Portfolio Risk Management using Reverse Optimisation

Jason MacQueen

Problems with Optimisation

- Very sensitive to small changes in the inputs
  - Especially return forecasts
- Tends to error-maximise
  - See various papers by R.O. Michaud, FAJ etc.
- Almost always delivers counter-intuitive results
  - Constraints are usually required to produce an acceptable solution, so managers have little faith in “quant techniques”

Optimisation to an Engineer

- The basic process of optimisation is simple :

  RETURNS + RISKS => HOLDINGS

- But now check the tolerances involved :

  Very fuzzy ranges + Bit fuzzy estimates => 2 decimal places accuracy !?

- Note that the problem lies in the ***implementation***, not in the ***concept*** (which won Markowitz a Nobel prize and is still the standard paradigm for portfolio design).

R-Squared Ltd

Reverse Optimisation

- Rather than asking :
  - **What holdings make the portfolio efficient ?**
- It makes more sense (in engineering terms) to ask :
  - **What returns make the portfolio efficient ?**
- The ***concept*** is still the same, but now the process is :

HOLDINGS + RISKS => RETURNS

R-Squared Ltd

Implied Returns

- This process begins by analysing the **risk structure** of the portfolio, which shows whether the manager
  - has the right **amount** of risk overall
  - has the right **kinds** of risk
  - has the right **proportions** of risk.
- It then gives the **implied returns** required for efficiency, which can be compared with the managers's **expected ranges** to determine which holdings are not efficiently reflecting his views, given his forecasting ability.

Point Forecast of Return

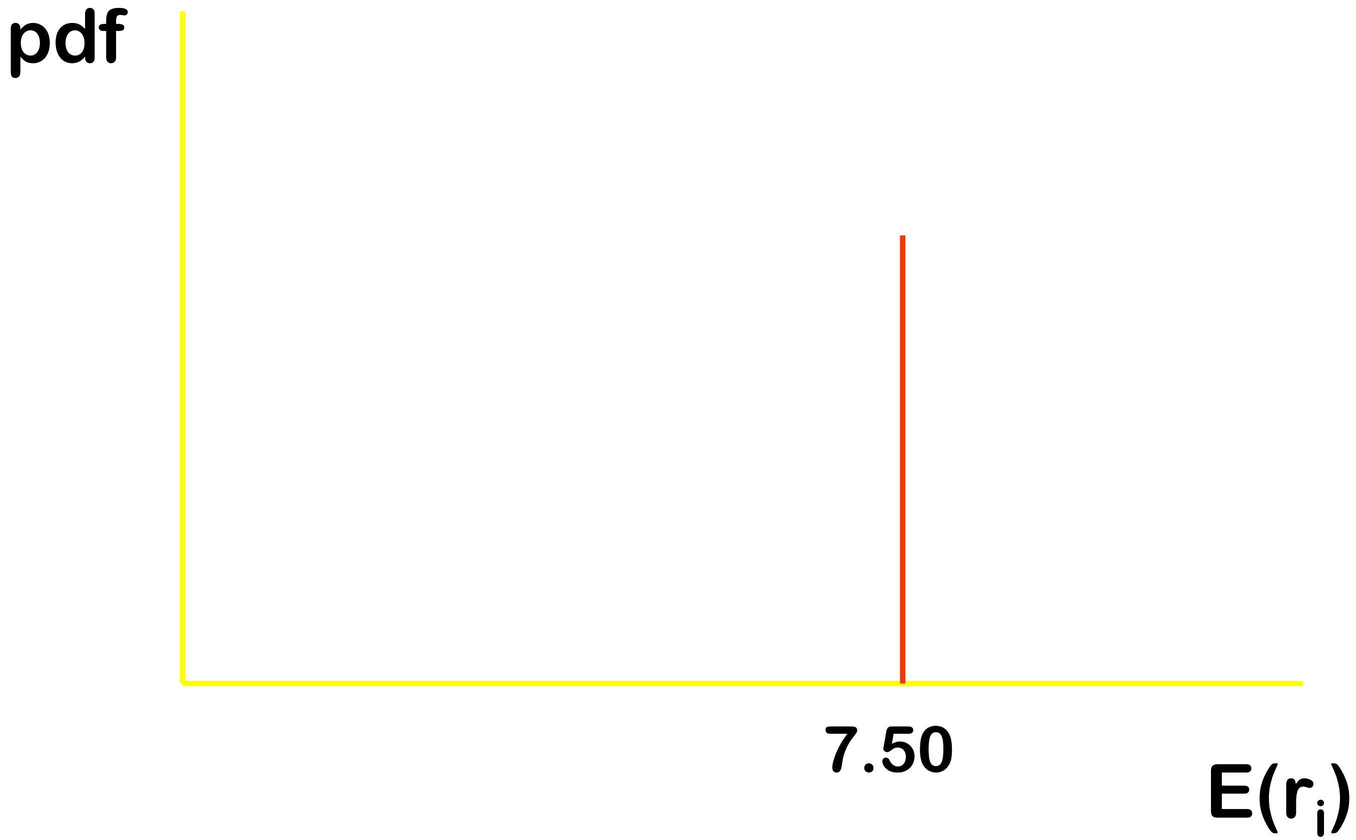

R-Squared Ltd

Range Forecast of Return

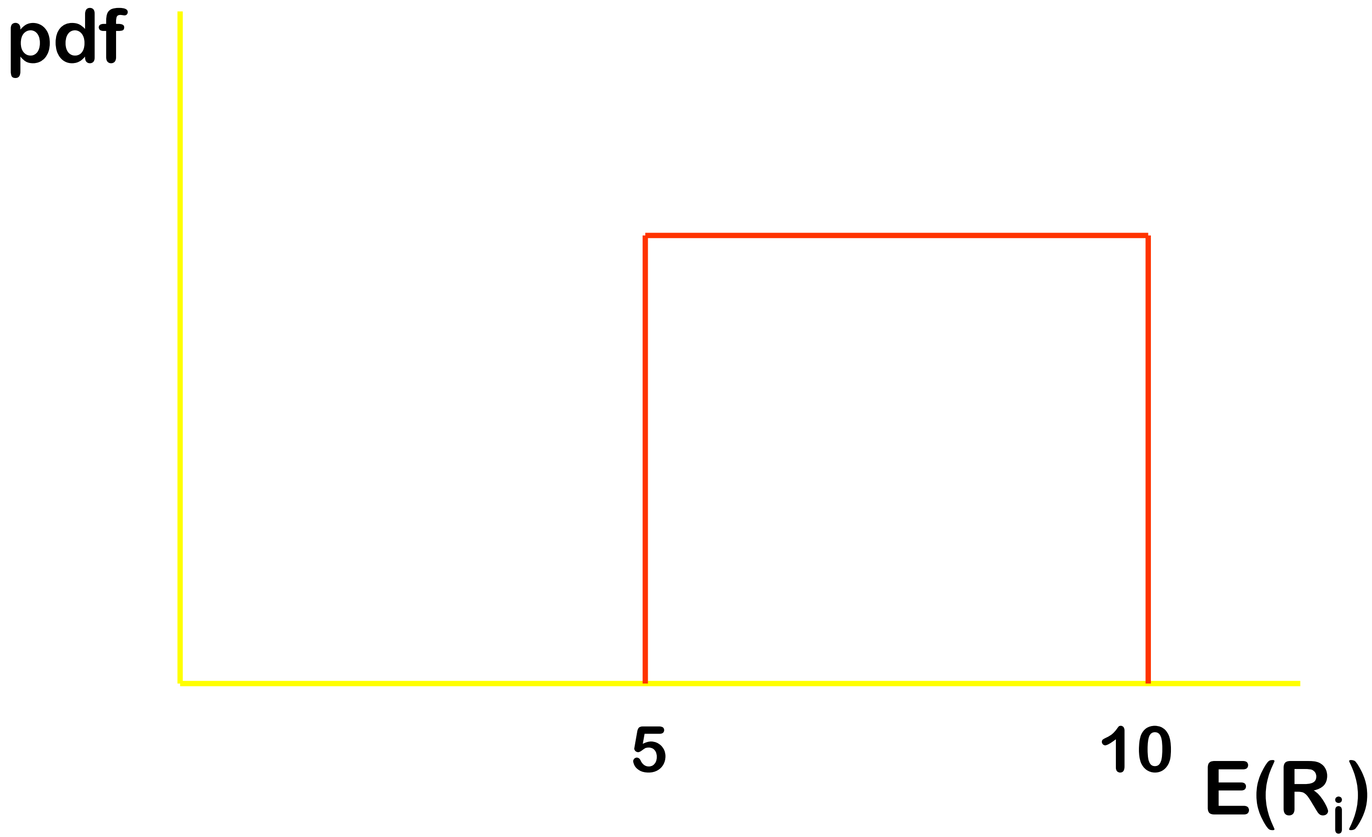

R-Squared Ltd

Actual Return Expectation

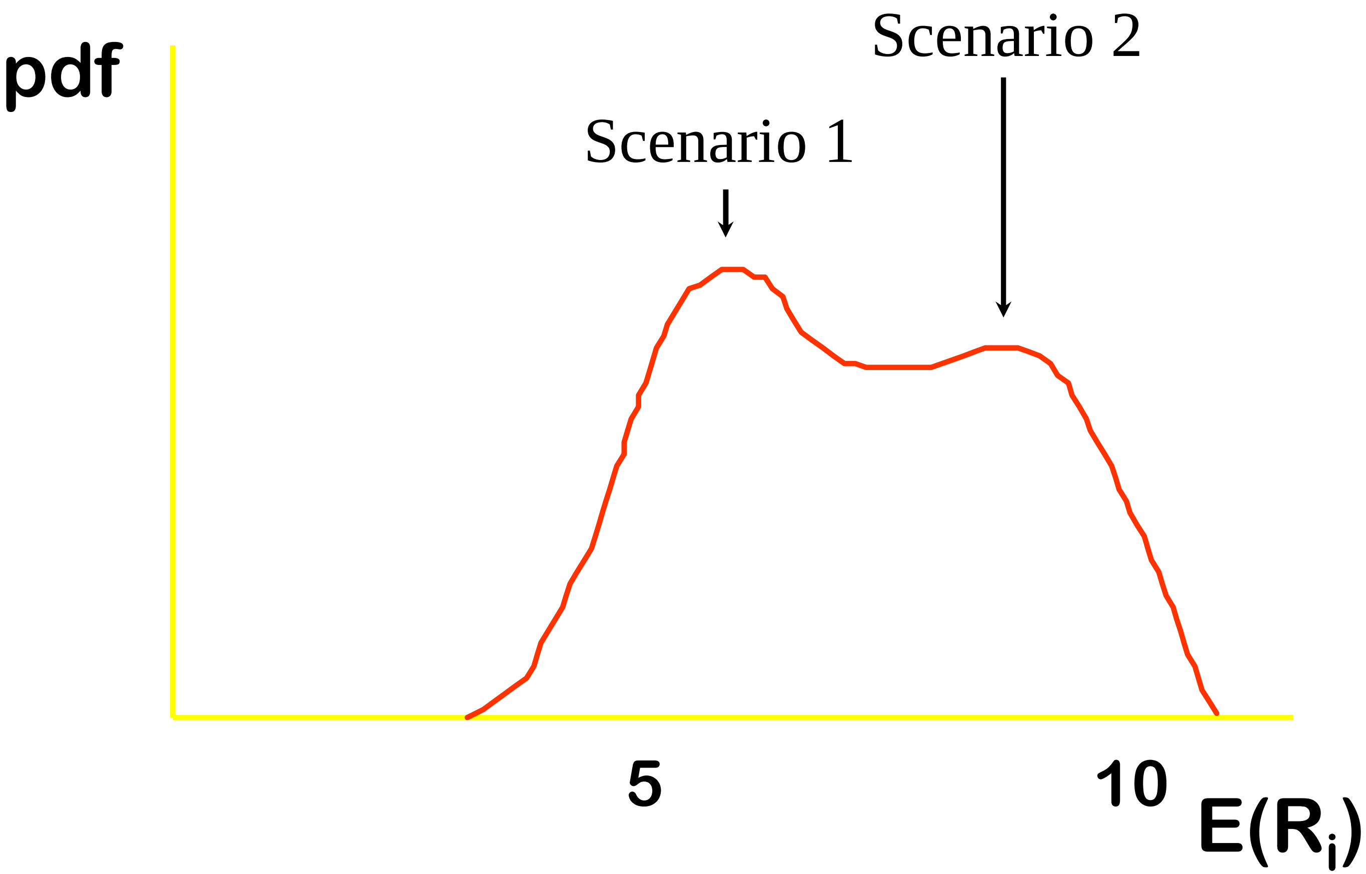

Point Forecast vs. Implied Return

Range Forecast vs. Implied Return

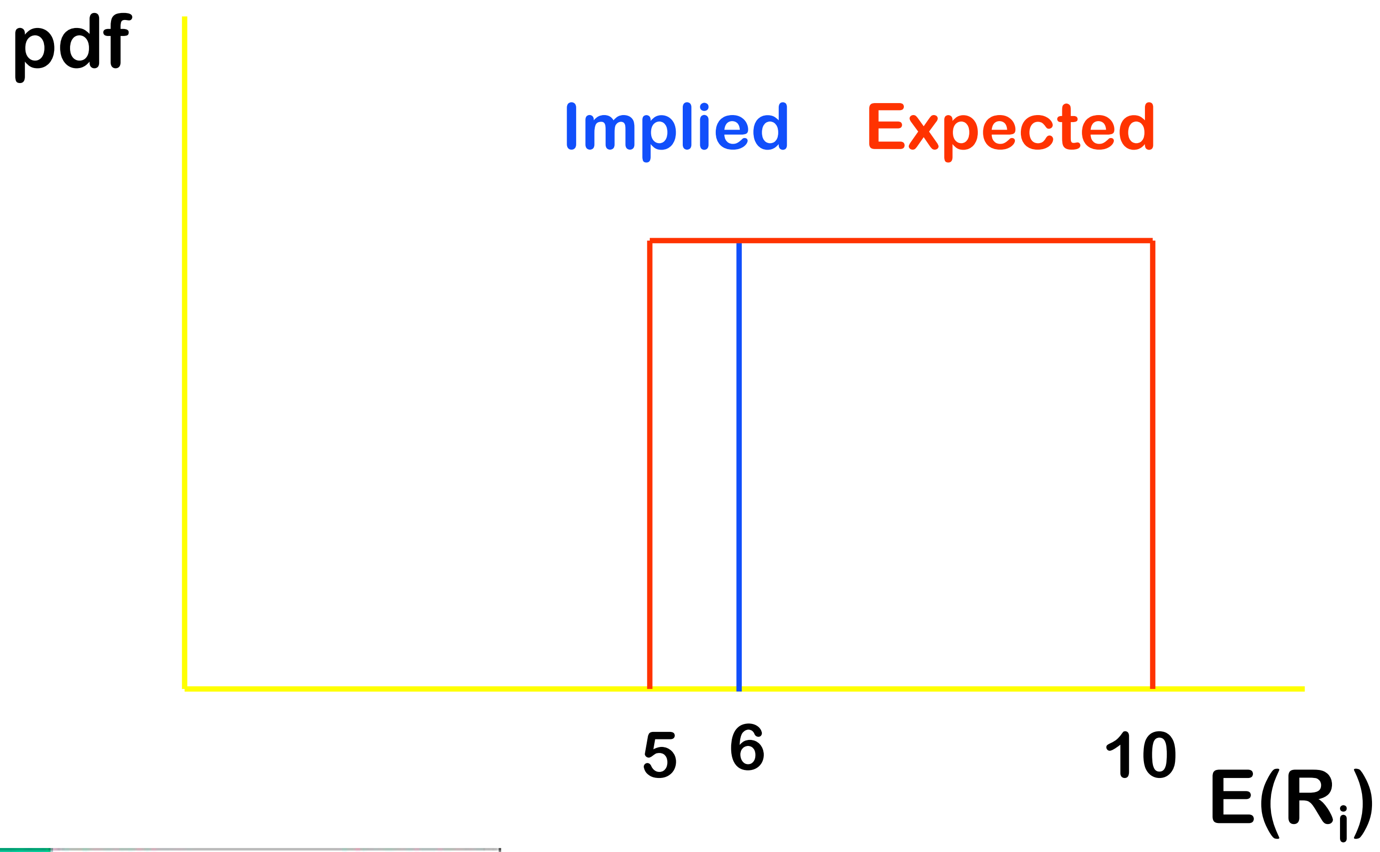

R-Squared Ltd

Actual Forecast vs. Implied Return

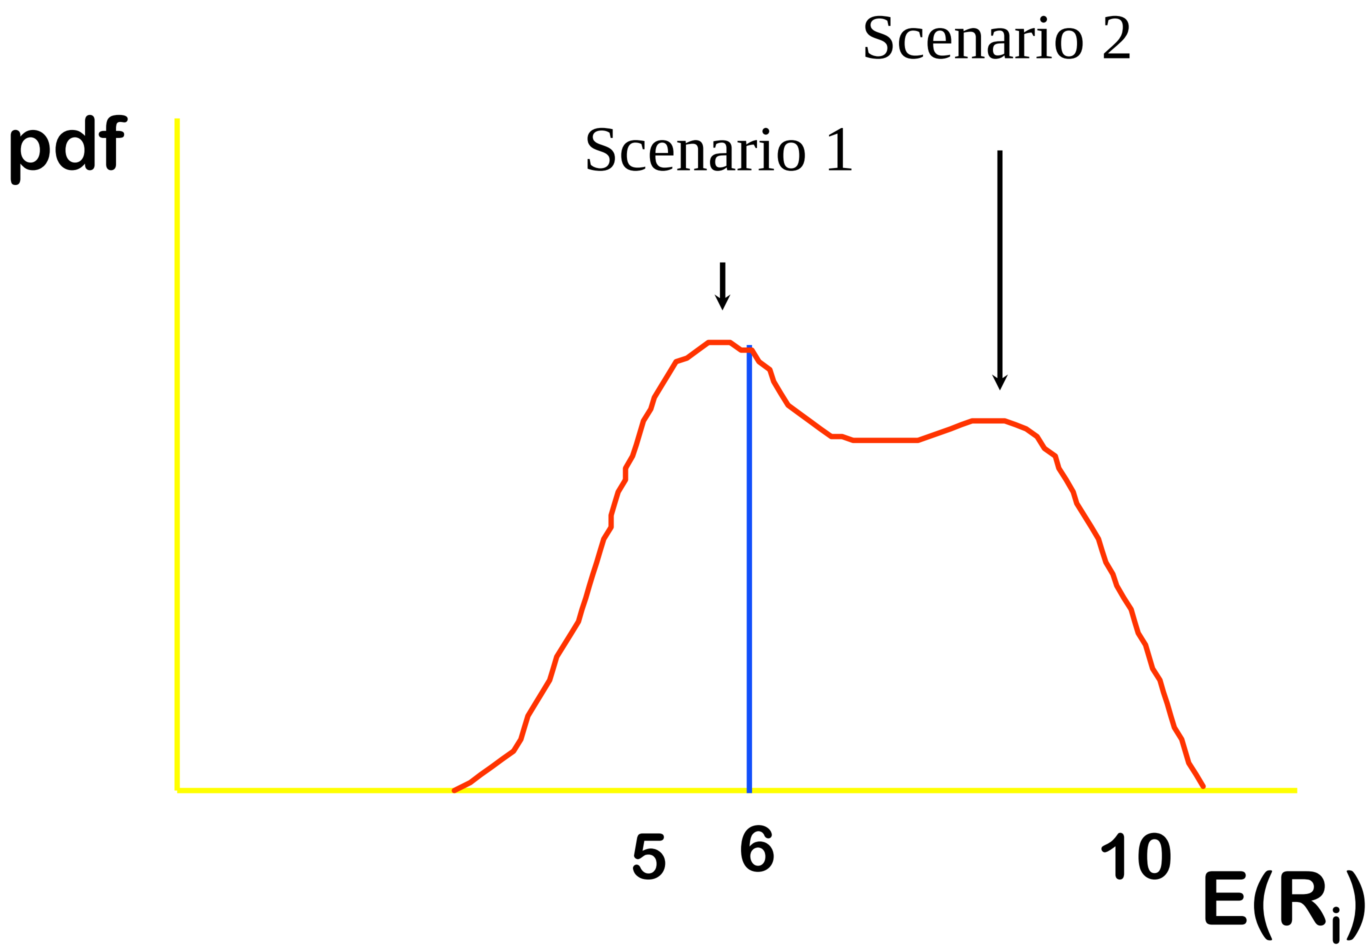

Practical Benefits of this Method

- Practical tool to identify which holdings contribute more or less than their “fair share” of risk.
- Allows fine tuning of portfolio holdings without the need for a full optimisation
- Allows use of imprecise or fuzzy ‘range forecasts’
- Incorporates use of <u>all</u> the manager’s information
- In the limit, managers don’t even need to quantify their return forecasts, as the analysis will generate an implied preference ranking of assets

The Risk Structure of a Portfolio

- We begin by breaking down the total variance of a portfolio into contributions from individual holdings
- We have

$$V_p = \sum_i^N \sum_j^N x_i x_j C_{ij}$$

- From which we derive individual contributions to variance as

$$ACV_i = \sum_j^N x_i x_j C_{ij}$$

R-Squared Ltd

Risk Contribution from each holding

$$ACV_i = \sum_j^N x_i x_j C_{ij} = x_i \sum_j^N x_j C_{ij}$$

$$= x_i \sum_j^N x_j \, cov(R_i, R_j) = x_i \, cov\left(R_i, \sum_j^N x_j R_j \right)$$

$$= x_i \, cov(R_i, R_P) = x_i C_{iP}$$

where Cip is the covariance of asset i with the portfolio P.

R-Squared Ltd

Derivation of Relative Imbalance

- Consider how risky each holding is relative to the whole portfolio
- Intuitively, if a 10% holding is contributing 10% of the portfolio risk, that is, in some sense, just its ‘fair share’ of risk, whereas a 5% holding contributing 10% of the portfolio risk is giving more than its ‘fair share’, and hence is a more extreme bet
- We define the Relative Imbalance of each asset as the ratio of its Percentage Contribution to Variance to its percentage holding size :-

$$Relative\ Imbalance_i = \frac{PCV_i}{x_i 100}$$

R-Squared Ltd

Meaning of Relative Imbalance

- Clearly, for the portfolio as a whole, 100% of the risk is coming from 100% of its holdings, so this ratio for the whole portfolio will be 1.
- Relative Imbalances for individual assets may be above or below 1, showing whether the bet on each asset is more or less extreme than average for the portfolio.
- These Relative Imbalances tell us how much return we should expect from each asset, and give us the implied ranking of assets in the portfolio

Relative Imbalance Unwrapped

- We can more clearly understand what the Relative Imbalance actually is by expanding the above expression, and then simplifying :

$$Relative\ Imbalance_i = \frac{PCV_i}{x_i 100} = \frac{ACV_i * 100}{x_i V_P * 100} = \frac{ACV_i}{x_i V_P}$$

$$= \frac{x_i C_{iP}}{x_i V_P} = \frac{C_{iP}}{V_P} = \beta_{iP}$$

R-Squared Ltd

Relationship to Implied Returns

- We may also derive an expression for the implied returns (required for portfolio efficiency) by starting with the definition of efficiency as maximising utility.
- This gives :-

$$I(R_i) = E(R_P) + \psi S_P(\beta_{iP} - 1)$$

where $I(R_i)$ is the Implied Return for each asset i.

The Formal Result

- This equation shows that, for efficiency, the expected return on an asset should be equal to the expected portfolio return (i.e. the average return) plus an adjustment that will be positive or negative depending on whether the asset is more risky $(\beta_{iP} > 1)$ or less risky $(\beta_{iP} < 1)$ than average.
- Note that the size of the adjustment depends on ψ , the incremental risk aversion parameter. The more risk averse we are, the larger the value of ψ , and hence the higher the returns required on the more risky holdings (and, of course, the lower the returns required on the less risky holdings).

R-Squared Ltd

Practical Applications - 1

- Ranking the assets in a portfolio from high to low by their Relative Imbalance immediately gives us the implied ranking of assets by their relative attactiveness
- This is particularly useful for managers who are unable (or unwilling) to quantify their return expectations
- For such managers, portfolio **inefficiency** consists of having their favourite stocks too far down the list, and their no-so-favourite stocks being near the top
- The analysis will usually suggest some obvious pairs trades to improve the overall efficiency of the portfolio

Practical Applications - 2

- For managers brave enough to quantify their forecasts, there are two obvious applications
- Entering a target Portfolio Return $E(R_p)$, and knowing the Portfolio Risk S_p, we can generate Implied Returns for any value of the Risk Aversion Parameter ψ .
- Alternatively, we can use this parameter to adjust the scale of the Implied Returns so that they are on the same scale as the actual Expected Returns, and then compare them directly

The Black-Litterman Methodology

- Essentially a special case of Reverse Optimisation
- **Assume** the World Markets portfolio is efficient (International Capital Asset Pricing Model).
- **Assume** that currency influences aren't a problem (Universal Currency Hedging Constant is valid and can be determined accurately enough).
- **Then,** given any known return such as US T-bills, we can derive the **implied equilibrium returns** for all other markets.

Implied and Equilibrium Returns

- **Implied returns** allow managers to use information about their expected ranges of returns rather than having point forecasts misrepresent their real views.
- **Equilibrium returns** give the manager additional information about the extent to which his views differ from those implied by theoretical considerations.
- Note that both the Implied and Equilibrium Returns effectively provide a **reality-check** on the consistency of the original forecasts.

R-Squared Ltd

Conclusion

- Portfolio managers often find the output of quantitative investment technology counter-intuitive or puzzling.
- The good news is that their intuition is often correct.
- But most of the time the problem is simply that their intuition has access to more information than has been used in deriving the strange result.
- The solution is to use as much information as possible.
- *Remember, managers are smarter than optimisers !!*

R-Squared Ltd