predictability: theory
- The dividend yield as a predictor
- Other predictors

2. Cross-sectional evidence


- Motivation
- Review of APT and factor models
- Quantifying risk premia
- The factor zoo

Time series evidence

Stylized facts from historical data


- From 1926–2024, average annual mean excess return $\approx 8\%$ and standard deviation $\approx 20\%$
 - What's the 95% confidence interval for the equity risk premium (ERP)?
- Sharpe ratio ≈ 0.42
 - U.S. historical returns look very attractive, also compared to most other countries.
- Typical features (more pronounced at higher frequency):
 - fat tails
 - left-skewed
 - autocorrelation close to zero
 - Stocks tend to decline in bad economic times, e.g., during the Global Financial Crisis or the COVID-19 pandemic.

U.S. stock market returns



Source: own calculations. Data: CRSP U.S. Stock Market Index, Fama-French 3-month T-Bills, 1928–2024.

Histogram



Source: own calculations. Data: CRSP U.S. Stock Market Index, Fama-French 3-month T-Bills, 1928–2023.

Why should we care about predictability?

- Remember from last week the basic allocation problem between the risky and the risk-free asset.
- The optimal weight is given by

$$w^* = \frac{\mu}{A\sigma^2}$$

- Keeping everything else fixed:
 1. Invest more in stocks if the risk premium $\mu = E[R] - R_f$, increases
 2. Invest less in stocks if they become riskier, $\sigma \uparrow$
 3. Invest less in stocks if you become more risk averse, $A \uparrow$
- Understanding $E[R]$ is crucial. Is it time-varying?

Are returns predictable? (1)

- “Predict” does not mean with certainty, of course, but is there a way to know that the odds are in your favor on some days and against you on others?
- To find out if returns are (a bit) predictable, we can simply run a regression

$$R_{t+1} = a + b \cdot x_t + \epsilon_{t+1},$$

where R_{t+1} = return over next period, x_t = explanatory variable observable today, and ϵ_{t+1} is an error term.

Are returns predictable? (2)

- We run regressions of tomorrow's return R_{t+1} on any variable we can see today x_t .
- If we find a “big” regression coefficient b , or a large R^2 , you might be able to make money, buying when x_t is high and vice versa. If you find a small b , and low R^2 , returns are not predictable, and you can't make much money.
- Thus, the regression measures the question “are returns (somewhat) predictable?”
- Equivalently, this regression model implies that the expected return at time t is $\mathbb{E}_t[R_{t+1}] = a + b \cdot x_t$.
- The “expected return” can vary over time, being higher when the signal x_t is higher and vice versa.

The classic efficient market view

- In the classic efficient-markets view, stock prices are not predictable, (loosely, the “random walk” view) so we should see $b = 0$ and $R^2 = 0$, for any variable.
- If you could predict good/bad days, what would you do? On days when a signal x_t predicts good returns, you buy, and vice versa.
- Of course, we can't all do this. If everyone sees x_t , they also try to buy, driving prices up until today's price is the same as the price we expect tomorrow (less a small discount).
- Hence, competition in stock markets should drive out any predictable movement in stock prices. The information in the signal x_t will quickly be impounded in today's price.

Predicting returns with past returns

Regression of returns on lagged returns, annual data 1928 - 2024

$$R_{t+1} = a + b \cdot R_t + \epsilon_{t+1}$$

	b	$t(b)$	R^2	$\mathbb{E}[R]$	$\sigma(\mathbb{E}_t[R_{t+1}])$	$\sigma(R_{t+1})$
Stock	-0.02	-0.15	0.000	11.7	0.30	19.86
T-bill	0.91	21.24	0.826	3.3	2.81	3.09
excess	-0.01	-0.06	0.000	8.3	0.13	20.16

Source: own calculations based on CRSP data

- Stock returns are basically just unpredictable.
 - $b = -0.02$ means that if returns go up 100% this year, you "expect" a decrease of 2% next year. Note that b is statistically not significant.
- T-bills are predictable
 - $b = 0.91$ and $t(b) > 2$.

The facts until the 80s

- Predictability of T-Bills does not mean that the market is inefficient
- If you know stock returns will be high next year, you can borrow money and invest in the market. If you expect T-bill returns to be high next year, you'd have to borrow at the same high rate to invest, which doesn't do you any good.
- This is why we run forecasting regressions for “excess returns”: $R_{t+1}^e = R_{t+1}^{stock} - R_{t+1}^{bill}$
- Until the 1980's thousands of similar regressions have been estimated, with different frequencies (days, weeks, months etc. – but always less than a few years).
- Surprising finding: excess returns are unpredictable!

Predictability and time-varying risk premia

- Subtle point: violation of the “random walk hypothesis” does not always mean market inefficiency.
- The required risk premia may change over time. For example, after market crashes, investors have lost wealth, financial institutions have lost equity capital and they may then require higher risk premia (remember the liquidity spiral).
- In this case expected returns vary over time and may be predictable.
- But this does not imply inefficiencies!
- Market efficiency correctly interpreted means “There is no free lunch”. I.e. in expectation everyone just earns a “fair” risk premium.
- If the “fair” risk premium changes over time, then excess returns may be predictable, even in efficient markets.

Two paths to progress

- Predicting longer-horizon returns, i.e. 3 – 5 years
- Using other forecasting variables $\mathbb{Z}$ (semistrong-form return prediction)
 - valuation ratios, dividend yield
 - interest rates and spreads
 - macroeconomic variables
 - many other variables such as equity vs debt issuance, net payout, consumption-wealth ratio, technical indicators
- General form of the regressions:

$$r_{t,t+k} = \alpha_k + \beta'_k \mathbb{Z}_t + \epsilon_{t+k}$$

In-sample vs out-of-sample


- The evidence for the predictability of stock returns comes primarily from in-sample predictive regressions
 - Econometric difficulties such as lack of exogenous and non-persistent regressors; overlapping observations
 - Concerns of data mining
- Out-of-sample tests use observations that are not used in the estimation of the statistical model itself
 - Out-of-sample analyses provide some protection against data mining
 - The evidence for return predictability gets much weaker (e.g., Goval, Welch, and Zafirov, 2024)
 - There is a debate that both in-sample and out-of-sample tests are susceptible to data mining

What does out-of-sample mean?


- Simulated (“pseudo”) out-of-sample (OoS) forecasts seek to mimic the “real time” updating underlying most forecasts
- Method splits data into an initial estimation sample (in-sample period) and a subsequent evaluation sample (OoS period)
- Forecasts are based on parameter estimates that use data only up to the date when the forecast is made
- As the sample expands, the model parameters get updated, resulting in a sequence of forecasts
- Why out-of-sample forecasting?
 - control for data mining - harder to “game”
 - feasible in real time

Main methods for OOS estimation

Expanding window



Rolling window



- Out-of-sample R^2 of **Campbell and Thompson (2008)** measures the proportional reduction in mean squared prediction error for the predictive regression forecast relative to the historical average benchmark:

$$R_{OS}^2 = 1 - \frac{MSE_{model}}{MSE_{bench}} = 1 - \frac{\sum_{t=1}^T (r_t - \hat{r}_t)^2}{\sum_{t=1}^T (r_t - \bar{r}_t)^2}$$

$\hat{r}_t$ and $\bar{r}_t$ are the fitted value from a predictive regression and the historical average, respectively, estimated through period $t - 1$.

- Diebold and Mariano (1995)** test for the equal predictive ability of forecasts. Recipe:

$$d_{t+1} = (y_{t+1} - \hat{f}_{1,t+1|t})^2 - (y_{t+1} - \hat{f}_{2,t+1|t})^2$$

Then run an OLS regression of the MSE difference, d_{t+1} , on a constant:

$$d_{t+1} = \mu + \varepsilon_{t+1}$$

Statistical vs. economic significance

Beyond the **statistical** significance of forecasts, what about their **economic** magnitude/relevance?

- **Sharpe Ratio gains:** build trading strategies based on some optimal portfolio allocation (e.g., maximize expected utility of a mean variance utility function) using the information coming from the predictive model vs. a benchmark (e.g., prevailing mean). Do you gain from using the forecasting model in terms risk-adjusted performance (e.g., Sharpe Ratio)?
- **Certainty Equivalent gains:** Assume a utility function (e.g., usually MV) and calculate the CE

$$CE = \bar{r}_p - \frac{\gamma}{2} \sigma^2(r_p)$$

for the portfolios constructed with and without the forecast of $\hat{r}_{p,t+1}$. Take the difference ($CE_{model} - CE_{bench}$). The CE gain can be interpreted as the fee an investor would be willing to pay to exploit the information in the predictor.

Return predictability: dividend yields

Dividend Growth of S&P 500


Period end	nom. DPS	real DPS
1925	3.4	-2.2
1935	-2.4	0.2
1945	3.5	0.6
1955	9.5	5.4
1965	5.2	3.4
1975	3.1	-2.5
1985	7.9	0.9
1995	5.7	2.2
2005	4.9	2.3
2015	6.9	5.0
2025	6.2	3.0

- The Table shows annualized growth rates in percent over 10-year periods.
- Past real dividend growth has been reasonably stable. Even in nominal values these growth rates were reasonably stable.

Own calculations.

Data source: [Prof. Robert Shiller](#)

Historical dividend-price ratios




Dividend yields and future returns

- The level of multiples provides important information about future expected rates of return:

Future 10-Year Real Rates of Return
When Stocks Are Purchased at Alternative Initial D/P
(from 1871 - 12/2015)

Quintile	Initial Dividend Yield Range	Real Return
Cheapest 20%	5.7% – 13.8%	10.6%
Second 20%	4.8% – 5.7%	7.2%
Third 20%	3.9% – 4.8%	6.4%
Fourth 20%	3.0% – 3.9%	5.7%
Most expensive 20%	1.1% – 3.0%	5.0%

Dividend yields and subsequent returns



Own calculations. Data source: Robert Shiller.

The red line is the dividend yield, the blue line is the cumulative return for the 7 years following the dividend yield.

- Result from John Cochrane (2011), Discount Rates, Journal of Finance 66 (4):

Table I
Return-Forecasting Regressions

The regression equation is $R_{t \rightarrow t+k}^e = a + b \times D_t/P_t + \varepsilon_{t+k}$. The dependent variable $R_{t \rightarrow t+k}^e$ is the CRSP value-weighted return less the 3-month Treasury bill return. Data are annual, 1947–2009. The 5-year regression t -statistic uses the Hansen–Hodrick (1980) correction. $\sigma[E_t(R^e)]$ represents the standard deviation of the fitted value, $\sigma(\hat{b} \times D_t/P_t)$.

Horizon k	b	$t(b)$	R^2	$\sigma[E_t(R^e)]$	$\frac{\sigma[E_t(R^e)]}{E(R^e)}$
1 year	3.8	(2.6)	0.09	5.46	0.76
5 years	20.6	(3.4)	0.28	29.3	0.62

Predictability: interpretation of regression output

- “5 Year” returns are a regression of the total return for the next 5 years on the initial dividend yield.
- 1% more dividend, 3.8% higher return over the next year!

Interpretation (1)

- The regression model

$$R_{t+1}^e = a + b \cdot (D_t/P_t) + \epsilon_{t+1}$$

means that at time t , the expected return is

$$\mathbb{E}_t [R_{t+1}^e] = a + b \cdot D_t/P_t$$

- “Returns are predictable” $\Rightarrow$ “Expected returns vary over time.”

Interpretation (2)

- The coefficient 3.8 means that in order to see the expected return at any point in time, just look at the D/P ratio and multiply by 3.8 (plus a constant).
- In the postwar era the DP ratio has varied from about $\approx 1\%$ to $\approx 5\%$, about 4 percentage points.
- That means that expected returns vary a lot! Magnitude:
 $4 \times 3.8 \approx 15$ percentage points. The long-term average expected return is only about 7 percentage points. Returns ARE predictable.

Long-run predictability (1)

- Long-run prediction seems to work even better.
- Reason: D/P ratios are fairly stable:

$$D/P_t = a + 0.94 \cdot D/P_{t-1} + \epsilon_t$$

- Define $x_t = D_t/P_t$

$$r_{t+1} = b \cdot x_t + \epsilon_{t+1}$$

$$x_{t+1} = \phi \cdot x_t + \delta_{t+1}$$

- The *two-year* regression implied by a pair of one-year regressions:


$$\begin{aligned} r_{t+1} + r_{t+2} &= (b \cdot x_t + \epsilon_{t+1}) + (b \cdot x_{t+1} + \epsilon_{t+2}) \\ &= b \cdot x_t + b \cdot (\phi \cdot x_t + \delta_{t+1}) + \epsilon_{t+1} + \epsilon_{t+2} \\ &= b \cdot (1 + \phi) \cdot x_t + (b \cdot \delta_{t+1} + \epsilon_{t+1} + \epsilon_{t+2}) \end{aligned}$$

- Similarly

$$r_{t+1} + r_{t+2} + r_{t+3} = b \cdot (1 + \phi + \phi^2) \cdot x_t + (\text{error})$$

Long-run predictability (2)

Why D/P forecasts long horizon returns



Source: John Cochrane, Lecture Notes in Asset Pricing

Predictability and market efficiency (1)

- Does the D/P ratio predict returns or dividend growth?

Horizon k (years)	$R_{t \rightarrow t+k}^e = a + b \frac{D_t}{P_t} + \epsilon_{t+k}$			$\frac{D_{t+k}}{D_t} = a + b \frac{D_t}{P_t} + \epsilon_{t+k}$		
	b	$t(b)$	R^2	b	$t(b)$	R^2
1	2.83	2.95	0.047	-0.19	-0.22	0.000
2	5.39	2.75	0.076	-0.99	-0.61	0.005
3	8.15	2.16	0.117	-1.27	-0.51	0.006
4	11.26	2.04	0.142	-1.38	-0.51	0.005
5	12.85	2.3	0.139	-1.93	-0.70	0.008

Own calculations based on CRSP data (1927 – 2024). t-values are obtained from Newey-West corrected standard errors. Dependent variables for horizons > 1 are obtained from compounding annual variables.

Predictability and market efficiency (2)

- Dividend growth should be forecastable. When prices are high relative to current dividends, this should be a signal that investors expect dividends to be higher in the future.
 - A regression of ΔD_{t+1} on P_t/D_t should give a positive coefficient; a regression of ΔD_{t+1} on D_t/P_t should give a negative coefficient.
 - As the table shows, the coefficients are not significant and R^2 's are basically 0!
- Variation in D/P ratios doesn't reflect much news about cash flow growth, but instead means that discount rates vary over time a lot.

Other variables to predict stock market returns

- **Valuation measures.** E.g., CAPE (cyclically adjusted price-earnings ratio), dividend yield plus buyback adjustment, earnings yield (sometimes compared to treasury yields), ratio of GDP to equity market capitalization.
- **Money and credit.** Most prominent: steepness of the yield curve. Other variables: short rate momentum, long rate momentum, broker-dealer leverage growth.
- **Real activity.** Examples for predictors are business confidence, the unemployment rate, historic real GDP growth.
- **Risk indicators.** Typically there is a contemporaneous relation, but results for prediction are weak.
- **Sentiment and expectations.** There is a negative predictive correlation between various sentiment indicators and equity returns. Lettau and Ludvigson (2001) propose the consumption/wealth ratio to back out return expectations.

Predictability of the stock market and market efficiency

- Predictability does not necessarily imply market inefficiency
 - Efficiency does not require that expected returns (= “required returns” = “discount rates”) remain constant over time.
- Question: do the forecasts correspond to a “rational variation in the equilibrium risk premium,” or to something else, like “irrational exuberance” or “bubbles”?
- If expected returns were 1% on Monday, Wednesday and Friday, but 13% on Tuesdays and Thursdays, we'd have a hard time cooking up a model of “equilibrium risk premium.”

Caveats

- Data mining (=fishing):
- Thousands of assistant professors plus hedge fund quants are running regressions similar to the ones above
 - We are bound to find some high t-stats
→ economic story is important!!!
- Peso problem:
 - Statistical relations may hide a small chance of a huge crash (lesson of the decade!)
- Estimation errors are likely to be large!
- Structural shifts (maybe the equilibrium D/P shifts over time)

Cross-sectional Evidence

Motivation: factor models for portfolio choice

- Investors want to construct mean-variance optimal portfolios. They might also want to take into account risk from other investments or liabilities.
- This means investors need estimates of assets' expected returns and the covariance matrix (or equivalently standard deviations and correlations).
- Remember the analytical solutions for the tangency portfolio

$$\mathbf{w}^* = \frac{\Sigma^{-1} \mathbb{E} \mathbf{r}^e}{\mathbf{1}' \Sigma^{-1} \mathbb{E} \mathbf{r}^e}$$

and the minimum variance portfolio

$$\mathbf{w}^{MVP} = \frac{\Sigma^{-1} \mathbf{1}}{\mathbf{1}' \Sigma^{-1} \mathbf{1}}$$

Reducing the complexity

- CAPM. The Capital Asset Pricing Model uses an equilibrium argument to identify the mean-variance efficient portfolio.
 - Important insights and easy to estimate expected returns, but many return anomalies.
- APT: The Arbitrage Pricing Theory is consistent with several risk factors earning risk premia.
 - Richer factor structure in the covariance matrix of risky returns
 - Nevertheless, important reduction in dimensions.
 - But what are the factors?

Reduction of dimension in factor models

- To obtain a full covariance matrix for n stocks, we need to estimate n variances and $n(n-1)/2$ covariances.
 - For $n = 50$, this equals 1,275 estimates.
- Risk decomposition in a single factor model:

$$\sigma_i^2 = \underbrace{\beta_i^2 \sigma_M^2}_{\text{systematic risk}} + \underbrace{\sigma_{e_i}^2}_{\text{idiosyncratic risk}}$$

- The covariance between assets is determined by their betas:

$$\sigma_{i,j} = \text{Cov}[\beta_i r_M^e + e_i, \beta_j r_M^e + e_j] = \beta_i \beta_j \sigma_M^2$$

- How many parameters do we need to estimate now for $n = 50$?
- Multifactor model:
$$\sigma_i^2 = \beta_i \Sigma_F \beta_i + \sigma_{e_i}^2 \text{ and } \sigma_{i,j} = \beta_i \Sigma_F \beta_j$$
 - How many parameters (2 factors, $n = 50$)?

Why might we need more factors:

Risks beyond the aggregate market return may matter

- Investors have background risk (e.g., from a job or holdings of private equity).
- Compare two stocks with the same beta:
 - Stock A does poorly in a recession, when the representative investor takes a hit from his non-marketable holdings.
 - Stock B does well in a recession.
- CAPM says we are indifferent between these two stocks.
- But if investors prefer stock B, it will have a higher price (hence lower expected returns). We might need to add a second risk factor related to the business cycle (recession risk).

Why might we need more factors: Cash flow and discount rate risks

- Why could stocks with the same beta be different over the business cycle?
- We know that the value of a financial asset is the discounted value of future expected cash flows.
- There are two sources of changes in stock prices:
 - news about expected future cash flows, and
 - news about discount rates.
- Campbell and Vuolteenaho (AER, 2004) argue that investors care more about the sensitivity to cash flow news (bad beta) than to discount rates (good beta).
 - Lower valuations due to higher discount rates are alleviated by improved investment opportunities, which is not the case with cash flows.

Why might we need more factors: Effects from investor heterogeneity

- Suppose some investors have portfolios that are biased towards large, fast growing firms.
 - This could be driven by less sophisticated investors who have a preference for glamorous firms.
- Then, other investors have to overweight the small, slow growth firms in their portfolio of risky assets.
 - This could be sophisticated investors who require compensation (i.e., higher expected returns) for doing so.
- No investor group holds the market portfolio. In fact, the market portfolio is no longer efficient.
 - Additional factors – value and size – can help in explaining expected returns.

Multifactor models & APT

- We now assume that the return data are generated by

$$r_i^e = \alpha_i + \beta_{i1}F_1 + \beta_{i2}F_2 + \cdots + \beta_{iL}F_L + e_i,$$

for all stocks $i = 1, \dots, N$. The $F_j, j = 1, \dots, L$ are uncorrelated with each other: $\rho_{F_m, F_n} = 0 \ \forall m, n$.

- E.g., F_1 is the excess return of the market index;
- E.g., F_2 is a portfolio that is long small stocks and short large market-cap stocks, constructed in a way to have zero correlation with the market.
- The residuals are uncorrelated with each other
 $\rho_{e_i, e_j} = 0 \rightarrow$ all asset correlation via factors!
- The residuals are uncorrelated with each factor:
 $\rho_{e_i, F_j} = 0$, i.e. $\text{Cov}[e_i, F_j] = \mathbb{E}[e_i (F_j - \mathbb{E}[F_j])] = 0$

Portfolio expected return and variance

- Consider portfolio P with weights $\mathbf{w}$, where elements of $\mathbf{w}$ are $w_i, i = 1, \dots, N, \sum_{i=1}^N w_i = 1$.
- The expected excess return is given by the weighted average of the securities' expected excess returns:

$$\mathbb{E}[r_P^e] = \mathbf{w}'\mathbb{E}[\mathbf{r}^e].$$

- The variance is given by

$$\sigma_P^2 = \mathbf{w}'\Sigma\mathbf{w}$$

Factor betas

- $\mathbf{B}$ is the $(N \times L)$ matrix of the assets' factor betas.
 - Each element (i, j) contains the beta β_{ij} of asset i to factor F_j .
 - Row i contains the vector β'_i of asset i 's betas to all factors.
- The portfolio's factor betas β_{Pj} to F_j are given by:

$$\beta_{Pj} = \sum_{i=1}^N w_i \beta_{ij} = \mathbf{w}' \beta_j$$

- The $(1 \times L)$ row vector β'_P of P 's factor betas is:

$$\beta'_P = (\beta_{P1} \ \beta_{P2} \ \dots \ \beta_{PL}) = \mathbf{w}' \mathbf{B}$$

- Portfolio variance is given by:

$$\begin{aligned} \sigma_P^2 &= \beta'_P \Sigma_F \beta_P + \mathbf{w}' \Sigma_e \mathbf{w} \\ &= \mathbf{w}' \mathbf{B} \Sigma_F \mathbf{B}' \mathbf{w} + \mathbf{w}' \Sigma_e \mathbf{w} \end{aligned}$$

Multifactor models example

Example with 2 assets and two factors

- Securities $i \in \{A, B\}$. Factors F_1 and F_2 .
- Information on betas and residual variance:

	β_{i1}	β_{i2}	σ_i^2
Security A	0.60	0.20	0.05
Security B	0.90	0.10	0.02

- Assume that $\sigma_{F_1}^2 = 0.12$, $\sigma_{F_2}^2 = 0.1$, and $\rho_{F_1, F_2} = 0$.
- Compute the variance of portfolio P with weights $w_A = 0.5$; $w_B = 0.5$.

Solution

Zero-beta portfolio

Often, we may want to isolate factor exposure to just one factor (or no factor at all) to have a “clean exposure”

- We can construct portfolios that have exactly the desired factor exposure.
- In particular, a zero-beta portfolio has factor exposures of zero.
 - Example: Single factor model, assets A and B, $\beta_A = 1.5$, $\beta_B = -0.5$.
 - With $w_A = 0.25$ and $w_B = 0.75$, we obtain $\beta_P = 0$.
- This works more generally in multifactor models. We have to set portfolio weights $\mathbf{w}$ such that $\mathbf{B}'\mathbf{w} = \mathbf{0}$ and $\mathbf{1}'\mathbf{w} = 1$.

Portfolio with a specific factor structure

- If we desire a portfolio with factor betas $\beta_P = \beta^*$, the conditions are

$$\mathbf{B}'\mathbf{w} = \beta^*$$

and

$$\mathbf{1}'\mathbf{w} = 1.$$

- To construct a long-short (zero-cost) portfolio, use instead

$$\mathbf{1}'\mathbf{w} = 0.$$

Constructing a zero-cost portfolio

Example with 3 assets and 2 factors

- Securities A, B, C . Factors $F1$ and $F2$.

Beta matrix:

	β_{i1}	β_{i2}
Security A	-0.400	1.750
Security B	1.600	-0.750
Security C	0.667	0.250

- How can you construct a zero cost portfolio P with $\beta_{P1} = 1$ and $\beta_{P2} = 0$?

Solution

Arbitrage pricing in equilibrium

Being able to easily compute covariances is nice – but we want to understand prices (and expected returns).

- **If** returns follow linear model (assumption) for $i = 1, \dots, N$:

$$r_i^e = \alpha_i + \beta_{i1}F_1 + \beta_{i2}F_2 + \dots + \beta_{iL}F_L + e_i,$$

- **Then:** Pricing in equilibrium: There exist risk prices λ_j for each factor F_j such that the expected return on any security i can be stated as:

$$\mathbb{E}[r_i^e] = \beta_{i1}\lambda_1 + \beta_{i2}\lambda_2 + \dots + \beta_{iL}\lambda_L,$$

- The theory puts no restrictions on the risk prices $\lambda_1, \dots, \lambda_L$.


Example: Arbitrage in expectations (1)

Assume there are well-diversified portfolios A , B , C , and D with the following characteristics:

	β_i	$\mathbb{E}[r_i]$ (%)
A	0.50	2.00
B	1.00	3.20
C	1.50	4.00
D	2.00	4.40

Example: Arbitrage in expectations (2)

Non-Feasible Risk-Return Relationship



Example: Arbitrage in expectations (3)


- Form a zero-beta portfolio $z1$ consisting of assets/portfolios A and C
- Form a zero-beta portfolio $z2$ consisting of assets/portfolios B and D
- If the expected returns of $z1$ and $z2$ differ, you can go long the one with the higher expected return and short the one with the lower expected return.
- There is no systematic risk, and if the portfolios are well-diversified, no unsystematic risk either.
- Hence, there would be an arbitrage opportunity, in contrast to our assumption.

Example: Arbitrage in expectations (4)

Portfolio	β_i	$\mathbb{E}[r_i]$
$z1 = 1.5A - 0.5C$	0.0	1 %
$z2 = 2B - 1D$	0.0	2 %

- Long $z2$, short $z1$!
- Hence, prices of B and C will increase, prices of D and A decrease.
- Expected returns of B and C will decrease, and of D and A increase.

Non-Feasible Risk-Return Relationship



Recovering risk prices

- We have seen that the expected excess return of asset i equals $\mathbb{E}[r_i^e] = \beta_i \lambda$.
- If we have several assets, we can write this conveniently in matrix form:

$$\mathbb{E}[\mathbf{r}^e] = \mathbf{B}\lambda.$$

- If $N = L$ (and $\mathbf{B}$ has full rank) we can solve for risk prices:

$$\lambda = \mathbf{B}^{-1} \mathbb{E}[\mathbf{r}^e].$$

- When we have a set of factor mimicking assets (i.e. portfolios or assets with unit exposure to exactly one factor), the corresponding beta matrix is the identity matrix. The risk prices are the expected excess returns of the factor mimicking portfolios.

Example

Recovering risk prices

- Suppose APT holds and consider the following information:

	β_{i1}	β_{i2}	$\mathbb{E}[r_i]$
Security A	1.0	1.6	16.2
Security B	3.0	2.8	21.6
riskfree asset	0.0	0.0	10.0

- What is the risk premium for factor 1?
- What is the risk premium for factor 2?

Solution

- *Pricing in equilibrium:* There exist risk prices λ_k for each factor k such that the expected return on any security i can be stated as:

$$\mathbb{E}[r_i] = \beta_{i1}\lambda_1 + \beta_{i2}\lambda_2 + \cdots + \beta_{iL}\lambda_L,$$

$$\forall i = 1, \dots, N$$

- Very general formulation.
- The theory does not actually tell you which factors to use!
- For a larger set of securities, you probably need more factors.

Implementing the APT

Suppose we want to test the model

$$\mathbb{E}[r_i^e] = \beta_{i1}\lambda_1 + \beta_{i2}\lambda_2 + \cdots + \beta_{iL}\lambda_L.$$

- Typically, to conduct such tests, one has to
 1. Decide on the factors and test assets,
 2. Estimate the betas, and
 3. Estimate the risk prices.
- Classic approaches in empirical asset pricing are
 - Time-series regressions,
 - Cross-sectional regressions, and
 - Fama-MacBeth regressions.

Time-series regressions

- This works only if the factors are returns.
- Follow the approach:
 1. Estimate the loadings $\{\beta_{ij}\}$ and the alphas α_i for each asset $i = 1, \dots, N$, on factors $j = 1, \dots, L$ in N time-series regressions with T observations each:

$$r_{i,t}^e = \alpha_i + \beta_{i1}F_{1,t} + \beta_{i2}F_{2,t} + \dots + \beta_{iL}F_{L,t} + e_{i,t}. \quad (1)$$

2. Estimate risk premiums by taking sample averages of the L factors:

$$\hat{\lambda}_j = \frac{1}{T} \sum_{t=1}^T F_{j,t}. \quad (2)$$

3. Test for statistical significance of the intercepts.
Ideally joint test: Gibbons, Ross, Shanken (1989).

Cross-sectional regressions

- Two-step procedure

1. Perform N time series regressions to estimate the betas:

$$r_{i,t}^e = \alpha_i + \beta_{i1}F_{1,t} + \beta_{i2}F_{2,t} + \cdots + \beta_{iL}F_{L,t} + e_{i,t}. \quad (3)$$

2. Use sample means as estimates for expected returns:

$$\bar{r}_i^e = \frac{1}{T} \sum_{t=1}^T r_{i,t}^e.$$


Estimate risk premiums by running a cross-sectional regression:

$$\bar{r}_i^e = \hat{\beta}_{i1}\lambda_1 + \hat{\beta}_{i2}\lambda_2 + \cdots + \hat{\beta}_{iL}\lambda_L + \eta_i. \quad (4)$$

Note 1: You use the estimated betas from first step as explanatory variables in second step to estimate the lambdas.

Note 2: Usually estimated without constant, but you could include one and test for significance.

Cross-sectional regression (Cochrane, Fig 12.1)



Fama-MacBeth procedure

1. Estimate the loadings $\{\beta_{ij}\}$
for each asset $i = 1, \dots, N$, on factors $j = 1, \dots, L$
in N time series regressions with T observations each:

$$r_{i,t}^e = \alpha_i + \beta_{i1}F_{1,t} + \beta_{i2}F_{2,t} + \dots + \beta_{iL}F_{L,t} + e_{i,t}. \quad (5)$$

2. Now estimate a cross-sectional regression for each of the T months with N observations each:

$$r_{i,t}^e = \hat{\beta}_{i1}\lambda_{1,t} + \hat{\beta}_{i2}\lambda_{2,t} + \dots + \hat{\beta}_{iL}\lambda_{L,t} + \eta_{i,t}. \quad (6)$$

3. Take averages $\hat{\lambda}_j = \frac{1}{T} \sum_{t=1}^T \hat{\lambda}_{j,t}$.

To obtain standard errors use $\sigma^2(\bar{x}) = \frac{1}{T}\sigma^2(x)$:

$$\text{Calculate } \sigma^2(\hat{\lambda}_j) = \frac{1}{T^2} \sum_{t=1}^T (\hat{\lambda}_{j,t} - \hat{\lambda}_j)^2.$$

Portfolios as Risk Factors

The Fama-French model

Fama and French (1993) propose the following **three-factor model** (FF3) for stock returns:

$$r_{i,t}^e = \alpha_i + \beta_{i,M} r_{M,t}^e + \beta_{i,SMB} r_{SMB,t}^e + \beta_{i,HML} r_{HML,t}^e + e_{i,t}.$$

- Implication of APT:

$$\mathbb{E}[r_i^e] = \beta_{i,M} \lambda_M + \beta_{i,SMB} \lambda_{SMB} + \beta_{i,HML} \lambda_{HML}.$$

- The expected return on stocks is only determined by the exposure to each of the factors.


The Fama-French risk factors

FF3 uses the returns on **portfolios** of traded assets as factors.

- In particular, the factors are **constructed** using six portfolios next to the market portfolio.
- Thereby, they focus on three risk factors known as
 - **Market** (M),
 - **Size** (SMB , or “*Small Minus Big*” market equity), and
 - **Value** (HML , or “*High Minus Low*” book-to-market).
- The **market factor** is defined as $r_M - r_f$, i.e., the excess return on the market, which is the value-weighted return on all NYSE, AMEX, and NASDAQ stocks (from CRSP) minus the one-month Treasury bill rate (from Ibbotson Associates).

Excess returns of 25 portfolios

Small → Big, Growth → Value



The six Fama-French portfolios

Market equity "Size" Median ME (NYSE)	Big Growth	Big Neutral	Big Value
	Small Growth	Small Neutral	Small Value
	30th B/M		70th B/M
	Book to Market "Value"		

Strategy to form the six portfolios

The details of the **portfolio construction** are:

- Portfolios are constructed at the end of each June.
- Intersections of
 - 2 portfolios formed on size (market equity, ME), and
 - 3 portfolios formed on the ratio of book equity to market equity (BE/ME).
- The size breakpoint for year t is the median NYSE market equity at the end of June of year t .
- BE/ME for June of year t is the book equity for the last fiscal year end in $t - 1$ divided by ME for December of $t - 1$. The BE/ME breakpoints are the 30th and 70th NYSE percentiles.

There is also an intercept α in the FF3 model, i.e.,

$$r_{i,t}^e = \alpha_i + \beta_{i,M} r_{M,t}^e + \beta_{i,SMB} r_{SMB,t}^e + \beta_{i,HML} r_{HML,t}^e + e_{i,t}.$$

- Only if factors are **traded returns**, then we can use α , the intercept in a time series regression of returns on the factors, as a measure of the “risk-adjusted excess return”.

Chasing α in 3-factor model is zero-sum game

- The representative investor's portfolio has to be the market.
- If not, the market would not clear. The entire supply of risky assets has to be held in equilibrium.
- Expected returns are such that she does just that.
- Hence, the representative investor's portfolio P_A has zero loadings on SMB and HML :

$$r_{P_A,t} - r_f = \alpha + 1 \times (r_{M,t} - r_f) + 0 \times r_{SMB,t}^e + 0 \times r_{HML,t}^e.$$

- The representative investor's portfolio has zero α because

$$r_{P_A,t} = r_{M,t}.$$

What does the FF3 model mean?

- Investors are rewarded for taking on **covariance** with *HML* and *SMB*
- Investors are **not** rewarded for investing in small stocks or value stocks.
- There could be large stocks that have large positive loadings on *SMB*, e.g., there could be high B/M stock that could have large negative loadings on *HML*.
- In this view of the world, *HML* and *SMB* capture other sources of macro-economic risk that affect the representative investor.


Fishing in the factor zoo

- Harvey, Liu, and Zhu (2016) argue that the usual t -statistic of 2 is not sufficient to justify proposing a new factor in asset pricing.
- They look at 313 papers from top journals proposing 316 factors!
- They classify factors according to the risk type into common factors (financial, macro, microstructure, behavioral, accounting, and other) and characteristics (financial, microstructure, behavioral, accounting, other).
- Introducing a new multiple testing framework, they propose t -values above 3.0.
- But qualitatively, a factor derived from a theory should have a lower hurdle than purely empirical factor fishing.

Factors

(Harvey, Liu, Zhu, Fig. 2)

Figure 2. Factors and publications



Some other factors

- The market return, i.e., the **CAPM** is usually included and extended.
- The Fama and French (2015) five-factor model, where they add **profitability** and **investment** factors.
- Jegadeesh and Titman (1993) uncovered the tendency for good or bad performance of stocks to persist over the following months. Based on this, **momentum** was added to factor models by Carhart (1997).
- Pástor and Stambaugh (2003) conclude that **liquidity** risk is a priced factor.
- The **q-factor** model by Hou, Xue, and Zhang (2015) (and the five-factor variant from Hou, Mo, Xue, and Zhang, 2021).