Rear-View Mirror or Windshield?

Generating Sound Return Estimates while Controlling Estimation Error

White Paper · Portfolio Selection GmbH · 17 October 2026

Introduction

Every asset allocation rests on return forecasts – explicitly, when a model reports expected returns; implicitly, when a portfolio is assembled from fixed weights, rules of thumb or past performance. Whoever holds 60 % in equities is assuming a risk premium, even if they have never put a number on it. The question is therefore not whether returns are forecast, but how.

A widespread practice is the naive application of estimation statistics to historical returns: one simply takes the average of past annual returns and treats it as the expected value of a distribution from which future returns are drawn. This presupposes that returns are the outcomes of a repeated experiment under identical conditions – like throws of the same die.

Realised returns consist of income and price changes. Price gains may reflect higher earnings, but they may also reflect investors paying more for the same income. This second effect shifts the starting point for the future. An example: a perpetual bond with no maturity costs 100 euros and pays a coupon of 10 euros a year. Its running yield is 10 %. If its price rises to 200 euros, the existing holder earns a one-period return of 110 % including the coupon. For a new buyer, however, the running yield is now only 5 %. No subsequent price decline is required for this. The same valuation logic applies to equities, other things being equal; total return additionally depends on earnings growth and on further changes in valuation. The level of interest rates also affects valuations. The steep decline in rates since the early 1980s was a major tailwind behind past price gains. An extrapolation of historical returns must not treat such a tailwind as repeatable without examination.

A historical average can serve as a reference. As the sole forecast, however, it ignores today's starting valuation and the current level of interest rates. It describes the outcome of varying historical conditions; the investor needs an expectation under today's conditions. This white paper therefore examines structured, fundamentally grounded return estimates: their economic logic, the empirical evidence, their uncertainty and their use in a decision process that explicitly accounts for possible estimation errors.

Summary

Historical average returns are not a sufficient basis for today's return expectations. Structured forecasts decompose expected return into income, growth and valuation change. Their principal value lies in making assumptions and uncertainties visible and verifiable. The literature reviewed in the list of sources provides theoretical arguments and empirical indications in support; it does not establish a general or always realisable excess return.

The benefit depends on the horizon. In the studies discussed, valuation measures offer little guidance over one year; over longer horizons, stronger relationships emerge. Their statistical significance, however, depends on the sample and on overlapping observation windows. For 2015–2024, Böni et al. report model-internal expected risk-adjusted advantages of 51–88 basis points p.a.; realised results over the same period provide no statistically robust evidence of an excess return.

Structured forecasts are thus a foundation for long-term planning, not a reliable timing tool. They require transparent assumptions and an examination of the consequences of possible errors. Conviction contributes reasoned return expectations; Concern limits the allocation's dependence on their being correct.

How Historical Averages Mislead

A 60/40 portfolio of US equities and bonds returned 7.6 % nominal and around 5 % real p.a. from 1871 to 2010 (Brightman 2012). These figures underlie the 7.5–8 % that many pension plans and endowments assume. Yet they are not a measure of the future but the outcome of a starting point with a 4.6 % dividend yield and a 5 % bond yield.

The mechanism is simple: past returns – especially after good decades – consist to a large extent of valuation increases. The higher the valuation today, the less of it remains for the future. Markowitz puts it succinctly in his interview with Haas: the greater the historical returns, the higher current valuations and the lower the probability that the same assets will deliver the same returns in future.

History bears this out. A portfolio of 60 % equities, 10 % commodities and 30 % bonds returned 11.6 % p.a. over 1970–1999 and only 2.9 % over 2000–2009 (J.P. Morgan). Conversely, at the end of 1980 the historical average forecast around 7–8 % for 60/40 while bonds yielded 12.8 %; the realised return was 14 % p.a. Naive extrapolation is wrong in both directions: too optimistic after good decades, too pessimistic after bad ones.

Böni et al. add the behavioural side: surveys show that retail investors and equity analysts extrapolate the recent past. Their expectations are procyclical – highest when markets are expensive. Institutional capital market assumptions behave the opposite way.

Ilmanen (AQR, 2025) compares both approaches over roughly 150 years. In this analysis, the cyclically adjusted earnings yield correlates with the return of the following decade at +0.52, whereas the return of the past decade correlates at –0.37. One possible explanation is that good market phases leave behind high valuations that may subsequently normalise. The historical pattern, however, is no certain forecast for any individual decade.

Brooks (AQR, 2023) makes this concrete for the last decade. The 11.9 % excess return of the S&P 500 over cash from mid-2013 to mid-2023 consisted of 2.1 % dividends, 4.5 % real earnings growth, 3.6 % valuation expansion (CAPE from 24 to 30) and 1.7 % from negative real interest rates. A repeat would require a CAPE of 51 even with 6 % real earnings growth – above the tech-bubble peak of 44. Whoever naively extrapolates the past is assuming an unprecedented valuation.

The Building-Block Methodology

The models examined decompose expected return into economic building blocks. For bonds, the starting yield is a central point of departure; for equities, it is payouts, earnings growth and valuation changes. Only some of these quantities are observable today. Future growth and future valuations have to be estimated. The models differ in their design and in the judgement calls they require.

Building block	Brightman (2012)	Invesco CMA (2024)	J.P. Morgan LTCMA (2025)
Equity income return	Dividend yield	Dividends + buybacks (US, 10-yr smoothed)	Dividends + buybacks – gross dilution
Earnings growth	1.7 % real EPS growth (140 yrs of history)	real GDP per capita + expected inflation	Revenue from trend GDP + margin normalisation
Valuation change	deliberately ignored	Reversion to regression P/E, scaled by distance	Normalisation to target P/E over 10–15 years
Bonds	Starting yield	Average of current and expected yield + roll-down – credit losses	cycle-neutral yield after normalisation + curve + roll-down – defaults
Risk	–	historical quarterly returns, exponentially weighted correlation	History from 2006, regime-weighted (15 % stress, 10 % high inflation)

Brightman is deliberately minimalist: three verifiable numbers, no assumption about valuations. Invesco formalises the normalisation of the P/E ratio and dampens it when the deviation is small. J.P. Morgan relies on analyst judgement about margins, dilution and target valuation. What all models share is transparency: every building block can be questioned and replaced individually – an advantage that J.P. Morgan explicitly cites as the most important benefit in client dialogue.

What matters is what these models are not: error-free images of the future. The capital market assumptions of J.P. Morgan and Invesco are fundamentally grounded, component-wise estimates whose building blocks differ in their estimability by orders of magnitude. The payout yield is observable today. Trend growth can be roughly estimated from long series. The valuation change is largely an assumption – about the target level and about the period of normalisation. The value of the decomposition therefore lies less in the forecasting accuracy of the sum than in its attributability: it makes visible where the uncertainty sits, and it forces the sum into an identity that can be tested. Brooks' calculation that a repeat of the last decade would require a CAPE of 51 is precisely this test – it would be impossible with a simple average of realised returns.

The decomposition does not automatically diversify estimation errors, because the components are dependent. A high payout yield may reflect a low valuation, but also a high payout ratio or particular risks. Whether a later normalisation raises returns depends on further assumptions. Earnings growth and target P/E may likewise rest on shared assumptions, for instance about margins and their persistence. The same error can then hit several building blocks at once. Such dependencies belong in the sensitivity analysis.

Empirical Evidence

Over a decade horizon the models have a track record that – consistent with the component-wise reading – is measured less by the hit rate of the sum than by the origin of the errors. Invesco's back-calculation for the S&P 500 since 1973 shows a correlation of 0.84 between forecast and realised ten-year return with a mean error of 3.1 percentage points; for US investment-grade bonds since 2001, a correlation of 0.89 with 0.6 percentage points. The difference is instructive: for bonds the starting yield is observable and dominates, for equities the error stems almost entirely from the valuation building block. Brightman's minimalist model – dividend yield, 1.7 % real growth, inflation, no valuation assumption – matched the realised 60/40 return over 14 decades to within 0.5 percentage points on average. That is remarkable, but it is an average across decades that offsets individual decade errors (the World War decade of the 1910s, the tech bubble of the 1990s), and no reliable indication of forecasting accuracy in any single decade.

Model	Forecast object	Period	Result
Brightman	60/40, 10-year return	1871–2010	expected avg. 7.2 %, realised avg. 7.6 %
Ilmanen (AQR)	US equities, 10-year return: earnings yield vs. rear-view mirror	1871–2024	Correlation +0.52 (earnings yield) vs. –0.37 (past return)
Invesco CMA	S&P 500, 10 years	1973–2023	Correlation 0.84, mean error 3.1 percentage points
Invesco CMA	US IG bonds, 10 years	2001–2023	Correlation 0.89, mean error 0.6 percentage points
J.P. Morgan	60/40 ACWI/US Agg, 2010 forecast	2010–2024	realised always within the 5–95 % band, near the median until 2022
Invesco (Vig)	S&P 500 price return from Shiller P/E	1983–2015	R^2 0.74; since 1881 only 0.10

For high-quality bonds, the starting yield often provides a more reliable point of departure than for equities. How well it describes the subsequent return depends, among other things, on the investment horizon, duration, reinvestment and possible defaults. Price losses from rising rates can be offset partially, or over sufficiently long horizons largely, by higher reinvestment income. For equities, the future valuation in particular remains a substantial source of uncertainty.

The reported long-term statistics should be read with caution. Many rest on overlapping ten-year windows and thus on strongly dependent observations. A period from 1983 to 2015 contains only about three non-overlapping ten-year periods. This does not yield a blanket number of independent statistical observations; it does, however, show how limited the data base is. High correlations and R^2 values alone therefore establish neither statistical certainty nor reliable forecasting performance outside the sample examined (Boudoukh, Richardson and Whitelaw 2008).

The forecast vintages reviewed illustrate the differences: at the end of 2023 Invesco expected 6.4 % p.a. for the S&P 500; in September 2025 J.P. Morgan cited 6.7 % for US large caps, 7.0 % for global equities, 4.6 % for ten-year US Treasuries and 5.2 % for investment-grade credit. A CAPE-based analysis in February 2025, by contrast, implied around 2.3 % annual price return. Because of differing data dates and return definitions, these figures are not directly comparable. They do, however, illustrate the influence of growth, interest-rate and valuation assumptions. The choice of a historical target P/E is not a neutral input; its transferability to today's conditions must be justified.

The Role of the Forecast Horizon

Long-term valuations can be more informative for strategic planning than for short-term price forecasts. In the Invesco study reviewed, the relationship between CAPE and subsequent return is practically zero over one year; in selected ten-year analyses the R^2 reaches 0.74–0.78. These values are subject to the sampling and overlap problems described above. Earnings growth and valuations fluctuate strongly in the short term; even in the long term there is no guarantee of a return to any particular normal level.

This is why J.P. Morgan and Invesco work with 10- to 15-year horizons and spread the valuation normalisation over that period. Invesco's 5-year variant doubles the valuation effect and is explicitly offered only as a half-cycle reading for tactical purposes.

Ilmanen differentiates the horizons more finely still. Over months, extrapolation is empirically even defensible, because returns exhibit short-term trends. Over three to ten years it is most dangerous: these are the reversal horizons at which valuations normalise – whoever uses the rear-view mirror here behaves like a momentum investor precisely where momentum does not work. For trend growth and risk premia, by contrast, the longest possible histories are needed; a single decade yields "anecdotes, not statistics". In 2000 and in 2010 the rear-view mirror delivered catastrophic forecasts for the following decade – far too high in one case, far too low in the other.

Findings of Böni, Brüggemann and Kröncke (2025)

The study combines the institutional capital market assumptions of the Horizon Actuarial Survey (41 institutions, 10-year horizon) with a rule-based Black-Litterman model and tests whether static portfolios such as 60/40 or 1/N are efficient ex ante. Three results are central to the horizon question:

1. **Ex ante**, the forecast-based allocation achieves 51 to 88 basis points of risk-adjusted excess return p.a. over the static benchmarks – statistically significant and stable across the years 2015–2024. The advantage is largest against 60/40 (88 bp) and smallest against 1/N (51 bp).
2. **Ex post**, realised alphas over the same period lie between −30 and +62 basis points and are not distinguishable from zero. The main cause is the consistent underweight of US large caps, which the assumptions classified as expensive throughout – and which nevertheless remained the best asset class for ten years.
3. The dispersion of expected returns across asset classes is markedly smaller than that of realised returns. A substantial part of the subsequent differences thus arises from unexpected developments. The Brinson studies examine the importance of investment policy for the variation of portfolio returns; taken on their own, they do not demonstrate that return differences between asset classes can be reliably steered in advance.

These findings do not render the forecasts automatically worthless, but they limit the conclusions that can be drawn. Dahlquist and Ibert (2024) show that institutional assumptions are consistent with well-founded, countercyclical long-term expectations. That is a plausibility check, not proof of a subsequent excess return. With small expected advantages, even a decade can be too short for statistically robust evidence. At the same time, a long phase of weak results can be grounds for re-examining assumptions and implementation.

For practice this means: the horizon of the forecast must match the horizon of the mandate. Whoever allocates over ten years can use structured forecasts; whoever is judged over one to three years can look wrong with them for long stretches.

Markowitz' Perspective

Harry Markowitz, whose portfolio theory is the foundation of every allocation optimisation, did not answer the question of historical versus forward-looking inputs late in life, but already in his principal work Portfolio Selection (1959) – on its first pages. There he writes that portfolio selection should be based "on reasonable beliefs about future rather than past performances per se"; whoever selects on past data alone is in effect assuming that past average returns are good estimates of future returns (p. 14). For him, the analysis of the past is only the first and most naive of three kinds of input; the second are the "probability beliefs" of experts, the third the predictions of a method together with their uncertainty (pp. 26–33). Forty-seven years later, in his interview with Ferdinand Haas (Institutional Money, 2006), asked whether it dismayed him that investors feed nothing but historical data into a portfolio optimiser, he replied tersely: "Yes, that is bad."

Four statements are decisive for this white paper; each appears in the 1959 book and is confirmed in the 2006 interview:

- **Return from forward-looking information.** For Markowitz the security analyst is "the meteorologist of stocks and bonds": he takes the past into account but does not extrapolate it naively. Past data are presented to him "for convenience", "to be considered or ignored as they saw fit" (1959, pp. 28, 31). In the interview: investment decisions should be made on forward-looking information rather than on historical data. There he explicitly confirms the relationship that high past returns lead to high valuations and thus to lower future risk premia.
- **Risk from history.** For volatility and correlation he recommends in the interview the longer data series, because it contains events like 1929 and gives the user an understanding of what uncertainty means. In the book he is stricter still: variances and covariances, too, are "of course" beliefs and not simply past values (1959, p. 297, fn. 1) – the long history is material for judgement, not a substitute for it. The separation – returns forward-looking, risk from long history – is practised by J.P. Morgan and Invesco in exactly this form.
- **Estimates, not truths.** "No method of prediction is infallible"; both predictions and the uncertainty of predictions are inputs to the analysis (1959, p. 33). Whoever holds beliefs need not believe they are "good beliefs" – he may concede that he is systematically too optimistic or too pessimistic; all that is presupposed is the ability to learn over time (pp. 272–273). In the interview: one should confine oneself to conservative estimates. As a Bayesian, he knows nothing with final certainty.
- **Simulation instead of point forecast.** As early as 1959 Markowitz explicitly checks the result of his optimisation against what "unaided judgment" would have chosen, calling this "some check on the suitability of the results" (pp. 296–297). In the interview he goes further: mean-variance is only one part of the analysis; in practice it must be combined with simulation analysis, because there is far more uncertainty about what may happen in ten or twenty years than anyone is thinking of today.

For long-term outcomes, the distinction between the arithmetic average return and the geometric growth rate also matters. For the same arithmetic mean, larger fluctuations typically reduce long-term wealth growth. Leverage can raise expected return, but it also magnifies loss risks; excessive leverage can reduce the growth rate. Added to this are extreme events and illiquidity, which simple distributional assumptions may capture inadequately. Forecasts should therefore be viewed as a band of possible outcomes and together with stress scenarios.

Limits and Judgement Calls

Structured forecasts offer a comprehensible economic foundation. They are, however, neither objective nor proven to be generally superior. The literature reviewed and the opposing positions identify in particular the following limits.

Not a timing instrument. Invesco's CAPE model said the same thing in February 2024 as in February 2025 – US equities expensive, low ten-year return. In between, the S&P 500 rose by 19 %. The author concedes that the market has "defied gravity for a long time" and may continue to do so. Böni et al. show the same over an entire decade: the underweight of expensive US large caps was justified under the model assumptions but proved a major drag on performance in the period examined.

Sample choice is decisive. The CAPE regression explains three quarters of return dispersion since 1983, but only a tenth over the full history since 1881, because the normal valuation level rose in the 1980s from an average of 15.6 to 24.5. Which period serves as the benchmark is a judgement call with great leverage. Markowitz' anecdote of the broker who wavered between data from 1927 and data from 1950 describes exactly this problem.

The building blocks sometimes contradict each other in detail. Brightman considers buybacks in aggregate a myth, because new issuance has exceeded them by a net 1.25 % p.a. since 1960; Invesco, by contrast, adds a buyback yield of 1.2 %. Brightman ignores valuation changes, the CAPE regression counts them in full, Invesco and J.P. Morgan sit in between. The results cited for US equities differ markedly. Data dates, methodology and return definition must be borne in mind; in particular, price return and total return are not directly comparable.

Ex-ante advantages are model-dependent. The 51–88 basis points reported by Böni et al. measure an expected advantage on the basis of the capital market assumptions used. Studies by Dahlquist and Ibert (2024) and by Couts, Gonçalves and Loudis (2025) help to place such expectations in an external context. They do not, however, replace evidence that the specific allocation actually performs better. The short backtest and the coarse breakdown of asset classes further limit what can be inferred.

The opposing position. Goyal and Welch (2008) show that the well-known predictors – including dividend and earnings yield – forecast the equity premium out of sample no better than the historical average; anyone who had used the regressions in real time would in fact have done worse for long stretches. The finding applies to short horizons and to individual predictors in regression form. The models described here respond in two ways: they work on ten-year horizons, and they use valuation not as a regressor but as the starting point of an identity of payout, growth and valuation change. That this yields better results is not thereby proven – Böni et al. show how long the evidence can take – but the critique hits the regression logic harder than the component logic.

A historical average has one practical merit: given a fixed data base, it is simple and reproducible to compute. It is nevertheless not free of judgement, since period, market, currency and return definition must be chosen. Structured forecasts add further decisions. Whoever uses them should document these assumptions and examine their effects.

Conclusion and Consequences for the Investment Process of Portfolio Selection GmbH

For long-term mandates we prefer structured, fundamentally grounded return estimates to an unexamined extrapolation of historical averages. The reason is their economic traceability: they take today's starting point into account and make assumptions and sources of error visible. This does not amount to a general proof of superior investment performance. The decisive next step is to control the effects of possible estimation errors on the specific allocation.

From this follow our five principles for capital market assumptions:

1. Assess returns and risks forward-looking. Derive expected returns from starting income, growth and valuation. Use long risk histories including crises as the basis and examine their transferability to the future.
2. Adopt a long-term horizon. Align capital market assumptions to ten years and longer. Do not confuse this period with a guarantee of valuation normalisation or with a short-term trading signal.
3. **Conservative and transparent.** Every building block is documented and justified; in case of doubt the more cautious assumption applies (Markowitz). No sample is chosen because it delivers the desired result.
4. **Expected value with a band, not a point forecast.** Forecasts enter simulations as distributions; the central value is an orientation, not a promise.
5. Implement with discipline and remain verifiable. An underweight in expensive asset classes can lag for years. That justifies neither a reflexive change of method nor unlimited persistence. Assumptions, confidence and implementation are to be reviewed regularly against criteria defined in advance.

Naive extrapolation remains useful as a reference: it shows what an investor expects who looks only in the rear-view mirror, and it makes the difference to the structured forecast visible.

Implementation in the Investment Process of Portfolio Selection GmbH

Portfolio Selection GmbH implements the five principles with the software of Portfolio Generator GmbH. The starting point is a principle that follows from everything above: capital market assumptions are intelligent estimates, not truths. This is why the entire logic places not the model but the decision-maker at the centre. The relevant question is not which return vector is correct – nobody knows that – but which assumption would cost the decision-maker dearly if it were wrong. The process has three stages that make this question answerable; each addresses one of the sources of error described above.

Placing the decision-maker at the centre is not a house philosophy but the original reading of the theory by its author. Markowitz wrote in 1959, when the computer was still a rare exception: “The computer supplements rather than supplants the human analyst. The analyst, or analyst team, must decide on the objectives of the analysis, select securities to be analyzed, decide on a model for derived covariances, and make judgments pertaining to the expected returns, variance of returns, and covariances of returns of securities” (Portfolio Selection, p. 114). The machine takes over the computation; judgement remains with the human. The process of the Return Estimator module of Portfolio Generator GmbH described below implements this division of labour.

1. **Everything in the space of return vectors.** Every perspective is expressed as a vector of expected returns. For forecasts this is immediately the case; for portfolios – market capitalisation, equal weighting, minimum variance – so-called reverse optimisation determines the return vector that would make precisely that portfolio appear efficient.
2. Conviction and Concern. The perspectives are divided into two groups. Conviction comprises return expectations to which the decision-maker attributes forecasting power: for instance the capital market assumptions of J.P. Morgan and Invesco, or the decision-maker’s own views. Concern has a different task: it is meant to control the consequences of possible estimation errors. This is served by return vectors derived from reference portfolios, such as a market-implied vector, a vector consistent with equal weighting and a minimum-variance reference. These quantities are not additional independent forecasts. They define deliberately chosen shrinkage targets towards which the allocation can tend when confidence in the Conviction assumptions is lower. The kinship with multi-fund approaches such as Kan, Wang and Zhou (2016) serves methodological classification; it is no proof of equivalence or effectiveness. Conviction determines which expectations are followed; Concern limits dependence on their being correct.
3. Weighting in two stages and comparison of alternative perspectives. First the perspectives are weighted within the groups, then the groups against each other. In Rashomon mode, the resulting allocation is re-assessed under each individual assumption. This makes visible which expectations support it and under which alternatives it appears unfavourable. Broadie (1993) shows in simulations how estimation errors can make the estimated efficient frontier appear too optimistic. The comparison of perspectives takes up this problem, but knows neither the true return distribution nor the true efficient frontier. It shows sensitivities relative to the chosen assumptions and does not test an unknown truth.

Methodological guiding idea. Classical optimisation asks for an efficient portfolio given a set of expectations. The Return Estimator adds the question of how sensitive this decision is to other expectations. Conviction contributes reasoned return expectations; Concern acts as a deliberately chosen corrective against estimation errors. The sensitivity analysis makes alternative perspectives comparable and can reduce dependence on individual assumptions. It guarantees neither a general limitation of losses nor an excess return. Judgement about the weighting remains with the decision-maker and is to be documented. Even such a process can lag a simple benchmark for years; this possibility must be explicitly taken into account in its design and ongoing review.

Sources

Source	Type	Key contribution
Brightman, C. J. (2012): Expected Return. Investments & Wealth Monitor, Jan./Feb. 2012, pp. 24–29, 34	Journal article, Research Affiliates	Simple building-block model; 140-year comparison of expected vs. realised 60/40 return
Broadie, M. (1993): Computing Efficient Frontiers using Estimated Parameters. Annals of Operations Research 45, pp. 21–58	Journal article	True, estimated and actual efficient frontier; optimism bias through error maximisation; means from models, risks from history
Brooks, J. (2023): Driving with the Rear-View Mirror. AQR Alternative Thinking, 4Q 2023	Research note	Decomposition of the 11.9 % excess return 2013–2023; conditions for a repeat
Ilmanen, A. (2025): How Do Investors Form Long-Run Return Expectations? Part 1. AQR, April 2025	Research note	Earnings yield vs. rear-view mirror 1871–2024 (+0.52 vs. −0.37); subjective vs. objective expectations; reversal horizons
J.P. Morgan Asset Management (2024): Long-Term Capital Market Assumptions – Methodology Handbook	Methodology handbook	Complete building-block and normalisation model; track record of the 2010 forecast
J.P. Morgan Asset Management (2025): 2026 Long-Term Capital Market Assumptions, as of 30 September 2025	Annual publication	Current forecasts: US large caps 6.7 %, global equities 7.0 %, 10-yr US Treasuries 4.6 %
Invesco Investment Solutions, M. Shmerling (2024): Capital Market Assumptions, presentation for DVAG	Methodology presentation	Formula-based building blocks; correlation CMA vs. realised since 1973
Invesco Global Market Strategy Office, A. Vig (2025): The Shiller P/E and S&P 500 returns revisited, 2 March 2025	Market commentary	CAPE regression; horizon dependence of forecasting power; sampling problem
Haas, F. (2006): “Ich wollte, dass die Leute die Theorie verstehen” [“I wanted people to understand the theory”] – Interview with Harry M. Markowitz. Institutional Money 1/2006, pp. 54–60	Interview	Forward-looking return estimation, long risk history, conservative estimates, simulation
Kan, R., Wang, X., Zhou, G. (2016): On the Value of Portfolio Optimization in the Presence of Estimation Risk: the Case with and without Risk-free Asset. Working Paper, University of Toronto	Working paper	Three-fund rule (tangency, minimum-variance, 1/N portfolio); estimation error in μ vs. Σ; correction of DeMiguel et al.
Markowitz, H. M. (1959): Portfolio Selection: Efficient Diversification of Investments. Cowles Foundation Monograph 16, John Wiley & Sons, New York	Monograph	Inputs as beliefs about the future rather than past data (pp. 14, 26–33); the analyst judges, the machine computes (p. 114); conditions of applicability (p. 297)
Böni, P., Brüggemann, M., Kröncke, T. A. (2025): How Efficient are Static Multi-Asset Portfolios? Evidence from Institutional Capital Market Expectations. SSRN 5315806, 23 June 2025	Working paper	Ex-ante alpha 51–88 basis points significant, ex post not; dispersion of expected vs. realised returns
Boudoukh, J., Richardson, M., Whitelaw, R. F. (2008): The Myth of Long-Horizon Predictability. Review of Financial Studies 21(4), pp. 1577–1605	Journal article	Overlapping observation windows overstate the statistical significance of long-horizon predictability
Brinson, G. P., Hood, L. R., Beebower, G. L. (1986): Determinants of Portfolio Performance. Financial Analysts Journal 42(4), pp. 39–44; Brinson, G. P., Singer, B. D., Beebower, G. L. (1991): Determinants of Portfolio Performance II: An Update. Financial Analysts Journal 47(3), pp. 40–48	Journal article	Dominant influence of asset allocation on the dispersion of realised returns

Source	Type	Key contribution
Couts, S. J., Gonçalves, A. S., Loudis, J. (2025): The Subjective Risk and Return Expectations of Institutional Investors. Working Paper (SSRN)	Working paper	External validation of institutional capital market assumptions
Dahlquist, M., Ibert, M. (2024): Equity Return Expectations and Portfolios: Evidence from Large Asset Managers. Review of Financial Studies 37	Journal article	Institutional assumptions are countercyclical and consistent with econometric long-term estimates
Goyal, A., Welch, I. (2008): A Comprehensive Look at the Empirical Performance of Equity Premium Prediction. Review of Financial Studies 21(4), pp. 1455–1508	Journal article	Opposing position: predictors do not beat the historical average out of sample

The empirical figures and references follow the sources cited and their respective data dates. Historical forecasts are to be read as time-bound examples, not as current capital market assumptions.