Università Commerciale
Luigi Bocconi

Performance Measurement and Attribution in Asset Management

Prof. Massimo Guidolin

Portfolio Management

Second Term 2019

Outline and objectives

- The problem of isolating skill from luck
- Simple risk-adjusted performance measures
- Performance decomposition
- The bad habits: multi-year return chasing (in the face of mean-reversion); underdiversification; seeking comfort
- One last, really dangerous bad habit: alpha-seeking as the cornerstone for portfolio construction

Key Concepts/1

- The goal of performance analysis is to distinguish skilled from unskilled investment managers, in an ex-post realized perspective
- The mere existence of positive returns does not prove skill because important issues remain: how much risk was taken on in generating that return?
- Two kinds of skill matter: ① Superior risk-adjusted returns can be derived through either superior timing or superior security selection; ② The ability to diversify the portfolio completely to eliminate all unsystematic risk, relative to some benchmark
- The first, key step is to adjust for the risk taken any reported time-weighted average performance
- Total risk is sensible for large, already well-diversified portfolios
- Systematic risk is instead sensible when the portfolio is smaller either in size or in breadth, or because constraints have been imposed on its overall diversification

Key Concepts/2

- Peer-group comparison performs a benchmarking with randomly drawn managers that select from identical asset menu and are subject to identical constraints
- Risk-adjustment is applied by selecting peers that have similar risk exposures, total and/or systematic
- The Sharpe ratio is $$SR_{FMV} \equiv \frac{(T^{-1}\sum_{t=1}^{T} R_{FMV,t}) - R^f}{\sqrt{T^{-1}\sum_{t=1}^{T}[R_{FMV,t} - (T^{-1}\sum_{t=1}^{T} R_{FMV,t})]^2}}$$ and it is appropriate for well-diversified, large ptfs.
- The Capital Market Line Alpha is the plus/minus of a manager vs. a naïve (leveraged) investment in the market ptf., say an index mutual fund
- In the plot, manager A has generated a plus

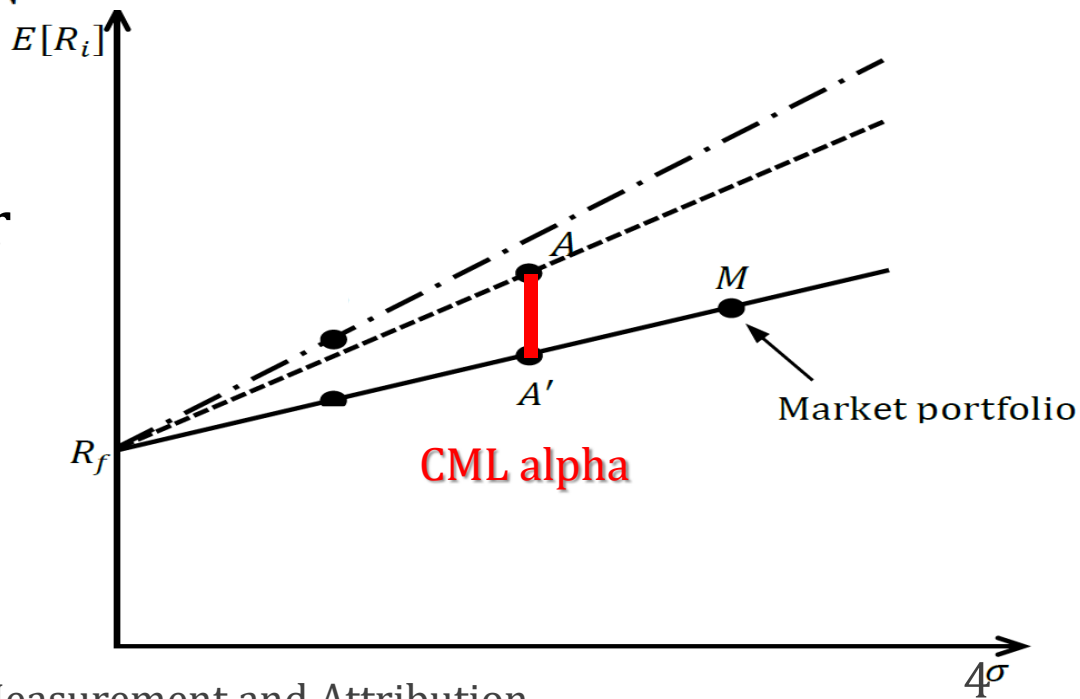

Key Concepts/3

- Treynor ratio is the analog of Sharpe scaling by systematic risk:

$$TR_{FMV} \equiv \frac{(T^{-1}\sum_{t=1}^{T} R_{FMV,t}) - R^f}{\hat{\beta}_{FMV}}$$

- Unfortunately, Treynor relies rather heavily on the CAPM
- Jensen's alpha extends CML alphas to a CAPM world: the return difference between the actual mean realized return of a given ptf. and the return that would be expected if the ptf. were on the SML:

$$\hat{\alpha}_{FMV}^{JENSEN} \equiv T^{-1}\sum_{t=1}^{T}\left(R_{FMV,t} - R^f\right) - \hat{\beta}_{FMV}\left[T^{-1}\sum_{t=1}^{T}(R_t^m - R^f)\right]$$

- Fama's performance decomposition is based on the identities:

Overall Performance = Mean Excess Return = (Portfolio Risk) + (Selectivity)

Overall Performance = [(Investor's Risk) + (Manager's Risk)] + (Selectivity)

- A net selectivity component reveals what part, if any, of a ptf. return comes from selecting securities in % different from the market
- The diversification contribution comes from the fact that in order to increase performance, a manager tends to under-diversify

Skill vs. Luck, that's the question…

- We now adopt and ex-post perspective: was asset management successful, given assigned goals?
  - Assigned goals == preferences of the investors, presumed or effective
- The goal of performance analysis is to distinguish skilled from unskilled investment managers. But, how do you tell them apart?
  - In a population of 100 investment managers, say 5 percent, or 5, should have exceptional performance by chance alone
  - None of the successful managers will admit to being lucky; all of the unsuccessful managers will cite bad luck
  - Those with both skill and luck are blessed; they deserve to thrive, and they will
  - Those with neither skill nor luck are doomed and will be wiped out by natural selection. But what about the two other categories?
  - Managers with skill but no luck will be unjustly expelled from the industry because their historical performance will not reflect true skill
- The mere existence of positive returns does not prove skill because important issues remain: how much risk was taken on in generating that return?

Two types of rewarded skills

- Two kinds of skill:

  ① Superior risk-adjusted returns can be derived through either superior timing or superior security selection

  ② The ability to diversify the portfolio completely to eliminate all unsystematic risk, relative to the portfolio's benchmark

  - When restrictions have been imposed on the operations pursued by delegated ptf. management (e.g., social responsibility limits to the asset menu, or specialization constraints in terms of asset classes, sectors, countries, etc.), performance measurement represents a crucial step to understand the realized, ex-post costs deriving from such restrictions
- The evaluation of portfolio performance is essentially concerned with comparing the return earned by some portfolio with the return earned on one or more benchmarks
- It is important that ptfs chosen for comparison are truly comparable, by carrying similar risk and being bound by similar constraints
- A first important choice consists of whether total or systematic risk should be used to perform any risk-adjustment

Two types of rewarded skills

- Even if the money manager herself may be risk-neutral, it seems appropriate that risk discounting be applied
- Total risk is sensible for large, already well-diversified portfolios
  - For instance, sovereign and endowment funds may easily be ranked in this way because they will normally find very little comfort in the fact that part of the risk could be diversified away if they held other assets when the portfolio under consideration contains their total asset
- Total risk means that indicators based on (say) portfolio variance or standard deviation, for instance, the Sharpe ratio, should be used
- Systematic risk is instead sensible when the portfolio is smaller either in size or in breadth, or because constraints have been imposed on its overall diversification
  - For instance, a specialized domestic small cap equity fund faces limits to diversification; risk-adjusting on the basis of its market beta or the set of its betas estimated vs. a range of well-recognized factors is sensible
  - A first, simple technique works in the following way: a synthetic portfolio of assets having approximately the same amount of risk is randomly built by drawing securities for a given, wide universe of assets

Peer-group comparisons

- Benchmark must incorporate the same restrictions that applied ex-ante
- The performance of the ptf. under consideration is then compared to the performance—over the same period and horizon—of the average across the a large number of random benchmarks that have been constructed
- E.g., suppose that Aldebaran has achieved a 5-year annualized return of 9%, has a total annualized standard deviation of 22%, and a beta of 1.3
- Form 10,000 random portfolios of international domestic stocks—the same asset menu from which Aldeberan has been selecting—that have approximately a standard deviation of 22% and a beta of 1.3
- Suppose that the average time-weighted return across such 10,000 portfolios is 7.5%: 0.5% represents a management plus by Aldebaran

- This technique is called peer-group comparison
- Apart from direct comparisons, there are four different one-parameter performance measures that have been proposed in the literature:

① Sharpe ratio ② Capital Market Line Alpha

③ Treynor ratio ④ Jensen's Alpha

- The first two are total risk measures, to be analyzed in the mean-variance space

Sharpe ratio and CML alpha

- The **Sharpe ratio** is defined as:

$$SR_{FMV} \equiv \frac{(T^{-1}\sum_{t=1}^{T} R_{FMV,t}) - R^f}{\sqrt{T^{-1}\sum_{t=1}^{T}[R_{FMV,t} - (T^{-1}\sum_{t=1}^{T} R_{FMV,t})]^2}}$$

  - Sharpe ratio represents the slope of the capital market line and it measures the most favorable, achievable tradeoff between risk premia and risk—as measured by total portfolio standard deviation
  - The Sharpe measure looks at the decision from the point of view of an investor choosing a ptf. to represent the majority her investments
  - If an investor desired a risk different from that offered by a given FMV, she would modify the risk by lending and/or borrowing

- The **CML alpha** is the percentage return difference between the actual mean realized return of a given ptf. and the expected return that would be anticipated if the ptf. were efficient and fell on the CML:

$$\alpha_{FMV}^{CML} \equiv (T^{-1}\sum_{t=1}^{T} R_{FMV,t} - R^f) - \frac{(T^{-1}\sum_{t=1}^{T} R_t^m) - R^f}{\sqrt{T^{-1}\sum_{t=1}^{T}[R_t^m - (T^{-1}\sum_{t=1}^{T} R_t^m)]^2}} \times \sqrt{T^{-1}\sum_{t=1}^{T}[R_{FMV,t} - (T^{-1}\sum_{t=1}^{T} R_{FMV,t})]^2}$$

Treynor ratio

- CML alpha measures the plus/minus vs. a naïve strategy of investing in the market ptf.
  - CML alpha > 0 implies that in mean-standard deviation space, a given portfolio lies above the CML
  - Only ptfs. with positive CML alpha should be bought
  - In the figure, two pts. receive a different ranking from the Sharpe ratios and CML alpha criteria
  - $SR_B > SR_A$, but the CML alpha of A exceeds that of B

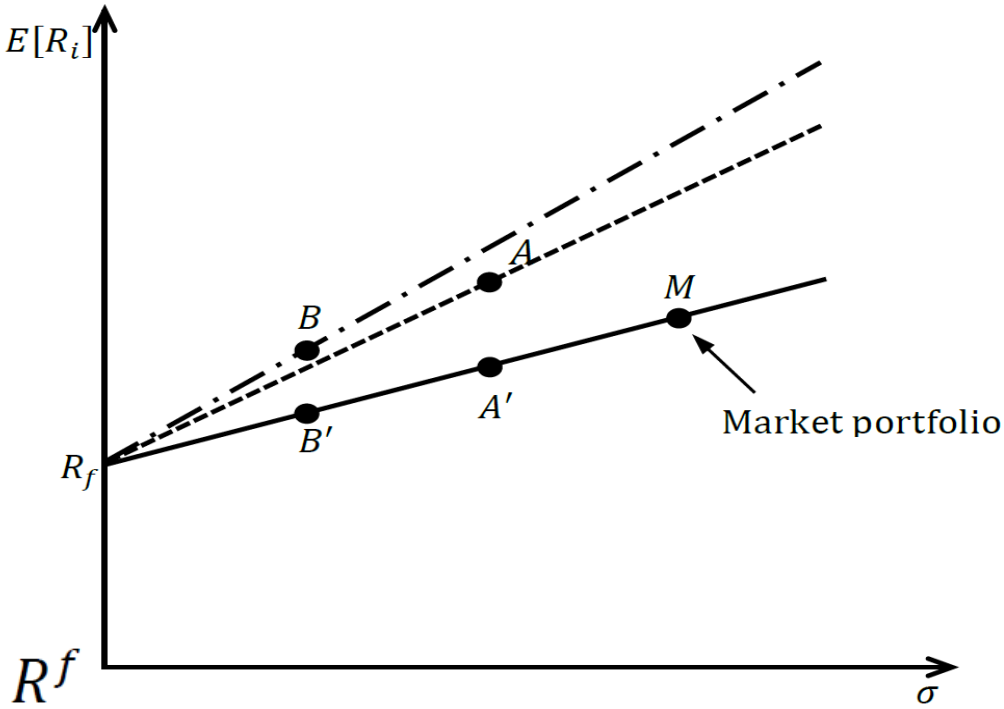

- The Treynor ratio is:

$$TR_{FMV} \equiv \frac{(T^{-1}\sum_{t=1}^{T} R_{FMV,t}) - R^f}{\hat{\beta}_{FMV}}$$

  - The logic of is that a ptf. should be as close as possible to the Security Market Line, ideally positioned on it or even above it
  - The expression measures the portfolio's risk premium per unit of systematic risk

Jensen's alpha

- Comparing a portfolio's TR value to a similar measure for the market portfolio—which equals the market risk premium as the market's beta equals 1—indicates whether the portfolio would plot above the SML
- The Treynor measure looks at the decision from the point of view of an investor choosing a ptf. that represents a portion of her financial wealth, and for which the overall contribution to systematic risk needs to be as efficient as possible, i.e., rewarded in a way consistent with CAPM

- Jensen's alpha extends CML alphas to a CAPM/APT world: the return difference between the actual mean realized return of a given ptf. and the return that would be expected if the ptf. were on the SML:

$$\hat{\alpha}_{FMV}^{JENSEN} \equiv T^{-1}\sum_{t=1}^{T}\left(R_{FMV,t} - R^f\right) - \hat{\beta}_{FMV}\left[T^{-1}\sum_{t=1}^{T}\left(R_t^m - R^f\right)\right]$$

- Jensen's alpha > 0 implies that in risk premium-beta deviation space, a given portfolio lies above the SML
- Jensen's alpha can also be directly estimated and its statistical significance tested, when it is written as the intercept in a regression:

$$\left(R_{FMV,t} - R^f\right) = \alpha_{FMV}^{JENSEN} + \beta_{FMV}\left(R_t^m - R^f\right) + \epsilon_{FMV,t}$$

- Extensions to multi-factor, APT-style return models are obvious

The information ratio

- Another widely used performance measure is Goodwin's (1998) information ratio (also known as an signal-to-noise ratio)
- It measures a portfolio's average return in excess of that of a benchmark portfolio divided by the standard deviation of this excess return:

$$IR_{FMV} \equiv \frac{\hat{\alpha}_{FMV}^{JENSEN}}{\sqrt{T^{-1}\sum_{t=1}^{T}[(R_{FMV,t} - R^f) - \hat{\beta}_{FMV}(R_t^m - R^f)]^2}}$$

- Sometimes the numerator is simply stated as the difference between the sample mean excess return on a portfolio and the same quantity for some appropriate benchmark, the average tracking error
- Goodwin (1998) has noted that the Sharpe ratio is a special case of the IR where the risk-free asset is the benchmark portfolio and should be zero for passively managed portfolios

- Another important tool consists of performance decompositions: separating skill from luck and different types of skill
- The commonly employed performance decomposition is Fama's (1972): basic premise is that overall performance of a ptf., can be decomposed into measures of risk-taking and security selection skill

Fama's performance decomposition

Overall Performance = Mean Excess Return = (Portfolio Risk) + (Selectivity)

- If there is a difference between the risk level specified by the investor and the actual risk level adopted by the portfolio manager (in cases where these are separate individuals), this calculation can be further refined to

Overall Performance = [(Investor's Risk) + (Manager's Risk)] + (Selectivity)

- Skill = selectivity component, portion of the portfolio's actual return beyond that available to an unmanaged portfolio with identical systematic risk
- Luck = various forms of risk, compensated
- The Jensen's alpha for some ptf. A is the height above the SML, A – A'
- This is (total) return from selectivity
- However, A and A' do not have the same total risk: the risk of portfolio A' is completely non-diversifiable

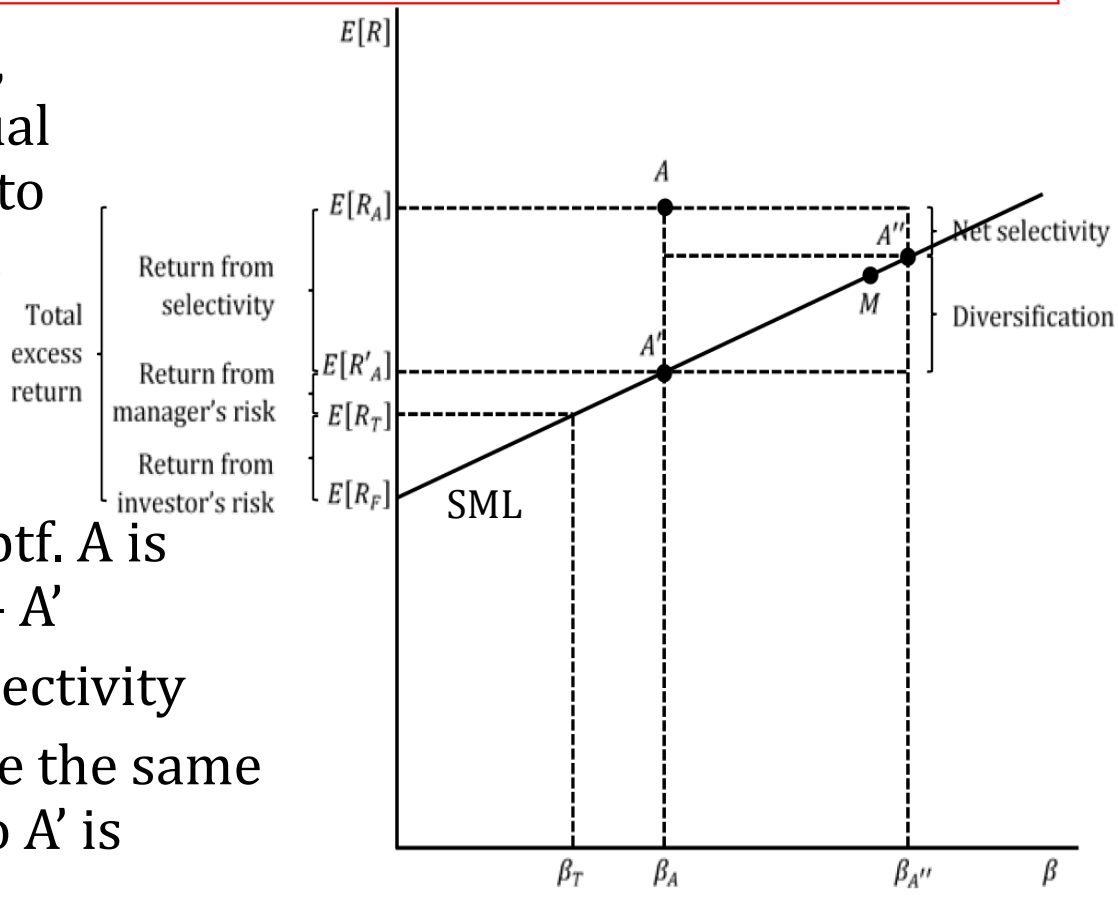

Fama's performance decomposition

- Is the extra return measured by the difference A – A' is worth the extra risk?
- Compare A with a portfolio with a ptf. with the same total risk, A''
- Only calculations can isolate the coordinates of A'', i.e., you are not supposed to spot it
- But A'' is efficient and hence on the SML
- If A'' has the same risk as A, that has greater total risk than portfolio A' has, then

$$Var[R_{A,t}] = \beta_{A''}^2 Var[R_t^m]$$

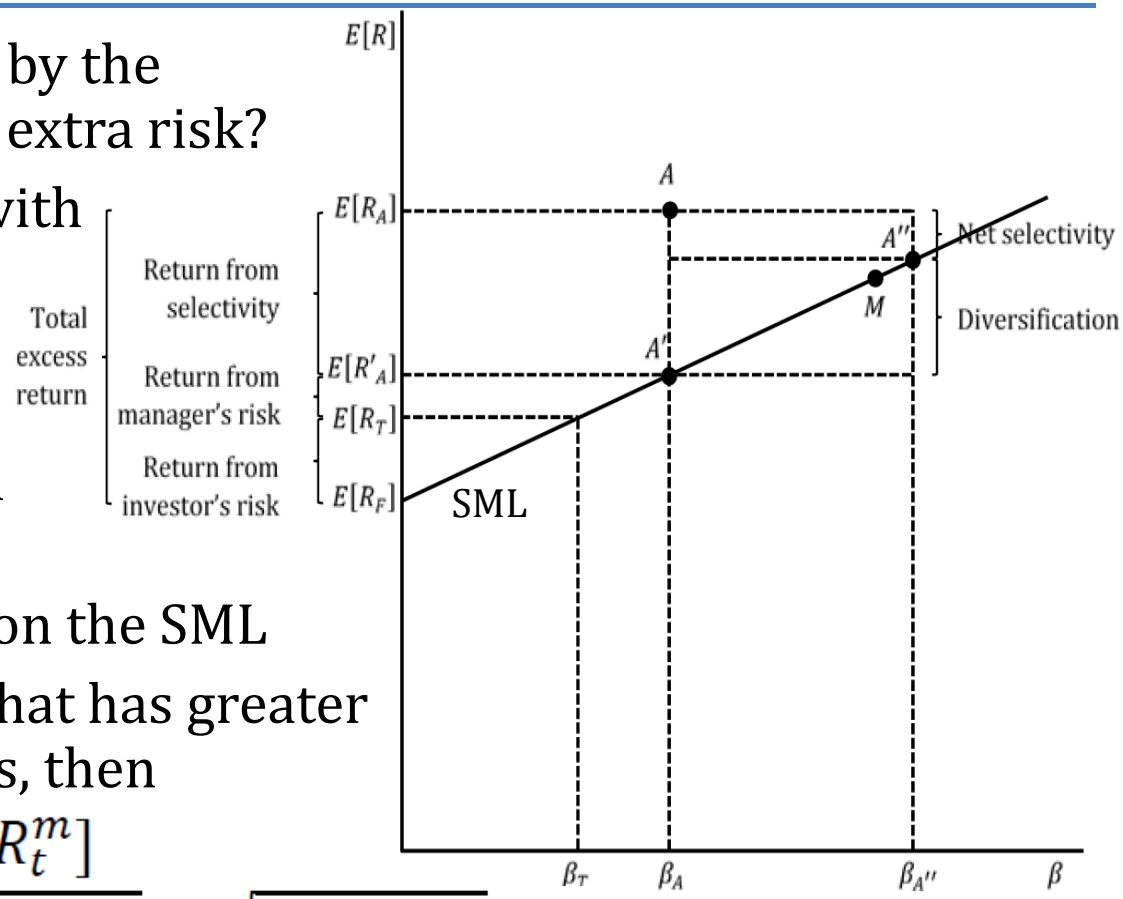

- This implies that

$$\beta_{A''} = \sqrt{\frac{Var[R_{A,t}]}{Var[R_t^m]}} > \sqrt{\frac{Var[R_{A',t}]}{Var[R_t^m]}} = \beta_A$$

- Hence portfolio A'' is located to the right of the initial ptf. A
- This show that

$$\underbrace{E[R_{A,t}] - E[R_{A',t}]}_{Total\ selectivity} = \underbrace{(E[R_{A,t}] - E[R_{A'',t}])}_{Net\ selectivity} + \underbrace{(E[R_{A'',t}] - E[R_{A',t}])}_{Diversification}$$

Fama's performance decomposition

- Net selectivity reveals what part, if any, of the managed ptf. return comes from selecting securities in % different from the market
- The diversification contribution comes from the fact that in order to increase performance, a manager tends to under-diversify
- A ptf. manager that over time tries and gets a large contribution from net selectivity is said to be an active manager in the cross section
  - The second part of a Fama-style decomposition occurs to the left
  - This further decomposition is possible only if the customers have specified a desired level of market risk, which is typical of pension funds
  - Assume the manager had been assigned a systematic target risk level, say, β_T
  - The difference in return between E[RA'] and E[RT] is the extra return due to the manager's risk appetite ("reach for yield")

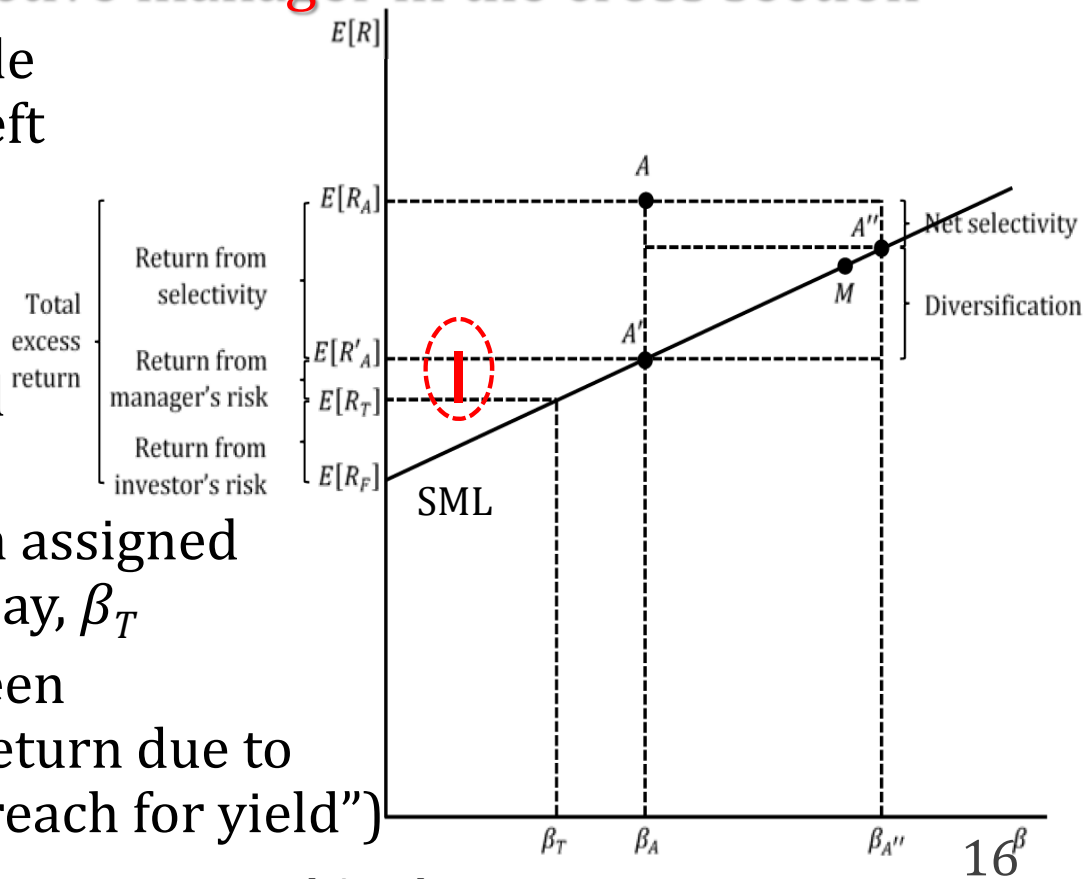

Tactical asset allocation

- A portfolio manager that over time receives a large contribution to her performance from risk reward is said to be an **active manager in the time series dimension**
- She would be actively changing over time the amount of risk exposure around a define mandate/target level β_T
- **Tactical asset allocation** (TAA) is a ptf. management strategy that aims at producing active surplus returns solely through allocation decisions
  - TAA managers adjust their asset class exposures based on perceived changes in the relative valuation of those classes
  - A typical TAA fund shifts money between three asset classes—stocks, bonds, and cash equivalents—although many definitions of these categories (e.g., large cap versus small cap, long term versus short term) are also used in practice
  - Sometimes, ptf. managers do "beta-timing": when the stock market is expected to surge (decline), the manager increases (decreases) her portfolio beta to obtain a portfolio with a greater responsiveness to market changes

Bad Habits and Good Practices

- The previous notions seem to set matters in a stark way: good portfoli managers ought to report high, competive Sharpe ratios, Treynor and information ratios, and especially Jensen's alphas deriving from selectivity and or/market timing
- Goyal, Ilmanen, and Kabiller (2015) warn us that some of these objectives may easily turn into "bad habits"
  - Of course, for each bad habit there is a contrasting good practice

① Chasing multi-year returns: many investors tend to buy multi-year winners and sell multi-year laggards—whether asset classes, strategy styles, single stocks, or funds

- This is not surprising, as the human tendency to extrapolate is one of our strongest behavioral biases
- While multi-year return chasing is harmful to ptf performance, evidence suggests that chasing winners over the past few months may be profitable (momentum) over monthly horizons up to a year
- Financial markets tend to exhibit more mean reversion at multi-year horizons, as opposed to the shorter-term continuation patterns

Bad Habits and Good Practices

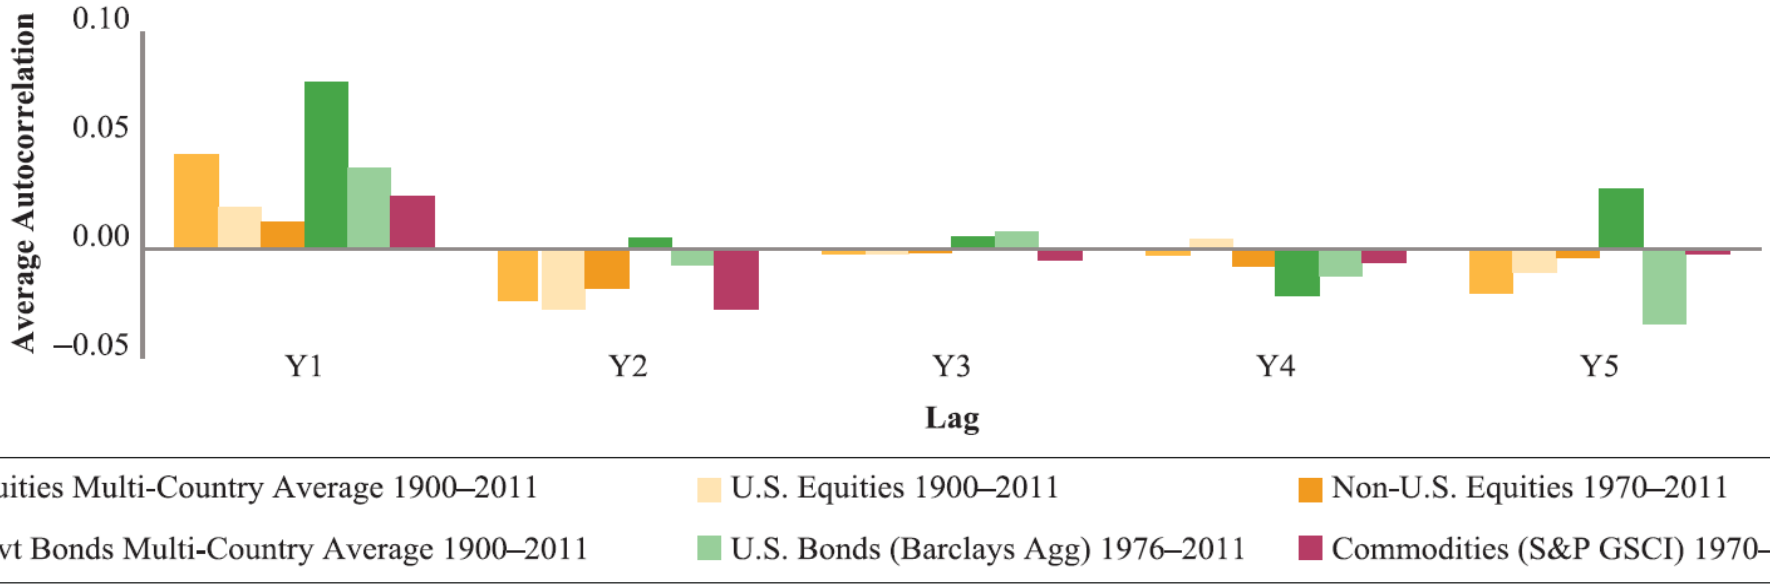

- Unfortunately, it is at this horizon that reallocation decisions tend to be made, making us momentum investors at reversal horizons
- Such pro-cyclicality for institution-nal investors at 3- to 5-year horizons may reflect typical performance evaluation periods

One-Year Momentum and Long-Term Reversal Patterns in U.S. Stock Returns, 1931–2013

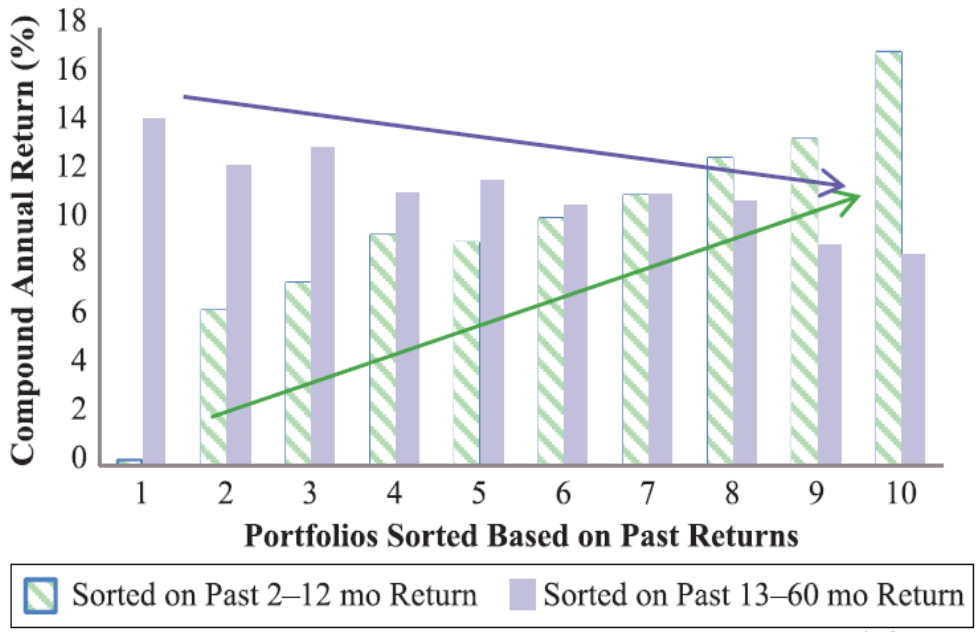

Bad Habits and Good Practices

- The best-known indirect evidence of procyclic investing's harmful effects comes from the gap between time-weighted investment returns and dollar-weighted investor returns
  - Studies show that the average returns investors historically experience are lower than the avg. returns for their allocations, because of investors' ill-timed activity (net inflows after high returns and before low returns)
  - Dichev (2007) shows that the dollar-weighted returns U.S. stock investors earned between 1926 and 2002 were 1.3% lower than the time-weighted (buy and hold) market returns of the NYSE/AMEX indices
- Well known academic studies indicate that large inflows predict low future returns for retail and institutional investors
  - Frazzini and Lamont (2008) show that retail investors tend to flow into mutual funds that hold stocks with low subsequent returns by 10%
  - The main underperformance occurs 6 to 30 months after inflow

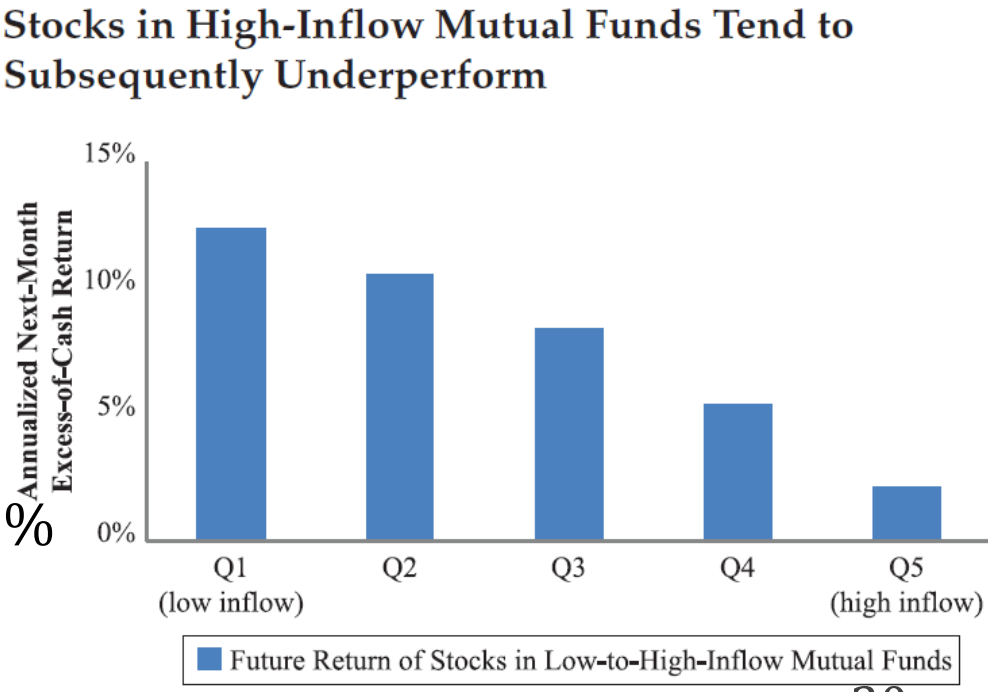

Bad Habits and Good Practices

- A rarely measured variant of this bad habit involves changes in acceptable investment universes or investable assets, as well as in benchmark or policy portfolios
- Even an investor who claims to follow a passive approach must decide which assets she deems investable: almost always, newly qualifying for investability follows strong multi-year performance
  - E.g., when did most investors extend their equity portfolios to include emerging markets or frontier markets?
  - The case of alternative asset classes is clearer: real estate, infrastructure, timber, farmland, commodities, private equity, private credit, and hedge funds all became increasingly widely held after extended benign periods, and these decisions were often reconsidered if persistent losses followed
- Why do we observe this behavior? Extrapolation implies the tendency to learn from patterns and expect their continuation
  - Humans tend to apply our instinctive desire to extrapolate even in instances when no pattern exists to be successfully extrapolated
- Procyclic actions are reinforced by social effects—herding, conventionality—and even by certain risk management rules

Bad Habits and Good Practices

② Under-diversification: although many investors value it, they may still have less diversity in their portfolios than they think

- Individual investors often hold just a few stocks; worse, their main holding may be the company they work for (stock options)
  - For proprietors of individual firms, risk concentration can remain extreme for decades
- Institutional investors rarely concentrate their risks among a few single stocks, but they may fall prey to Home Bias, when the weights of own-country assets exceed global market-cap
- Home bias has been declining over time but remains significant in every country
- A 2013 Towers Watson Study finds that the share of domestic equities has fallen from 65% in 1998 to 47%

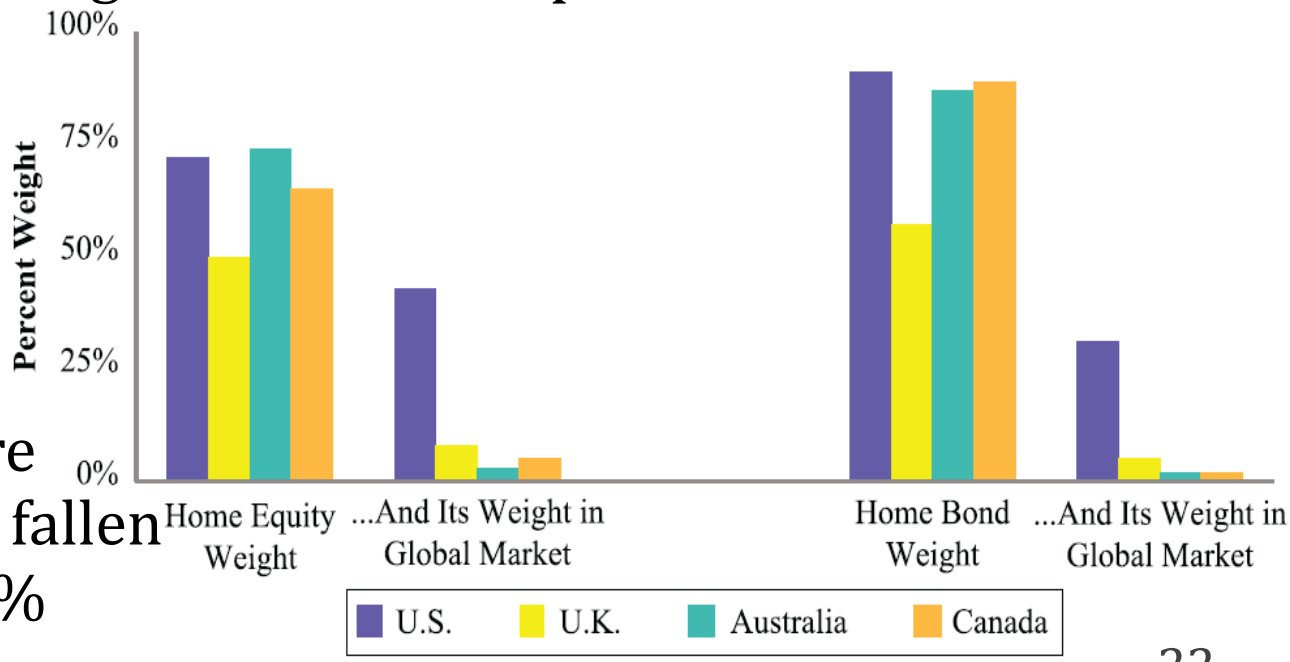

Bad Habits and Good Practices

- Most institutions allow their portfolios to be dominated by one source of risk: equity market direction
- This may reflect the illusion of diversification-- hold a large number of diverse assets, but the portfolio still has a single driving risk source
  - E.g., 60/40 stock/bond ptf generally have at least 90% risk concentration in equities, because equities are more volatile than other investments
  - Even alternative assets (private equity, hedge funds, etc.) or smart beta, do not materially help, as they often are highly correlated with equities
- One behavioral explanation for under-diversification is "narrow framing," a focus on single line-items or parts of the ptf
- The need for familiarity and aversion to ambiguity may be the primary drivers of home bias and the preference for own-company stock
- Another bias, overconfidence, is a key explanation for concentrated risks in single stocks and tactical timing bets

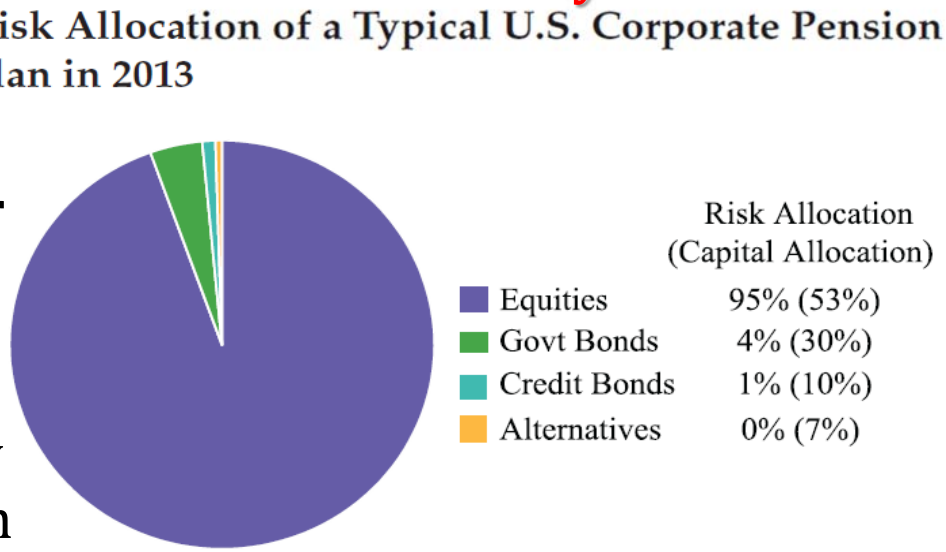

Bad Habits and Good Practices

- ③ Seeking comfort: some investors seek comfort when selecting investments instead of judging them purely on risk/reward merits
  - Such familiar investments can be structurally overpriced and thus deliver lower long-term returns
  - Investors underutilize comfort-challenging tools (leverage, shorting, and derivatives) that could be used to improve risk diversification
- Most ptfs are dominated by equity risk and, to many investors, the fact that their peers share this problem makes equity risk more bearable
- Many investors seem to overpay for embedded leverage to avoid direct leverage: leverage aversion can explain the higher risk-adjusted returns of low beta or low-volatility stocks
- Investors buy investments with a promise of large upside—assets like lottery tickets—but at the same time they buy protection against investments where the downside risk appears to loom large
- Not surprisingly, both strategies have delivered low long-run returns
- Many investors overpay for smooth returns, which explain why historical illiquidity premia on private assets are slimmer than warranted

Bad Habits and Good Practices

- It's the appearance of smoothness that comes from an inability to get timely marking to market, not actual smoothness

- Fortunately, each bad habit has a flip side: a good investment practice

Bad Habits	Good Practices
Chasing Multiyear Returns	Investing Strategically
Under-Diversification	Diversifying Risks Aggressively
Seeking Comfort	Accepting Discomfort When Paid To Do So

- All three bad habit/good practice pairs raise the question of macro-consistency: can everyone do it at the same time?
- All deviations from global market-cap ptf or buy and hold investing require that some investors take the other side of the trade
- In some cases, the proposed good practices correct bad habits in a macro-consistent way, that is, to equilibrium investment practices that every investor could enjoy
- Yet in suggesting practices such as pro-active diversification and leaving the comfort zone, one often implicitly assume that some non-equilibrium returns are available to a subset of investors only

The Limits to Alpha as Portfolio Construction Tool

④ Aggressive seeking alpha for the sake of it: not as an ex-post way to measure ptf. performance but as an absolute target to be pursued ex-ante

- There are 2 ways to view alpha: (i) as we have seen, as a measure of a fund manager's stock selection and market timing abilities, (ii) as a guideline for investors wishing to optimize their portfolios
- The alpha of an asset—calculated with respect to a given benchmark portfolio—measures the change in the portfolio's Sharpe ratio driven by a marginal increase in the asset's weight

$$\frac{\partial\left(\frac{\bar{R}_P - r_f}{\sigma_P}\right)}{\partial x_i^B} = \frac{\left(\bar{r}_i - r_f\right)\sigma_P - \left(\bar{R}_P - r_f\right)\frac{1}{\sigma_P}\sum_{j=1}^{N} x_j \sigma_{i,j}}{\sigma_P^2}$$

$$\boxed{\alpha_i = \bar{r}_i - \left[r_f + \beta_i\left(\bar{R}_P - r_f\right)\right]} \qquad = \frac{1}{\sigma_P}\left(\bar{r}_i - \left[r_f + \beta_i\left(\bar{R}_P - r_f\right)\right]\right) = \frac{1}{\sigma_P}\alpha_i$$

- The vector of alphas is thus the direction of marginal adjustment in portfolio-weight space that yields the maximal increase in the portfolio's Sharpe ratio

The Limits to Alpha as Portfolio Construction Tool

- Alphas tell investors how best to marginally adjust their portfolios relative to the benchmark: increase the weight of assets with positive alphas and decrease the weight of assets with negative alphas…
  - … and do so proportionately to the absolute size of the alpha
- Therefore, strictly speaking, alpha is a guideline only for marginal adjustments
- Levy and Roll (2016) examine the increase in Sharpe ratio obtained by shifting the ptf weights "in the direction" of the alpha vector
- They find that although alpha indeed indicates the best way to make an infinitesimal adjustment to ptf weights, it is not useful as a practical guideline in which small but finite adjustments are considered
- Rather than adjusting the weights according to alphas, the investor can do much better by directly optimizing the portfolio
- When finite adjustments are considered, adjustments according to alpha are not only suboptimal, but may be in the "wrong direction"

The Limits to Alpha as Portfolio Construction Tool

- Perhaps surprisingly, these effects take place even when the adjustments relative to the benchmark are small
  - E.g., if one allows ptf weights to deviate from the benchmark weights by only 2%, the Sharpe ratio obtained with the alpha-adjusted portfolio is about 30% lower than the Sharpe ratio obtained with the optimized portfolio with the same constraint on the distance from the benchmark
  - Alphas point in the "wrong direction" for 15% of the assets in this case
- It is hard to see alphas being useful in any practical context
- The reason is that alpha is an indication only about the best infinitesimal shift in ptf weights: once the ptf is shifted, alphas may change considerably
- This is closely related to the well-known fact that a small change in the benchmark ptf can produce a large change in the beta–expected return relationship—and thus to a large change in the assets' alphas
  - Although making a small shift in the portfolio weights in the direction of the alpha vector increases the portfolio's Sharpe ratio, those changes in the portfolio may also cause the alphas to change quite dramatically

The Limits to Alpha as Portfolio Construction Tool

- Continuing to move "in the direction" of the original alpha vector beyond the initial infinitesimal shift may be far from optimal—as this is no longer "the right direction"

- Levy and Roll use as a benchmark portfolio the value-weighted portfolio of the 100 largest U.S. stocks as of December 31, 2014
- The expected returns and covariances for these stocks are taken as their monthly sample values over January 1, 2005–December 31, 2014

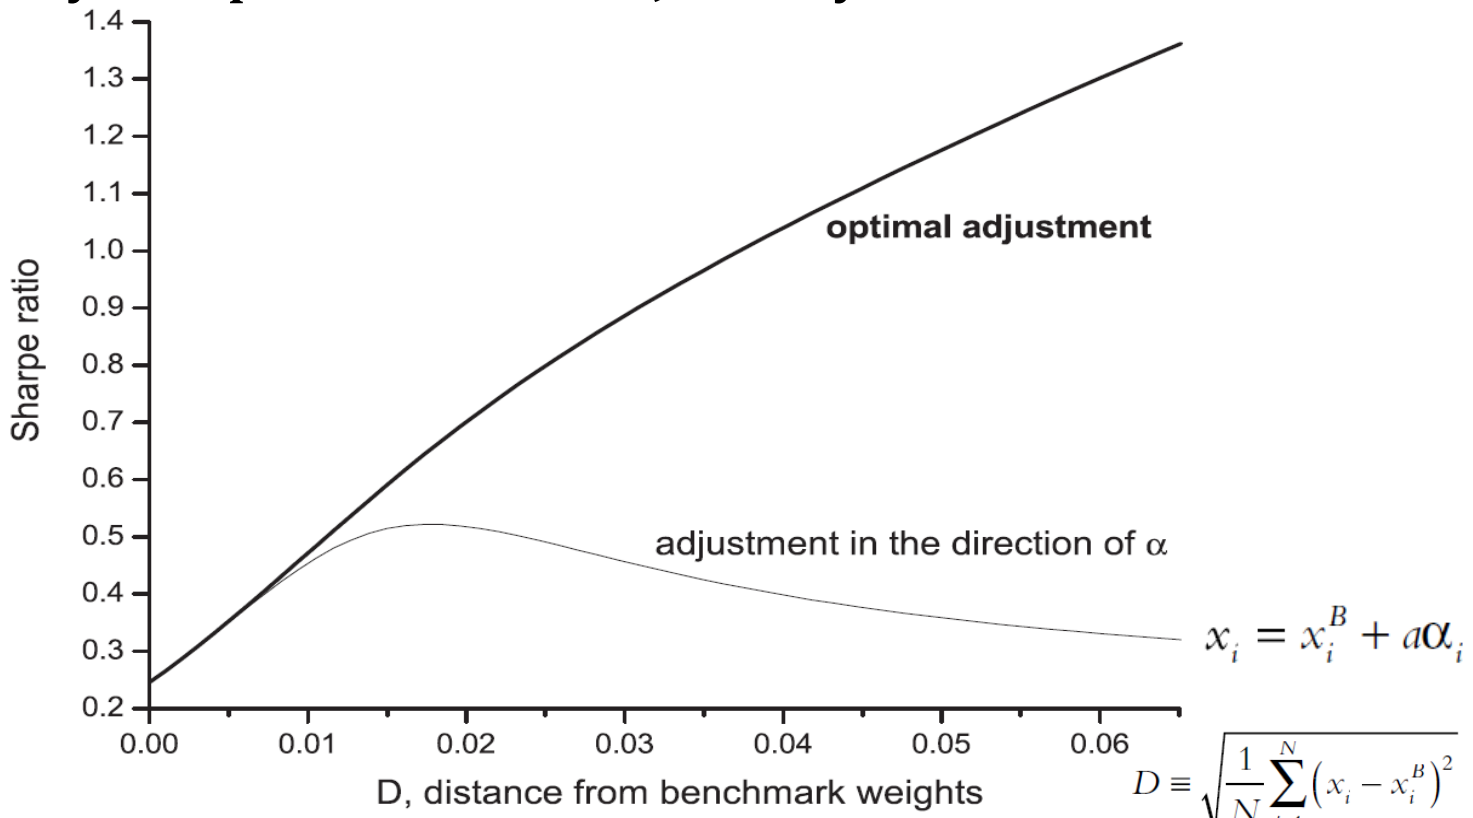

The Limits to Alpha as Portfolio Construction Tool

- When non-infinitesimal adjustments are considered, there is almost no relationship between an asset's alpha and its optimal adjustment
  - Betas can be quite sensitive to portfolio composition, and this implies that alphas are sensitive
  - When they examine the change in alphas as the portfolio is shifted, this sensitivity is exactly what they find
  - For a shift of D = 0.02, there is almost no relationship between the original alphas and the new alphas
- It is widely believed that over-weighting assets with positive alphas and underweighting assets with nega-tive alphas increase the Sharpe ratio
- This is only true for infinitesimal shifts in ptf weights!