From Signal to Weights: Quantitative Portfolio Construction

Dennis Hoffmann

Abstract

Factor-based investing rests on the premise that observable firm characteristics, such as valuation ratios, momentum, or profitability measures, carry information about expected returns. Yet the path from a raw characteristic to an investable portfolio is not a single step; it is a pipeline of engineering decisions: normalization, composite scoring, weighting, and constraint enforcement. Each decision materially affects realized performance. This article traces the signal-to-weights pipeline for equity portfolios, surveys the major weighting schemes (equal weight, score-proportional, minimum variance, mean-variance, and risk parity), and discusses the practical constraints that shape real portfolios. A concluding section positions machine learning as an extension of this traditional framework rather than a replacement.

1 Introduction

A substantial body of empirical research documents that certain firm characteristics predict cross-sectional stock returns. Value stocks (high book-to-market) have historically outperformed growth stocks (Fama and French, 1993), past winners have continued to outperform past losers over intermediate horizons (Jegadeesh and Titman, 1993; Carhart, 1997), and firms with higher profitability have earned higher average returns (Novy-Marx, 2013; Fama and French, 2015). These regularities, often called factor premia, have motivated a large industry of quantitative and factor-based investment strategies (see Ang, 2014, for a comprehensive treatment).

However, the existence of a predictive signal is only the starting point. Converting a signal into an investable portfolio requires a sequence of engineering decisions: how to normalize raw data across firms and time, how to combine multiple signals into a composite score, which weighting scheme to apply, and which constraints to enforce. These choices are not mere implementation details; they materially affect portfolio risk, return, capacity, and turnover. Portfolio constraints alone can substantially alter the relationship between signal strength and portfolio performance (Clarke et al., 2002), and the choice of weighting scheme can matter as much as the choice of signal (DeMiguel et al., 2009).

This article traces the full pipeline from raw firm characteristics to portfolio weights within a cross-sectional, single-period framework for equities. Section 2 briefly introduces factor signals and the distinction between tilts and pure factor portfolios. Section 3 covers signal processing: normalization, outlier treatment, and composite score construction. Section 4, the core of the article, surveys the major weighting schemes. Section 5 addresses constraints, transaction costs, and implementation. Section 6 positions machine learning as an extension of this traditional pipeline. The scope is limited to equities; fixed income, derivatives, and multi-asset portfolios raise additional considerations that fall outside the present discussion.

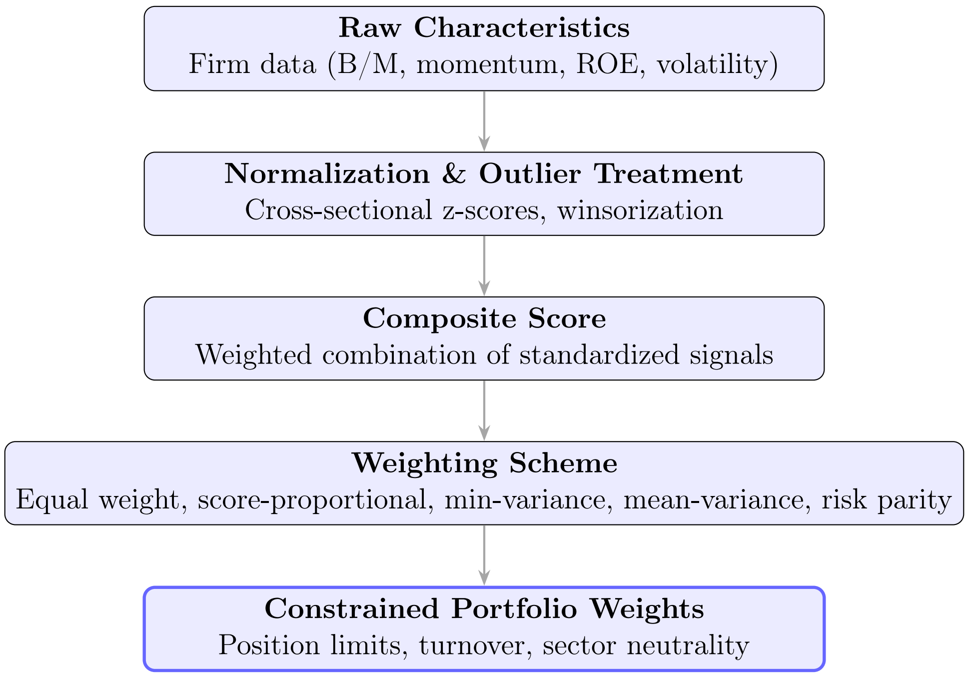

Figure 1: The signal-to-weights pipeline for factor-based equity portfolios.

2 Factor Signals in Brief

2.1 What Is a Factor Signal?

A factor signal is an observable firm characteristic hypothesized to predict cross-sectional differences in expected stock returns. The term "factor" originates from asset pricing theory, where it denotes a systematic source of risk or return variation, but in practice it has broadened to include any characteristic used to sort stocks into portfolios (Cochrane, 2011).

Canonical examples include value, typically measured as book-to-market equity (Fama and French, 1993); momentum, measured as the cumulative return over months $t-12$ through $t-2$ (Jegadeesh and Titman, 1993; Carhart, 1997); quality, where Novy-Marx (2013) uses gross profitability and Fama and French (2015) use operating profitability as distinct but related measures; low volatility, measured as realized return volatility or beta (Baker et al., 2011; Blitz and van Vliet, 2007); and size, measured as market capitalization (Banz, 1981; Fama and French, 1993). Each of these characteristics has been documented to carry a positive risk premium in long-short portfolio sorts, and several have been shown to persist across geographies and asset classes (Asness et al., 2013), though the magnitude and persistence of some premia remain debated (Harvey et al., 2016; Jensen et al., 2023).

2.2 Factor Tilts vs. Pure Factor Portfolios

An important distinction exists between pure factor portfolios and factor tilts. A pure factor portfolio is constructed directly from signal scores, typically by going long the top-ranked stocks and short the bottom-ranked stocks (the academic long-short portfolio). A factor tilt, by contrast, starts from a benchmark portfolio and introduces modest overweights and underweights to increase exposure to the desired characteristic while maintaining a controlled tracking error relative to the benchmark.

The choice between these approaches involves a trade-off between factor exposure purity and investability. Long-short portfolios isolate the factor premium more cleanly but face shorting constraints, borrowing costs, and limited capacity. Tilt-based portfolios are more practical for large allocations but dilute the factor exposure with benchmark-like holdings. A related development is fundamental indexation (Arnott et al., 2005), which weights stocks by economic fundamentals rather than market capitalization, producing a systematic value tilt relative to cap-weighted benchmarks. Both approaches pass through the same signal-to-weights pipeline described in the following sections; they differ primarily in how the final weights are anchored.

For a deeper treatment of which factors survive replication and the multiple-testing challenges of the "factor zoo," see Harvey et al. (2016) and Jensen et al. (2023).

3 From Raw Characteristics to Scores

3.1 Cross-Sectional Normalization

Raw firm characteristics are not directly comparable across time or across signals. Book-to-market ratios in 2005 and 2020 inhabit different ranges due to market-wide valuation shifts, and a given numerical value of momentum has a different distributional meaning than the same numerical value of return on equity. Cross-sectional standardization removes these level differences.

The standard approach is the cross-sectional z-score, computed at each point in time t across all N firms in the investment universe:

$$z_{i,t} = \frac{x_{i,t} - \bar{x}_t}{\sigma_t} \tag{1}$$

where $x_{i,t}$ is the raw characteristic for firm i at time t, and $\bar{x}_t$ and σ_t are the cross-sectional mean and standard deviation at time t. The resulting z-scores have zero mean and unit variance in every cross-section, making signals comparable across time and across different characteristics.

The cross-sectional nature of this transformation is critical. Time-series normalization (standardizing each firm's characteristic over its own history) would serve a different purpose, namely removing firm-specific level shifts. For portfolio construction based on relative rankings, the cross-sectional variant is the natural choice because portfolio weights depend on a firm's characteristic relative to its peers at a given point in time.

3.2 Outlier Treatment

Financial data is heavy-tailed. Extreme observations, whether due to genuine economic events (distressed firms, extreme earnings surprises) or data errors, can dominate z-scores and distort portfolio weights. Winsorization addresses this by clipping values beyond a specified percentile threshold before computing z-scores. A common choice is to clip at the 1st and 99th percentiles of the cross-sectional distribution.

The order of operations matters: winsorize first, then compute z-scores. If z-scores are computed on raw data, a single extreme observation can compress the z-scores of all other firms toward zero, reducing the effective signal variation. Winsorization before standardization ensures that the resulting z-scores reflect the bulk of the distribution rather than being driven by outliers.

3.3 Composite Score Construction

Most factor-based strategies combine multiple signals into a single composite score per firm. Given K standardized signals, the composite score for firm i is typically a weighted sum:

$$S_{i,t} = \sum_{k=1}^{K} w_k \cdot z_{i,k,t} \tag{2}$$

where w_k is the weight assigned to signal k. Equal weighting ($w_k = 1/K$) is the simplest and most common choice. Optimized factor weights, for instance those chosen to maximize the in-sample information ratio of the composite, introduce estimation risk and the potential for overfitting. In practice, equal weighting of signals is a common default. The broader principle that simple, estimation-free rules can outperform optimized alternatives out of sample, as DeMiguel et al. (2009) demonstrate for portfolio weights, provides an intuitive rationale, though their result applies to asset allocation rather than signal combination directly.

Signal decay is an additional consideration. Different characteristics have different half-lives: momentum signals decay within months, while value signals may persist for a year or more. When combining signals with different decay profiles, the rebalancing frequency should reflect the fastest-decaying component, or alternatively, decaying weights can be applied to older signal values.

4 Weighting Schemes

Given a vector of composite scores, the next decision is how to translate scores into portfolio weights. This section surveys five major approaches, each embodying different assumptions about what information is available and trustworthy.

Table 1: Signal processing pipeline.

Step	Input	Output	Key Decision	Common Pitfall
Raw characteristic	Firm-level data	Raw signal values	Which characteristic(s)	Data quality, covera
Winsorization	Raw signals	Clipped signals	Percentile thresholds	Over-aggressive clip
Cross-sectional z-score	Clipped signals	Standardized scores	Universe definition	Non-normal distribu
Composite construction	Multiple z-scores	Single score per firm	Signal weights	Correlated signals,

4.1 Score-Proportional Weights

The most direct approach sets portfolio weights proportional to composite scores: $w_i \propto S_i$. Firms with higher scores receive larger weights, directly expressing signal conviction in the portfolio. A rank-based variant replaces raw scores with percentile ranks, which is more robust to outliers but discards information about the magnitude of score differences.

Score-proportional weighting has the virtue of transparency: the portfolio is a direct, monotonic function of the signal. Its weakness is that it ignores the risk structure of the portfolio entirely. Two firms with identical scores but very different volatilities or high correlation will receive identical weights, potentially concentrating risk.

An additional subtlety arises when composite scores span both positive and negative values, as z-scores typically do. Setting $w_i \propto S_i$ directly would assign negative weights to below-average firms, producing a long-short portfolio. In practice, score-proportional weighting is usually applied either within a pre-selected long-only subset (e.g., the top quintile) or after shifting scores to be non-negative. The choice between these implementations materially affects the portfolio's factor exposure and risk characteristics.

4.2 Equal Weight (1/N)

The equal-weight portfolio assigns $w_i = 1/N$ to each of the N selected stocks. Despite its simplicity, equal weighting has a strong empirical track record. DeMiguel et al. (2009) show that the $1/N$ portfolio frequently outperforms optimized portfolios out of sample across a range of datasets, because it avoids the estimation error inherent in sample-based optimization.

Equal weighting is typically applied within a signal-selected subset: for example, equal-weighting the top quintile of firms ranked by composite score. This combines signal-based selection with naive diversification. The approach implicitly tilts toward smaller firms (which receive the same weight as large firms), introducing an unintended size bet that may or may not be desirable. This size tilt also creates capacity constraints: in broad universes, equal-weighting small and illiquid firms at the same level as large caps limits the total capital the strategy can deploy without excessive market impact.

DeMiguel et al. (2009) quantify the estimation challenge facing optimized alternatives: for a universe of 25 assets, approximately 3,000 months (250 years) of data would be required for mean-variance optimization to reliably outperform $1/N$. This result underscores why simpler weighting schemes remain competitive in finite samples.

4.3 Minimum Variance

The minimum-variance portfolio solves:

$$\min_{\mathbf{w}} \ \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w} \quad \text{s.t. } \mathbf{1}^\top \mathbf{w} = 1, \ \mathbf{w} \geq \mathbf{0} \tag{3}$$

where $\boldsymbol{\Sigma}$ is the covariance matrix of asset returns. This formulation requires only an estimate of the covariance matrix, not expected returns, making it attractive when return forecasts are unreliable. The minimum-variance portfolio is connected to the low-volatility anomaly: the empirical finding that low-volatility stocks have earned higher risk-adjusted returns than the CAPM predicts (Baker et al., 2011; Blitz and van Vliet, 2007).

The primary challenge is estimation error in $\boldsymbol{\Sigma}$. Sample covariance matrices estimated from historical returns are noisy, especially when the number of assets is large relative to the number of time periods. Shrinkage estimators address this by pulling the sample covariance toward a structured target, reducing estimation noise at the cost of some bias. Ledoit and Wolf (2004) propose linear shrinkage toward a single-factor model; more recent nonlinear shrinkage methods (Ledoit and Wolf, 2020) further improve upon linear approaches by applying different shrinkage intensities to different eigenvalues of the sample covariance matrix. The long-only constraint $\mathbf{w} \geq \mathbf{0}$, included here because it reflects the most common practical implementation, also acts as implicit shrinkage: Jagannathan and Ma (2003) show that imposing no-short-selling constraints on the minimum-variance portfolio improves out-of-sample performance by reducing the effect of estimation error, even when the constraints are not "correct" in any economic sense. In practice, long-only minimum-variance portfolios tend to be concentrated in a small number of low-volatility, low-correlation stocks, which may create liquidity and capacity constraints.

4.4 Mean-Variance (Markowitz)

The mean-variance framework of Markowitz (1952) jointly optimizes expected return and risk:

$$\max_{\mathbf{w}} \ \mathbf{w}^\top \boldsymbol{\mu} - \frac{\lambda}{2} \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w} \tag{4}$$

where $\boldsymbol{\mu}$ is the vector of expected returns, $\boldsymbol{\Sigma}$ is the covariance matrix, and λ is a risk-aversion parameter controlling the return-risk trade-off. When the true expected returns and covariance matrix are known, this framework yields the optimal portfolio under either of two sufficient conditions: normally (or elliptically) distributed returns, or quadratic investor utility. In practice, the true parameters must be estimated, which is the source of the difficulties discussed below.

In practice, mean-variance optimization is notoriously sensitive to estimation error in $\boldsymbol{\mu}$. Michaud (1989) describes it as an "error maximizer": small errors in expected return estimates are amplified into extreme portfolio weights. Best and Grauer (1991) show that mean-variance portfolios are more sensitive to changes in expected returns than to changes in covariances. Several robust extensions have been proposed: Bayesian shrinkage of expected returns, resampled efficiency (Michaud, 1998), and the Black-Litterman model (Black and Litterman, 1992), which blends market-implied equilibrium returns with investor views.

4.5 Risk Parity / Equal Risk Contribution

Risk parity seeks a portfolio in which each asset contributes equally to total portfolio risk. The equal risk contribution (ERC) condition requires:

$$w_i \cdot (\mathbf{\Sigma w})_i = w_j \cdot (\mathbf{\Sigma w})_j \quad \text{for all } i, j \tag{5}$$

where $(\mathbf{\Sigma w})_i$ is the i-th element of the vector $\mathbf{\Sigma w}$, so that $w_i \cdot (\mathbf{\Sigma w})_i$ is asset i's risk contribution to total portfolio variance (by the Euler decomposition, portfolio variance equals the sum of all such risk contributions). This condition has no closed-form solution in general but can be solved numerically (Maillard et al., 2010).

Like minimum variance, risk parity requires only a covariance matrix, not return forecasts. It produces more diversified portfolios than minimum variance because it allocates risk evenly rather than minimizing total risk. Maillard et al. (2010) show that ERC portfolios are robust to uncertainty about which assets will deliver the highest returns. Asness et al. (2012) provide a complementary theoretical perspective, arguing that leverage aversion among investors helps explain why risk parity strategies have historically performed well: if investors are reluctant to lever up low-risk assets, those assets become underpriced relative to their risk-adjusted expected returns. A limitation is that risk parity portfolios may require leverage to achieve competitive absolute return levels, since they tend to overweight low-volatility assets and underweight high-volatility assets relative to a capitalization-weighted benchmark.

Table 2: Weighting scheme comparison.

Scheme	Inputs Required	Key Assumption	Strengths	Weaknesses
Equal weight (1/N)	None	No informational advantage	Simple, diversified, robust OOS	Ignores risk a nal strength
Score-proportional	Factor scores	Signal predicts returns	Direct signal expression	Ignores risk st
Minimum variance	Covariance matrix	Variance is minimizable	Lower drawdowns, no return forecasts	Concentrated, mation error
Mean-variance	Returns + covariance	Returns are estimable	Theoretically optimal	Error amplifi extreme weigh
Risk parity (ERC)	Covariance matrix	Equal risk budget is fair	Balanced risk, robust to return uncertainty	May need le for competiti turns

5 From Weights to Portfolios: Constraints and Implementation

5.1 Long-Only vs. Long-Short

Academic factor research typically constructs long-short portfolios: going long stocks in the top quintile (or decile) of a characteristic and shorting stocks in the bottom quintile.

This design isolates the factor premium by hedging out market exposure, but it assumes frictionless short selling, which is unrealistic for many institutional investors.

Long-only strategies with factor tilts are the more common implementation in practice. A tilt-based portfolio starts from a benchmark (e.g., a capitalization-weighted index) and introduces active overweights and underweights proportional to factor scores, subject to a tracking-error budget:

$$\mathrm{TE} = \sqrt{\mathrm{Var}(r_p - r_b)} \tag{6}$$

where r_p and r_b are the portfolio and benchmark returns, respectively. The tracking-error budget controls how far the portfolio deviates from the benchmark, balancing factor exposure against the risk of significant underperformance relative to the index.

Table 3: Long-only tilt vs. long-short.

Dimension	Long-Short	Long-Only Tilt
Capital requirement	Margin for short leg	Fully funded
Shorting constraints	Borrowing costs, recall risk	None
Factor exposure purity	High (hedges market)	Diluted by benchmark holdings
Tracking error	High (benchmark-independent)	Controlled via TE budget
Capacity	Limited by short-side liquidity	Larger (long-only, broad universe)
Implementation complexity	High (prime brokerage, margin)	Moderate (standard custody)

5.2 Constraint Types

Unconstrained optimization frequently produces portfolios that are impractical to implement. Constraints encode real-world requirements into the optimization problem:

- **Position size bounds.** Upper bounds on individual weights prevent excessive concentration. Lower bounds (often zero for long-only portfolios) prevent short positions. Typical constraints might limit any single position to 5% of the portfolio.

- **Sector and industry neutrality.** Constraining sector weights to match the benchmark removes unintended sector bets. Without such constraints, a value signal might concentrate the portfolio in financials, while a momentum signal might overweight technology, introducing exposures unrelated to the intended factor.

- **Turnover limits.** Constraining the fraction of the portfolio that changes at each rebalance controls transaction costs. A turnover constraint can be implemented as a hard limit ($\sum_i |w_i^{\text{new}} - w_i^{\text{old}}| \leq \tau$) or as a penalty term in the objective function.

- **Leverage constraints.** For long-short portfolios, constraints on gross exposure ($\sum_i |w_i| \leq L$) limit the total amount of leverage.

5.3 Transaction Costs and Rebalancing

Every rebalance incurs transaction costs: bid-ask spreads, market impact, and brokerage commissions. The concept of implementation shortfall (Perold, 1988) captures the total cost of executing a trade as the difference between the decision price and the execution price. For large portfolios, market impact, the price movement caused by the trade itself, typically dominates other cost components (Almgren and Chriss, 2001).

Rebalancing frequency should be matched to the signal's half-life. A momentum signal with a half-life of one to three months calls for monthly rebalancing, while a value signal with a half-life of six to twelve months permits quarterly rebalancing. More frequent rebalancing captures signal changes more quickly but incurs higher transaction costs.

Cost-aware optimization incorporates transaction costs directly into the objective function, penalizing turnover rather than constraining it. This approach produces a smooth trade-off between signal freshness and trading costs, and naturally leads to partial rebalancing: the optimizer trades only when the expected benefit of updating a position exceeds its expected cost.

5.4 Backtesting Considerations

Any backtested strategy is subject to biases that can inflate apparent performance. Look-ahead bias arises when the backtest uses information that would not have been available at the time of the trading decision, for example by using point-in-time data that has been subsequently revised. Survivorship bias arises when the backtest universe excludes firms that delisted during the sample period, which tends to remove the worst performers and overstate returns. Lopez de Prado (2018) provides a comprehensive treatment of these and related pitfalls, including the dangers of multiple testing when evaluating many strategy variants on the same data.

6 Machine Learning as an Extension

The pipeline described in the preceding sections, from raw characteristics through normalization, scoring, weighting, and constraints, represents the traditional approach to factor-based portfolio construction. Machine learning enters this pipeline not as a replacement but as a set of tools that can potentially improve each stage.

6.1 ML for Signal Combination

The linear composite score of Section 3.3 can be replaced by nonlinear models: gradient-boosted trees, random forests, or neural networks. Gu et al. (2020) provide a comprehensive comparison of ML methods for predicting cross-sectional stock returns, finding that tree-based models and neural networks outperform linear models out of sample. The key advantage is the ability to capture interactions between signals (e.g., value may predict returns differently for high- vs. low-momentum stocks) without requiring the researcher to specify these interactions in advance.

6.2 Learned Risk Models

Traditional risk models estimate the covariance matrix using historical returns, possibly augmented with factor structure (e.g., the Barra model) and shrinkage. ML approaches can learn richer factor structures from the data. Gu et al. (2021) develop an autoencoder asset pricing model that simultaneously learns latent factors and their loadings from the cross-section of returns; the resulting factor structure can serve as input to both return prediction and risk modeling. These learned factor models can potentially capture time-varying risk dynamics that static factor models miss.

6.3 End-to-End Approaches

The most ambitious ML approaches bypass the explicit pipeline entirely, learning a mapping directly from raw data to portfolio weights. Zhang et al. (2020) explore deep learning models that take raw asset features as input and output portfolio allocations, with the loss function defined in terms of portfolio-level objectives (e.g., Sharpe ratio, maximum drawdown). This end-to-end framing is attractive in principle because it avoids the information loss inherent in each intermediate step of the traditional pipeline. A precursor to end-to-end ML approaches is the parametric portfolio policy framework of Brandt et al. (2009), which maps firm characteristics directly to portfolio weights without an intermediate scoring step. This framework demonstrates that collapsing the multi-step pipeline can be beneficial even with linear models, providing theoretical motivation for the more flexible ML extensions that followed. In practice, however, end-to-end approaches remain an active research frontier: they require large amounts of data, are difficult to interpret, and have limited evidence of reliable out-of-sample performance in live trading environments.

Table 4: ML at each pipeline stage.

Pipeline Stage	Traditional Approach	ML Extension	Key Reference
Signal generation	Firm characteristics	Learned features, alternative data	Gu et al. (2020)
Score combination	Linear composite	Tree/neural ensemble	Gu et al. (2020)
Risk model	Sample covariance + shrinkage	Deep factor models, autoencoders	Gu et al. (2021)
Weight optimization	Convex optimization	End-to-end learning	Zhang et al. (2020)

7 Conclusion

The path from a predictive signal to an investable portfolio is a sequence of engineering decisions, not a monolith. Each stage of the pipeline, from normalization and outlier treatment through composite scoring, weighting, and constraint enforcement, involves choices that materially affect realized portfolio characteristics. A strong signal processed through a poor weighting scheme, or a well-optimized portfolio implemented without regard to transaction costs, can easily underperform a simpler alternative.

The survey of weighting schemes highlights a recurring theme in empirical finance: theoretical optimality does not guarantee practical superiority. Mean-variance optimization is the theoretically correct framework under standard assumptions, yet its sensitivity to estimation error often makes it inferior to simpler approaches like equal weighting or minimum variance in out-of-sample tests. The choice of weighting scheme should be guided by what the investor knows with confidence: if return forecasts are reliable, mean-variance is appropriate; if only risk estimates are trusted, minimum variance or risk parity may be preferable; if neither is available, equal weighting within a signal-selected subset is a robust default.

Machine learning offers tools to improve each stage of this pipeline, from learning nonlinear signal combinations to estimating richer risk models to optimizing portfolio-level objectives end-to-end. However, the traditional pipeline remains the baseline that ML aims to improve. Understanding its structure, its decision points, and the assumptions embedded at each stage is a prerequisite for evaluating when and how ML can add value.

References

Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2):5–39, 2001. doi: 10.21314/JOR.2001.041.

Andrew Ang. *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press, New York, 2014. doi: 10.1093/acprof:oso/9780199959327.001.0001.

Robert D. Arnott, Jason Hsu, and Philip Moore. Fundamental indexation. *Financial Analysts Journal*, 61(2):83–99, 2005. doi: 10.2469/faj.v61.n2.2718.

Clifford S. Asness, Andrea Frazzini, and Lasse Heje Pedersen. Leverage aversion and risk parity. *Financial Analysts Journal*, 68(1):47–59, 2012. doi: 10.2469/faj.v68.n1.1.

Clifford S. Asness, Tobias J. Moskowitz, and Lasse Heje Pedersen. Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985, 2013. doi: 10.1111/jofi.12021.

Malcolm Baker, Brendan Bradley, and Jeffrey Wurgler. Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. *Financial Analysts Journal*, 67(1):40–54, 2011. doi: 10.2469/faj.v67.n1.4.

Rolf W. Banz. The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1):3–18, 1981. doi: 10.1016/0304-405X(81)90018-0.

Michael J. Best and Robert R. Grauer. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: Some analytical and computational results. *The Review of Financial Studies*, 4(2):315–342, 1991. doi: 10.1093/rfs/4.2.315.

Fischer Black and Robert Litterman. Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43, 1992. doi: 10.2469/faj.v48.n5.28.

David Blitz and Pim van Vliet. The volatility effect: Lower risk without lower return. *The Journal of Portfolio Management*, 34(1):102–113, 2007. doi: 10.3905/jpm.2007.698039.

Michael W. Brandt, Pedro Santa-Clara, and Rossen Valkanov. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *The Review of Financial Studies*, 22(9):3411–3447, 2009. doi: 10.1093/rfs/hhp003.

Mark M. Carhart. On persistence in mutual fund performance. *The Journal of Finance*, 52 (1):57–82, 1997. doi: 10.1111/j.1540-6261.1997.tb03808.x.

Roger Clarke, Harindra de Silva, and Steven Thorley. Portfolio constraints and the fundamental law of active management. *Financial Analysts Journal*, 58(5):48–66, 2002. doi: 10.2469/faj.v58.n5.2468.

John H. Cochrane. Presidential address: Discount rates. *The Journal of Finance*, 66(4): 1047–1108, 2011. doi: 10.1111/j.1540-6261.2011.01671.x.

Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The Review of Financial Studies*, 22(5):1915–1953, 2009. doi: 10.1093/rfs/hhm075.

Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993. doi: 10.1016/0304-405X(93) 90023-5.

Eugene F. Fama and Kenneth R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, 2015. doi: 10.1016/j.jfineco.2014.10.010.

Shihao Gu, Bryan Kelly, and Dacheng Xiu. Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273, 2020. doi: 10.1093/rfs/hhaa009.

Shihao Gu, Bryan Kelly, and Dacheng Xiu. Autoencoder asset pricing models. *Journal of Econometrics*, 222(1):429–450, 2021. doi: 10.1016/j.jeconom.2020.07.009.

Campbell R. Harvey, Yan Liu, and Heqing Zhu. ... and the cross-section of expected returns. *The Review of Financial Studies*, 29(1):5–68, 2016. doi: 10.1093/rfs/hhv059.

Ravi Jagannathan and Tongshu Ma. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance*, 58(4):1651–1683, 2003. doi: 10.1111/1540-6261.00580.

Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91, 1993. doi: 10.1111/j.1540-6261.1993.tb04702.x.

Theis Ingerslev Jensen, Bryan T. Kelly, and Lasse Heje Pedersen. Is there a replication crisis in finance? *The Journal of Finance*, 78(5):2465–2518, 2023. doi: 10.1111/jofi.13249.

Olivier Ledoit and Michael Wolf. Honey, I shrunk the sample covariance matrix. *The Journal of Portfolio Management*, 30(4):110–119, 2004. doi: 10.3905/jpm.2004.110.

Olivier Ledoit and Michael Wolf. Analytical nonlinear shrinkage of large-dimensional covariance matrices. *The Annals of Statistics*, 48(5):3043–3065, 2020. doi: 10.1214/19-AOS1921.

Marcos Lopez de Prado. *Advances in Financial Machine Learning*. Wiley, Hoboken, NJ, 2018. doi: 10.1002/9781119482086.

Sébastien Maillard, Thierry Roncalli, and Jérôme Teiletche. The properties of equally weighted risk contribution portfolios. *The Journal of Portfolio Management*, 36(4):60–70, 2010. doi: 10.3905/jpm.2010.36.4.060.

Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952. doi: 10.2307/2975974.

Richard O. Michaud. The Markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, 45(1):31–42, 1989. doi: 10.2469/faj.v45.n1.31.

Richard O. Michaud. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. Harvard Business School Press, Boston, 1998.

Robert Novy-Marx. The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1):1–28, 2013. doi: 10.1016/j.jfineco.2013.01.003.

André F. Perold. The implementation shortfall: Paper versus reality. *The Journal of Portfolio Management*, 14(3):4–9, 1988. doi: 10.3905/jpm.1988.409150.

Zihao Zhang, Stefan Zohren, and Stephen Roberts. Deep learning for portfolio optimization. *The Journal of Financial Data Science*, 2(4):8–20, 2020. doi: 10.3905/jfds.2020.1.042.