

Weight Determination and Portfolio Construction Based on Interval and Distributional Surface Forecasts in Financial Engineering

Yunkun Tu

*College of Liberal Arts and Sciences, University of Illinois Urbana-Champaign, Urbana, USA
jimmytu2@illinois.edu*

Abstract. If portfolio optimization uses only point estimates of returns or risk, estimation error often becomes extreme weights and too much trading. Interval forecasts give ranges for parameters or returns. Distributional forecasts give conditional distributions, quantiles, and tail features. Together they give fuller inputs for choosing weights. This review studies how those inputs enter a weight rule. Interval uncertainty sets, conformal coverage, distributionally robust optimization, and conditional value-at-risk (CVaR) can be written into the objective or the constraints and then mapped into tradable weights. Interval width κ acts as a penalty on effective expected return or as a bound on the feasible set, so a wider interval shrinks active positions. Distribution shape changes risk measurement: a thicker left tail raises CVaR and lowers the weight on that asset. The two channels can be placed in one program from a point forecast $\hat{\mu}$, an interval half width κ , and a conditional law F to a weight vector w under trading cost constraints. Putting interval width and distribution shape into the allocation rule can reduce overconfident concentration and improve risk budgeting and rebalancing. The argument is theoretical. It does not rest on a new backtest.

Keywords: interval forecast, distributional forecast, portfolio construction, CVaR

1. Introduction

In financial engineering, portfolio construction turns views about returns and risk into asset weights. Mean-variance analysis gives the basic tradeoff between return and risk, while conditional value-at-risk (CVaR) focuses on the left tail of losses. The hard part is often the inputs. Expected returns, covariances, and even return distributions are not seen directly and must be estimated or forecast. Once inside the optimizer, these inputs are often treated as known constants. A small revision to one sector mean can then cause large shifts in weights across industries. A portfolio that looks best on paper can amplify estimation error in practice.

This problem shows up clearly in empirical work. DeMiguel et al. compare many optimized rules with equal weights [1]. In out-of-sample tests, the optimized rules do not always beat equal weights in a stable way. Treating point estimates as true parameters can make portfolios unstable. Machine learning has later raised return predictability. Gu et al. compare regularization, trees, and neural nets [2]. Feng et al. integrate characteristic sorting and factor building in one deep learning setup that

aims at pricing errors [3]. Better forecasts still may not raise portfolio results after trading costs. Jensen et al. show that if evaluation and learning ignore trading costs, models lean on short-lived, small-scale signals [4]. Short-term reversal in small stocks is one clear case. Turnover rises and net returns fall. Weight rules therefore need both intervals for how unreliable a view is and distributions for how bad adverse outcomes can be.

Optimal weights change when interval width and distribution shape enter the objective or the constraints. Those changes matter for allocation, risk budgets, and rebalancing. Forecast model rankings by accuracy are outside the scope. The remainder of the paper reviews forecast improvement, interval-based robust allocation, and distributional and tail-risk approaches; introduces the corresponding analytical tools; examines mechanisms and comparative statics; and concludes with practical implications.

2. Literature review

2.1. Forecast improvement and portfolio performance

Return forecasting sits upstream of portfolio inputs. It sets the means, factor exposures, or risk prices that the optimizer uses. Gu et al. compare regularization, tree models, and neural nets for cross sectional risk premia in one out of sample design [2]. Flexible machine learning methods can catch nonlinear links and interactions that linear models miss. Gains differ across methods. Whether those gains become portfolio gains still depends on how forecasts are used. Feng et al. cast characteristic sorting, factor generation, and return fitting as one deep learning process trained on pricing errors [3]. This brings prediction closer to portfolio related pricing goals. These studies suggest that upstream forecasts can improve input quality; however, how uncertain forecasts should be translated into portfolio positions remains an open decision-design question addressed by interval- and distribution-based rules.

Even when predictability rises, portfolio results still depend on how forecasts are used. Kozak et al. show that turning high dimensional characteristics into risk price estimates without sound shrinkage can yield unstable estimates and poor portfolio inputs [5]. Their shrinkage stochastic discount factor (SDF) lowers the role of noisy principal components and gives a more stable base for cross sectional pricing and allocation. Giglio et al. warn from a multiple testing view that data snooping can create significant alphas too quickly [6]. Even in sample significance may shrink out of sample, so false discovery control is needed when screening strategies. Jensen et al. move the issue to trading [4]. They judge strategies by the risk return tradeoff after costs, called the implementable efficient frontier. When costs are ignored, learning favors high turnover, small scale, fast decaying signals. Once costs are counted, their net value falls fast. Beyond sharper point estimates, the range and shape of uncertainty must enter the weight rule. Evaluation should track both gross and net results.

2.2. Interval forecasts and robust allocation

One classic path from interval forecasts into portfolios is to write uncertainty about expected return estimates as a set of priors, then choose under a worst case prior. Garlappi et al. use confidence intervals around estimated expected returns as multiple priors [7]. They use ambiguity aversion to justify minimizing over those priors. The practice point is direct. Wider intervals lower an asset's effective appeal and shrink optimal holdings. In some designs the solution can be expressed as a shrunk combination of the mean variance and minimum-variance portfolios. Unreliability becomes

more conservative positions. Yin et al. discuss practical calibration [8]. Quadratic uncertainty sets are usually smoother than box constraints, the uncertainty matrix is preferably diagonal and built mainly from variances, and the uncertainty level can be set by asset Sharpe ratios. Their contribution is to show how the form and scale of $\mathbf{\Sigma}$ decide whether shrinkage is moderate or too strong. Bryzgalova et al. move uncertainty from one parameter to the model layer [9]. Facing many competing factor models, Bayesian model averaging gives posterior risk prices. When model uncertainty dominates, what is shrunk is often a set of model weights. One asset point estimate is only one piece.

Conformal prediction offers another way to build intervals that fits sequential forecasting better. Angelopoulos et al. cast online conformal prediction for time series as a control problem [10]. Quantile tracking delivers long run coverage, an error integral steadies the coverage path, and score forecasting looks ahead to systematic bias. Interval width can move with forecast errors and distribution drift, and steady errors are only a special case. For portfolios, $\mathbf{\Sigma}$ can act as a state-based risk budget control. Im et al. put conformal style uncertainty sets into constrained contextual optimization when constraint parameters are themselves learned by machine learning and thus uncertain [11]. They propose the Smart Predict-then-Optimize with Robust Constraints (SPO-RC) loss, a feasibility-sensitive loss function. The mechanism shows that uncertainty sets can rewrite the objective and tighten what is feasible. Moved to portfolios, the same logic can apply to leverage, risk, or exposure constraints. The original experiments concern knapsack and production type constraints.

2.3. Distributional forecasts and tail risk

Distributional forecasts focus on the shape of the conditional return law, of which the mean and variance are only two summary numbers. CVaR, widely used in practice, measures average losses beyond a quantile and reacts more than variance to left tail thickness. Two assets can share similar mean and variance point estimates yet get different CVaR optimal weights when left tails differ. Obłój and Wiesel show from a distributionally robust view that even local adverse changes of a baseline measure inside a Wasserstein neighborhood can change optimal investment policies and relative weights [12]. This response does not always rise or fall in a simple way with the baseline Sharpe ratio. Lassance and Vrms address the estimation cost of higher order comoments in large portfolios [13]. The number of comoment parameters grows fast with the number of assets. They project returns onto a few nearly independent factors with independent component analysis (ICA), so fat tails and asymmetry can enter optimization with fewer parameters. They also compare out of sample results after trading costs. Costa and Iyengar put a distributionally robust decision layer into end to end learning [14]. The prediction layer outputs a point forecast. Past forecast errors feed the decision layer to capture variance and model risk. The decision layer takes a worst case over an ambiguity set, and task losses send portfolio result gradients back to the prediction layer, so distribution uncertainty enters weight generation.

Distributional optimization is not automatically conservative. If the forecast distribution understates left-tail risk or the scenario set omits joint crashes, a tail-based optimizer may still select aggressive weights. When tail dependence rises in stress, portfolios that look diversified can still hide concentration risk. Whether the distribution path works depends on whether the conditional law is identified well and whether it truly enters the decision rule. Vuletić and Cont show that implied volatility scenarios under no arbitrage constraints can support option portfolio hedging [15]. If similar scenario ideas are moved to equity or multi asset allocation, they can enrich distribution inputs for the same decision process. That move goes beyond the direct claims of the original paper.

Upstream forecasts can improve point estimates and factor structure, while net results after costs still need separate care. Intervals, multiple priors, conformal coverage, and constraint uncertainty sets turn unreliability into shrinkage and feasible set limits. Distributions, higher order comoments, and distributional robustness put the severity of adverse states into risk measures and weight generation. The next section lays out the tools and how they enter one weight map.

3. Theoretical foundations

3.1. Mean variance analysis and parameter uncertainty

Let R denote the vector of asset excess returns and w the portfolio-weight vector. A mean-variance investor chooses the weight vector w^{MV} defined by

$$w^{MV} \in \arg \max_w (w^T \mu - \frac{\gamma}{2} w^T \Sigma w) \quad (1)$$

where $\mu = E[R]$ is expected excess return, Σ is the covariance matrix, and $\gamma > 0$ is risk aversion. The left side of equation (1) names the chosen weights. The right side is the mean variance program. w^{MV} maximizes expected excess return minus a quadratic penalty on risk. In practice μ and Σ are replaced by estimates.

If sample estimates of μ and Σ replace μ and Σ directly, optimal weights are often highly sensitive to errors in estimated μ . This sensitivity is especially pronounced for errors in expected returns. In the mean variance objective, expected returns determine the coefficients of the linear reward term, whereas the covariance matrix determines the curvature of the quadratic risk term. A small overstatement of the mean can be read by the optimizer as a reason for large active positions. Estimation error is therefore amplified by the objective. That amplification matters as much as the fact that error exists.

The literature offers two workable responses. The first admits that μ , and sometimes the model setup, lies in an uncertainty set, and then chooses under a worst case. Key work includes multiple priors and robust mean variance in Garlappi et al. [7] and Yin et al. [8]. The second uses the return distribution or its tail functional, including CVaR, and further distributionally robust forms in Obłój and Wiesel [12] and Costa and Iyengar [14]. Interval forecasts mainly serve the first path. Distributional forecasts mainly serve the second path.

3.2. Interval uncertainty sets and worst case decisions

Garlappi et al. write parameter uncertainty as confidence intervals around estimated expected returns and form a set of multiple priors [7]. Ambiguity aversion is modeled by minimization over that set. Uncertainty can differ across assets or asset groups, and several common forms admit closed form optimal portfolios. In a special case the multi prior solution equals shrinkage between the mean variance and minimum variance portfolios, so interval width pulls the portfolio away from aggressive mean variance solutions.

For mechanism analysis, write the uncertainty set for expected returns in box form as

$$U = \{\mu \in R^n : |\mu_i - \hat{\mu}_i| \leq \kappa_i, i = 1, \dots, n\} \quad (2)$$

where $\hat{\mu}_i$ is the point estimate for asset i and $\kappa_i \geq 0$ is the interval half width. The set U collects every mean vector in R^n whose i -th coordinate stays within κ_i of that point estimate. κ_i may

come from statistical confidence intervals or from conformal prediction intervals. The robust weight vector w^{rob} is defined by

$$w^{rob} \in \arg \max_w \min_{\mu \in U} (w^T \mu - \frac{\gamma}{2} w^T \Sigma w) \quad (3)$$

The left side of the robust problem in equation (3) names the solution. The outer problem chooses w . The inner problem searches for the least favorable μ in U for a given w . For the box set in equation (2) that inner search has the closed form

$$\min_{\mu \in U} w^T \mu = w^T \hat{\mu} - \kappa^T |w| \quad (4)$$

For a long position the least favorable move lowers the mean, and for a short position it raises the mean; the size of that move is set by κ . Thus κ acts like a penalty on effective expected return, and larger κ implies stronger shrinkage toward safer holdings. If κ rises for a stock with a high estimated mean but a short and noisy return history, its long weight shrinks even when the point forecast stays attractive.

Yin et al. give practical calibration that can be used [8]. Based on a review of mechanisms, they argue that quadratic uncertainty sets are preferable to box constraints. The uncertainty matrix is recommended to be diagonal and built mainly from asset variances. The uncertainty level can be set by asset Sharpe ratios. For the interval path to change positions, the shape and scale of κ must match asset risk features. Calibration that is too wide over-shrinks the portfolio. Calibration that is too narrow returns to point estimate hypersensitivity.

How intervals are built also affects the dynamics of κ . Angelopoulos et al. cast online conformal prediction for time series as a control problem [10]. Quantile tracking gives long run coverage, error integration steadies coverage, and scorecasting looks ahead to systematic score trends. They test the method on series that include individual stock returns of Amazon, Google, and Microsoft. For portfolios, κ may not be a fixed constant. When forecast errors stay large or volatility rises, conformal intervals can widen and automatically cut effective positions.

Im et al. push conformal type uncertainty sets into constrained contextual optimization [11]. They study problems where the constraint parameters themselves are predicted by machine learning and are therefore uncertain. They propose a feasibility sensitive SPO-RC loss and a convex proxy. The lesson for portfolio construction is that uncertainty sets can rewrite return terms in the objective and also enter feasibility checks. Risk, leverage, and exposure constraints are plausible portfolio analogues, although this extension requires separate theoretical or empirical validation.

3.3. Distributional functionals, distributional robustness, and CVaR

Let portfolio loss be $L(w) = -w^T R$. Conditional value at risk at level α is defined by

$$CVaR_{\alpha}(L) = \min_{\eta} \{ \eta + \frac{1}{1-\alpha} E[(L - \eta)_+] \} \quad (5)$$

where α is the confidence level and η is an auxiliary cutoff. Equation (5) defines CVaR as that minimized value in η . The investor then chooses $w^{CVaR} \in \arg \min_w CVaR_{\alpha}(L(w))$. This allocation is more sensitive to left tail thickness than variance. Once the full conditional distribution, or a rich enough set of scenarios, is available, weights can be chosen by minimizing CVaR or by a related tail functional. Assets with thicker left tails then face a stronger penalty.

Lassance and Vrms address the estimation barrier to using higher order comoments in large universes [13]. The number of comoment parameters grows fast with the number of assets, so direct estimation is often not feasible. They use independent component analysis to project returns onto a few nearly independent factors. Comoments can then be estimated with fewer parameters. They also show that after trading costs this approach can improve out of sample risk results relative to principal component analysis (PCA)-based setups.

Obłój and Wiesel show from another angle why distribution ambiguity rewrites weights [12]. They optimize against adverse probability measures within a Wasserstein ball centered at the baseline measure P . They then compute first order sensitivities of the investor value function, the optimal investment policy, and Davis marginal utility prices to model uncertainty. The results are nonparametric. They depend on both the baseline model and risk attitudes. Even local ambiguity can change relative asset weights, and the response may not be monotone in the baseline Sharpe ratio.

Costa and Iyengar embed a distributionally robust decision layer in end-to-end learning [14]. Relative to classical predict then optimize, their decision layer solves a minimax problem whose ambiguity set is centered on a nominal forecast. The prediction layer supplies the point forecast together with past prediction errors, which capture the joint effect of variance and model risk. Convex duality turns the minimax problem into a form that allows backpropagation, so task loss gradients can pass through the decision layer to the prediction layer. They also show that risk tolerance and robustness parameters can be learned from data. Distributional robustness is hard to use if its key parameters are set by guesswork.

3.4. Trading costs, implementability, and constraints

Even when intervals and distributions enter the objective, portfolios may still fail under trading frictions. Jensen et al. argue that strategies should be judged by the risk-return relationship after trading costs, the implementable efficient frontier [4]. If learning is blind to trading costs, models tend to rely on short-lived, small-scale features. A strategy that chases rapid reversals in illiquid names can look strong on paper returns while net returns after fees look weak. A practical response is to incorporate transaction costs into the portfolio objective and align learning with a net economic objective.

The interval channel mainly changes effective returns or constraint bounds through uncertainty sets. The distribution channel mainly changes risk measurement through tail functionals or distribution neighborhoods. Costs and constraints decide whether mapped results can become tradable weights. Im et al. show that when constraint parameters themselves are uncertain, uncertainty sets built by conformal and related methods can enter feasibility checks [11]. That mechanism can move to portfolio goals of target correction and feasible set shrinkage.

4. Model construction and mechanism analysis

4.1. A unified framework from forecasts to tradable weights

Weight determination is one process from forecasts to a tradable portfolio. Forecasting yields a point estimate $\hat{\mu}$, an interval half width K , and a conditional distribution F or a usable summary in the form of quantiles, higher moments, or scenarios. Representation then turns these into objects the optimizer can use. Intervals become box, quadratic, or other uncertainty sets $U(K)$. Distributions become scenarios, quantiles, comoments, or empirical error distributions. What stays invisible at this stage can hardly rewrite weights later. The mapping can be written as one program

$$w^* \in \arg \max_{w \in W} \{ \min_{\mu \in U(\kappa)} w^T \mu - \lambda \rho(L(w); F) \} \quad (6)$$

The left side of equation (6) names the chosen weights. The right side is the program that produces them. Here W is the feasible set after budget, leverage, exposure, and trading cost constraints. The inner minimum uses the interval channel: $\hat{\mu}$ is the point forecast of expected excess return, and κ is the interval half width from a confidence interval or a conformal interval, which forms $U(\kappa)$ as in the box uncertainty set in equation (2). The risk term $\rho(L(w); F)$ is the distribution channel. It may be variance, CVaR as defined in equation (5), or a distributionally robust functional, and F is the estimated conditional law of R , obtained from scenarios, quantiles, or comoments. The scalar $\lambda > 0$ converts risk into the same units as expected return. Each object must be estimated before the program is solved: $\hat{\mu}$ from a return model, κ from residual coverage or sampling error, F from a distributional forecast, and W from the trading desk constraints. If any of these stays outside the program, it cannot change w^* .

After representation, mapping turns uncertainty into first order conditions, penalties, or constraint bounds and solves for weights under the unified mapping in equation (6), or under a special case of robust mean variance, CVaR, model averaging, or an end-to-end economic goal. How this step uses κ and F decides the mechanism. Execution comes last. Turnover, leverage, sector, and cost constraints turn paper weights into a tradable portfolio. If mapping still uses only point estimates and execution later cuts positions with hard constraints, earlier interval and distribution information is often discarded. Uncertainty should enter the objective or constraints at mapping, before the trade list is made.

4.2. How intervals change optimal weights

Under the robust problem in equation (3), the worst-case mean of a long position equals

$$\mu_i^{wc} = \hat{\mu}_i - \kappa_i \quad (7)$$

With only the point estimate the optimizer turns optimistic forecasts into large positions. After κ enters, the effective mean of a long asset equals the point estimate minus κ , so the same point estimate maps to lower weight as κ rises. When κ is large enough, the asset may leave the portfolio. The interval thus acts as a continuous penalty or hard threshold on position size at decision time.

This matches the shrinkage reading in Garlappi et al. [7]. Ambiguity aversion pulls the portfolio from chasing estimated means toward more stable, often minimum variance, holdings. Yin et al. show that shrinkage strength depends on the form and level of the uncertainty set [8]. With their suggested quadratic form and Sharpe ratio-based uncertainty levels, shrinkage is smoother. It is also less likely to cut some assets fully while leaving others almost untouched, as can happen under box constraints.

If κ comes from conformal prediction, its time variation has economic content. Under the coverage-control framework of Angelopoulos et al., repeated coverage violations may increase the conformal quantile and thereby widen subsequent prediction intervals [10]. High volatility periods often come with more frequent misses, so intervals may widen as well. In a month with sharp market swings, wider conformal bands can cut active stock weights at the next rebalancing date. The effect is like raising risk aversion, or lowering risk tolerance. The basis is observed coverage error and interval width.

When uncertainty is mainly at the model level, Bayesian model averaging in Bryzgalova et al. [9] and shrinkage priors in Kozak et al. [5] give parallel mechanisms. They reduce concentrated exposure to one fragile factor setup or to high dimensional noisy signals. What is shrunk or averaged is mainly model weights or risk prices, not necessarily one asset point estimate. The shared feature of the interval path remains. Unreliability is written as a penalty on positions.

4.3. How distributions change optimal weights

Distributions change how risk is measured. Consider CVaR as defined in equation (5) again. Suppose two stocks have similar mean and variance point estimates, but one shows a thicker left tail after a market stress month. Under a CVaR goal, that stock has a higher marginal risk contribution. Optimal weight on it is usually lower, or extra hedges are needed. Lassance and Vrins show that once higher order comoments are estimated in a stable way and enter the objective, relative weights in large portfolios can shift clearly relative to mean variance only rules [13].

Distributional robustness pushes the logic further. The sensitivity results in Obłój and Wiesel imply that even local changes around a baseline distribution can change optimal relative weights [12]. Intuitively, the investor looks at nominal tail loss and also at nearby cases where the true distribution is a bit worse. Fragile concentrated holdings are then less likely to be optimal.

After Costa and Iyengar put the idea in a trainable system, the mechanism becomes concrete [14]. The prediction layer outputs a point forecast and uses past forecast errors to form a deviation risk. The decision layer takes a worst case over an ambiguity set and trades it off against expected return. Task losses then send out of sample portfolio result gradients back to the prediction layer. The prediction module is thus pushed to reduce errors that truly harm worst case risk. Directions of mean squared error are no longer treated the same. If the distribution is used only to report value-at-risk (VaR) or plot a histogram and never enters the objective, constraints, or task loss, its effect on weights is essentially zero.

4.4. How intervals and distributions interact, and when they conflict

Intervals address uncertainty in location parameters or forecast errors and suit shrinkage, screening, and constraint calibration. Distributions address asymmetry, fat tails, and tail dependence and suit risk budgeting and tail sensitive rebalancing. A useful order is to shape the tail first, then address thin confidence; or to shrink the feasible region first and then optimize a distributional objective inside it.

One common order is to obtain quantiles or CVaR from a distribution and then attach intervals to those quantile estimates to reflect estimation error in the tail statistics themselves. Another is to obtain a feasible weight set under interval constraints and then optimize a CVaR or distributionally robust objective inside that set. The constraint side uncertainty sets stressed by Im et al. align with the second sequence: first shrink the feasible set, and in portfolios leverage or exposure constraints can use the same idea [11].

Signal conflict should be avoided. If intervals are wide while the nominal distribution is too thin, shrinkage calls for less risk while the tail signal looks mild, and mapped weights become unstable. If intervals are narrow while the left tail is thick, the parameter layer is overconfident while the tail layer demands strong avoidance, and the portfolio swings across rules. Therefore, κ and the distribution are best drawn from a compatible forecasting system in which the same model produces a point forecast, conformal intervals, and quantiles, and these outputs are jointly calibrated inside one mapping rule.

4.5. Why prediction must align with decisions

If training minimizes only point forecast error, misalignment can arise. The model may overfit in directions that barely affect weights, while large errors remain in directions that do affect weights. Feng et al. connect characteristics, factors, and returns through a pricing error goal so that forecasts track portfolio related pricing structure [3]. Jensen et al. judge implementable strategies by net results, avoiding cases where forecasts look stronger but net portfolios look weaker [4]. Costa and Iyengar [14] and Im et al. [11] admit uncertainty into optimization and training through a distributionally robust decision layer and constraint feasibility, respectively.

After alignment, interval width is more likely to attach to parameters that move positions, and distribution quality is more likely to improve regions that enter CVaR or worst-case risk. The prediction loss should stay close to the objects the final weight rule uses.

4.6. Application implications, comparative statics and workflow

In multi asset allocation, intervals can limit exposure in highly uncertain directions by reducing active weights on themes with wide intervals. In tail management, distributions can set loss thresholds and fund protective overlays, and put spreads around an equity book are one concrete case. In rolling rebalancing, distributionally robust rules adjust risk sensitivity while intervals curb overtrading driven by short lived signals. Vuletić and Cont show that implied volatility scenario generation under no arbitrage constraints can support option portfolio hedging [15]. Using such scenarios as distribution inputs in equity or multi asset portfolios is a process level extension, not a setup result validated in the original paper.

Given point estimates, introducing intervals typically moves the portfolio toward shrinkage. Introducing distributions, especially thicker left tails or distributional robustness, typically moves it toward better tail profiles. Used together, the distribution can define a candidate set under a risk preference, and intervals can remove extreme solutions with weak confidence. The resulting weights may fall short of paper optimality when costs are ignored, whether they are easier to trade is an empirical question that this review does not test; the mechanism only says that cost and interval limits are written in before the trade list is made.

For the same central forecast, narrow intervals and thin tails can support larger active weights. Wide intervals and thick tails call for lower exposure or more hedging.

5. Conclusion

An interval and a distribution are added to the objective function or constraints to modify the portfolio weights. Assuming the point estimates are known will increase the estimation error and result in serious allocation problems and high turnover. Interval information deals with parameter uncertainty by penalising weak expected-return signals or narrowing the feasible set, and distributional forecasts consider the extent and dependence of bad outcomes through tail-sensitive risk measures. Assets with heavy left tails or strong downside dependence should be given less weight and may need to be hedged additionally. A coherent portfolio process should incorporate both sources of uncertainty throughout forecasting, representation, optimization, and implementation while maintaining an explicit expected-return objective.

In practice, interval width and distributional shape should be incorporated into the portfolio objective or constraints alongside transaction costs. Portfolio candidates should first be assigned

weights according to the investor's risk-aversion level, after which interval information can be used to penalize extreme positions that lack sufficient statistical support.

References

- [1] DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *The Review of Financial Studies*, 22(5), 1915–1953. <https://doi.org/10.1093/rfs/hhm075>
- [2] Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223–2273. <https://doi.org/10.1093/rfs/hhaa009>
- [3] Feng, G., He, J., Polson, N. G., & Xu, J. (2024). Deep learning in characteristics-sorted factor models. *Journal of Financial and Quantitative Analysis*, 59(7), 3001–3036. <https://doi.org/10.1017/S0022109023000893>
- [4] Jensen, T. I., Kelly, B., Malamud, S., & Pedersen, L. H. (2026). Machine learning and the implementable efficient frontier. *The Review of Financial Studies*, hhag022. <https://doi.org/10.1093/rfs/hhag022>
- [5] Kozak, S., Nagel, S., & Santosh, S. (2020). Shrinking the cross-section. *Journal of Financial Economics*, 135(2), 271–292. <https://doi.org/10.1016/j.jfineco.2019.06.008>
- [6] Giglio, S., Liao, Y., & Xiu, D. (2021). Thousands of alpha tests. *The Review of Financial Studies*, 34(7), 3456–3496. <https://doi.org/10.1093/rfs/hhaa111>
- [7] Garlappi, L., Uppal, R., & Wang, T. (2007). Portfolio selection with parameter and model uncertainty: A multi-prior approach. *The Review of Financial Studies*, 20(1), 41–81. <https://doi.org/10.1093/rfs/hhl003>
- [8] Yin, C., Perchet, R., & Soupé, F. (2021). A practical guide to robust portfolio optimization. *Quantitative Finance*, 21(6), 911–928. <https://doi.org/10.1080/14697688.2020.1849780>
- [9] Bryzgalova, S., Huang, J., & Julliard, C. (2023). Bayesian solutions for the factor zoo: We just ran two quadrillion models. *The Journal of Finance*, 78(1), 487–557. <https://doi.org/10.1111/jofi.13197>
- [10] Angelopoulos, A. N., Candès, E. J., & Tibshirani, R. J. (2023). Conformal PID control for time series prediction. *Advances in Neural Information Processing Systems*, 36, 23047–23074. <https://doi.org/10.52202/075280-1000>
- [11] Im, H., Benslimane, W., & Grigas, P. (2025). Smart surrogate losses for contextual stochastic linear optimization with robust constraints. *Advances in Neural Information Processing Systems*, 38, 83686–83711. <https://doi.org/10.52202/085713-2523>
- [12] Oblój, J., & Wiesel, J. (2021). Distributionally robust portfolio maximization and marginal utility pricing in one period financial markets. *Mathematical Finance*, 31(4), 1454–1493. <https://doi.org/10.1111/mafi.12337>
- [13] Lassance, N., & Vrin, F. (2021). Portfolio selection with parsimonious higher comoments estimation. *Journal of Banking & Finance*, 126, 106115. <https://doi.org/10.1016/j.jbankfin.2021.106115>
- [14] Costa, G., & Iyengar, G. N. (2023). Distributionally robust end-to-end portfolio construction. *Quantitative Finance*, 23(10), 1465–1482. <https://doi.org/10.1080/14697688.2023.2236148>
- [15] Vuletić, M., & Cont, R. (2024). VolGAN: A generative model for arbitrage-free implied volatility surfaces. *Applied Mathematical Finance*, 31(4), 203–238. <https://doi.org/10.1080/1350486X.2025.2471317>