Model Risk in Portfolio Optimization: A Multi-Agent Framework for Dynamic Portfolio Optimization

Omari Purtukhia
Department of Mathematics
Faculty of Exact and Natural Sciences and
Andrea Razmadze Mathematical Institute
Ivane Javakhishvili Tbilisi State University
Tbilisi, Georgia
o.purtukhia@gmail.com

Valeriane Jokhadze
Faculty of Economics and Business
Ivane Javakhishvili Tbilisi State University
Tbilisi, Georgia
jokhadze.valeriane@gmail.com

December 8, 2025

Abstract

Financial markets are complex, high-dimensional, and constantly evolving, posing significant challenges to traditional portfolio optimization methods. Classical mean–variance frameworks often rely on restrictive assumptions (such as Gaussian returns and static parameters estimated from historical data) that can lead to unstable and suboptimal portfolios when markets exhibit non-linear dependencies, heavy tails, or regime shifts. In this paper, we propose a novel multi-agent simulation and model selection system (MAS^2) for dynamic portfolio optimization that addresses two key issues: (i) robust parameter estimation under limited, noisy data, and (ii) adaptive model selection in non-stationary market conditions. In our framework, multiple agents, each based on a distinct stochastic model class, collaboratively learn and adapt by generating synthetic market data, updating their model parameters via Bayesian inference, and undergoing rigorous performance evaluation using Bayesian model selection criteria. Poorly performing models are iteratively pruned, and the process continues until convergence to a stable set of models and portfolio strategies. The resulting portfolio allocations inherently account for model uncertainty and adapt to changing market dynamics, offering a robust alternative to static single-model approaches. We present the framework and discuss its properties; empirical assessment is left for future work.

1 Introduction

Portfolio optimization is a central problem in quantitative finance, classically formulated in the mean–variance framework of [5]. In this paradigm, investors trade off expected return against variance under relatively simple distributional assumptions on asset returns, often Gaussian or more generally elliptical, and risk is summarized by the covariance matrix. A wide range of extensions—including alternative risk measures, constraints, and Bayesian formulations—remain conceptually anchored in this tradition; see [3] for an overview.

In practice, however, empirical return distributions deviate markedly from these assumptions. Financial time series exhibit skewness, excess kurtosis (fat tails), volatility clustering, and regime changes that are not well captured by static, low-dimensional Gaussian models. At the same time, the estimation of key inputs such as expected returns, volatilities, and correlations is notoriously unstable when the cross-sectional dimension of the portfolio is large relative to the available time series length. As emphasized by [4], such misspecification and estimation error give rise to *model risk*: optimal allocations computed under an assumed model can be highly sensitive to its parameters and structure, and thus perform poorly out-of-sample.

Two fundamental challenges arise from this perspective. First, there is the challenge of *robust parameter estimation under limited and unstable data*. High-dimensional financial systems with non-linear dependencies and heavy-tailed risks require estimating many parameters (for example, in Normal–Inverse Gaussian or Heston-type specifications) from samples that may be short, noisy, or non-stationary. Even modest estimation errors in such settings can propagate into extreme and unstable portfolio weights. Second, there is the challenge of *adaptive model selection in non-stationary markets*. Financial dynamics evolve over time: a model that fits well in a tranquil, low-volatility regime may fail abruptly during crises or structural breaks. No single model specification dominates across all conditions, so a static reliance on a single "best" model exposes the investor to substantial model risk. An effective portfolio optimization framework must therefore both stabilize parameter estimation and continually reassess its model assumptions as new data arrive and market regimes change.

Several strands of the literature address aspects of these issues. One line of work tackles model uncertainty through Bayesian model averaging and Bayesian model selection, in which multiple candidate models are entertained simultaneously and combined or compared using posterior-based criteria, such as information measures akin to DIC or WAIC; see, for example, [1]. Another line develops robust and distributionally robust portfolio optimization, where the investor guards against worst-case realizations over an ambiguity set of models or distributions [2]. These approaches provide formal tools to mitigate model risk, but they are typically formulated as static or one-shot optimization problems and do not explicitly exploit synthetic data generation or interactive learning between models.

In parallel, multi-agent systems and agent-based models have been widely used to study financial markets from a bottom-up perspective, including price formation, market microstructure, and emergent phenomena; see [6] for an overview. Most of this literature focuses on behavioral or market-level dynamics rather than on model-driven portfolio optimization under explicit model uncertainty.

In this paper, we propose a *Multi-Agent Simulation and Model Selection System* (MAS^2) that sits at the intersection of these strands. MAS^2 employs an ensemble of autonomous agents, each associated with a distinct stochastic asset model, that jointly learn from data and compete for relevance. At each iteration, agents are calibrated via Bayesian inference, generate synthetic market scenarios, and are evaluated using Bayesian model selection criteria in the spirit of [1]. Poorly performing models are pruned, while the surviving agents' portfolio recommendations are aggregated into a robust allocation. In this way, MAS^2 provides a dynamic, data-adaptive architecture for portfolio optimization that explicitly targets model risk: it combines Bayesian model comparison, synthetic data augmentation, and multi-agent interaction into a unified iterative framework. The remainder of the paper formally defines this framework and details its algorithmic implementation; empirical assessment is left for future work.

2 Methodology: Formal Definition of the System

We formally define the Multi-Agent Simulation and Model Selection System (MAS^2) as a tuple

$$\mathcal{S} = (\mathcal{A}, \mathcal{L}, \mathcal{D}, \mathcal{T}, \mathcal{E}, \mathcal{C}) ,$$

where each element represents a crucial aspect of the framework:

- $\mathcal{A} = \{A_1, A_2, \ldots, A_K\}$ is the set of K interacting agents. Each A_k is an autonomous software agent responsible for optimizing a portfolio according to a particular model class.

- $\mathcal{L} = \{L_1, L_2, \ldots, L_K\}$ is the set of K stochastic model classes considered (e.g., L_1 could be a Geometric Brownian Motion model, L_2 a Normal–Inverse Gaussian model, L_3 a Variance Gamma model, L_4 a Heston stochastic volatility model, etc.). Each agent A_k is associated with one model class L_k and uses it to simulate asset dynamics and estimate riskreturn profiles.

- $\mathcal{D} = \{D_{\text{real}}, D_{\text{synth}}^{(t)}\}$ denotes the family of data sources available at any iteration t. This includes the original real-world market data D_{real} and the synthetic data $D_{\text{synth}}^{(t)}$ generated by the agents during iteration t. At the start ($t = 0$), only real data are available, so $D^{(0)} = D_{\text{real}}$. For $t \geq 1$, $D^{(t)}$ denotes the combined dataset used for calibration at iteration t, typically

$$D^{(t)} = D_{\text{real}} \cup D_{\text{synth}}^{(t)}.$$

- $\mathcal{T}^{(t)} = \{\mu^{(t)}, \mathbf{\Sigma}^{(t)}, \mathbf{\Phi}^{(t)}, \mathbf{\Psi}^{(t)}\}$ is the collection of empirical moment tensors estimated from the current dataset $D^{(t)}$. Here $\mu^{(t)}$ is the vector of first moments (means of asset returns), $\mathbf{\Sigma}^{(t)}$ is the covariance matrix (second moments), $\mathbf{\Phi}^{(t)}$ represents third-order central moments (skewness and co-skewness), and $\mathbf{\Psi}^{(t)}$ represents fourth-order central moments (kurtosis and co-kurtosis). These moments up to the fourth order are tracked because the agents' utility functions (described below) depend on features of the return distribution beyond mean and variance.

- $\mathcal{E}$ is the simulation environment shared by all agents. This environment takes as input the parameter vectors (such as Θ_k for agent A_k) and generates simulated asset price paths or return scenarios via Monte Carlo methods. In essence, $\mathcal{E}$ codifies the common market conditions and mechanics under which the agents operate, providing a testing ground for their strategies using either real or synthetic data.

- $\mathcal{C}$ is the set of model evaluation criteria used for selection, such as the widely applicable information criterion (WAIC) or the deviance information criterion (DIC). These criteria allow the system to compare models on a statistically rigorous basis, accounting for both in-sample fit and model complexity. $\mathcal{C}$ guides the elimination of poorly performing agents over the course of the simulation.

This tuple $\mathcal{S}$ encapsulates the architecture of our system. The agents $\mathcal{A}$, each with their model in $\mathcal{L}$, interact through the environment $\mathcal{E}$ and update using data $\mathcal{D}$; they compute portfolio strategies based on the moments $\mathcal{T}^{(t)}$; and their continued participation is determined by the criteria $\mathcal{C}$. For notational simplicity, we will often omit the superscript (t) on $\mu^{(t)}$, $\mathbf{\Sigma}^{(t)}$, $\mathbf{\Phi}^{(t)}$, and $\mathbf{\Psi}^{(t)}$ when the iteration index is clear from the context.

2.1 Model Universe and Parameterization

Each model $L_k \in \mathcal{L}$ operates on a financial *universe* of N assets (in our empirical implementation, we consider $N = 10$ assets, though the framework generalizes to any N). The model L_k provides a probabilistic description of the joint dynamics of these N asset returns. For instance, one model might assume that asset returns follow correlated geometric Brownian motions (with drift and volatility parameters), while another might assume a jumpdiffusion or heavy-tailed distributional form. We denote by Θ_k the vector of parameters characterizing model L_k. This parameter vector could include, for example, drift and volatility for each asset (in a Gaussian model), or tail thickness and skew parameters (in a heavy-tailed model like NIG), or mean-reversion speed and volatility-of-volatility (in a stochastic volatility model like Heston), among others.

At each iteration t, agent A_k uses Bayesian inference to update its beliefs about Θ_k given the latest dataset $D^{(t)}$. Specifically, an MCMC routine is employed to draw samples from the posterior distribution $P(\Theta_k \mid D^{(t)})$. This posterior captures the range of plausible parameter values for model L_k given the observed data (both real and synthetic). The agent can then summarize this posteriorfor instance, by taking the mean or mode to obtain a point estimate $\bar{\Theta}_k^{(t)}$which will be used for subsequent steps such as simulation and optimization. Using the full posterior (rather than just a point estimate) in principle allows the agent to quantify parameter uncertainty; however, for tractability in the simulation environment $\mathcal{E}$, we typically propagate a representative parameter set such as $\bar{\Theta}_k^{(t)}$ forward.

2.2 The Objective Function (Utility)

Each agent A_k is tasked with selecting an optimal portfolio allocation (weight vector) for the N assets according to its model L_k and the current data-driven estimates of the return distribution. Let

$$\mathbf{w}_k = (w_{k1}, w_{k2}, \ldots, w_{kN})$$

denote the weight vector that agent A_k allocates to the N assets (where, for example, w_{ki} is the fraction of capital invested in asset i by agent A_k). The agent's goal is to maximize a utility function

$$U_k(\mathbf{w}_k; \mathcal{T}^{(t)}, \gamma, \lambda, \eta)$$

that depends on the portfolio weights, the empirical moments $\mathcal{T}^{(t)}$, and preference parameters (γ, λ, η).

We assume that each agent A_k evaluates portfolios through a higher-moment meanvarianceskewnesskurtosis preference. Given a weight vector $\mathbf{w}_k \in \mathbb{R}^N$, the agent's utility is defined as

$$U_k\big(\mathbf{w}_k; \mathcal{T}^{(t)}, \gamma, \lambda, \eta\big) \;=\; \mu_p(\mathbf{w}_k) \;-\; \frac{\gamma}{2}\,\sigma_p^2(\mathbf{w}_k) \;+\; \lambda\, S_p(\mathbf{w}_k) \;-\; \eta\, K_p(\mathbf{w}_k), \tag{1}$$

where $\mu_p(\mathbf{w}_k)$ is the expected portfolio return, $\sigma_p^2(\mathbf{w}_k)$ its variance, $S_p(\mathbf{w}_k)$ its skewness, and $K_p(\mathbf{w}_k)$ its kurtosis. The preference parameters $\gamma > 0$, λ, and $\eta > 0$ govern, respectively, risk aversion to variance, preference for skewness, and aversion to kurtosis (tail risk). For simplicity, we keep (γ, λ, η) common across agents, although they could in principle be agent-specific.

Let $R \in \mathbb{R}^N$ denote the vector of asset returns, and $R_p = \mathbf{w}_k^\top R$ the corresponding portfolio return. The empirical moments of R_p appearing in (1) are computed from $\mathcal{T}^{(t)}$ as follows:

$$\mu_p(\mathbf{w}_k) = \mathbf{w}_k^\top \mu, \tag{2}$$

$$\sigma_p^2(\mathbf{w}_k) = \mathbf{w}_k^\top \mathbf{\Sigma}\, \mathbf{w}_k, \tag{3}$$

$$m_{3,p}(\mathbf{w}_k) = \sum_{i,j,\ell=1}^{N} w_{ki} w_{kj} w_{k\ell}\, \Phi_{ij\ell}, \tag{4}$$

$$m_{4,p}(\mathbf{w}_k) = \sum_{i,j,\ell,m=1}^{N} w_{ki} w_{kj} w_{k\ell} w_{km}\, \Psi_{ij\ell m}, \tag{5}$$

where $m_{3,p}$ and $m_{4,p}$ denote the third and fourth central moments of R_p. The portfolio skewness and kurtosis are then given by

$$S_p(\mathbf{w}_k) \;=\; \frac{m_{3,p}(\mathbf{w}_k)}{\sigma_p^3(\mathbf{w}_k)}, \qquad K_p(\mathbf{w}_k) \;=\; \frac{m_{4,p}(\mathbf{w}_k)}{\sigma_p^4(\mathbf{w}_k)}. \tag{6}$$

The inclusion of these higher moments allows the agent to express preferences such as a liking for positive skewness and an aversion to fat tails, going beyond classical meanvariance utility.

Given this specification, the optimal portfolio for agent A_k solves

$$\mathbf{w}_k^* \;=\; \arg\max_{\mathbf{w}_k}\; U_k\big(\mathbf{w}_k; \mathcal{T}^{(t)}, \gamma, \lambda, \eta\big), \tag{7}$$

subject to appropriate portfolio constraints, such as full investment $\sum_{i=1}^{N} w_{ki} = 1$ and, if desired, no short-selling $w_{ki} \geq 0$. By tuning the parameters (γ, λ, η), the framework can represent different investor preferences or regulatory requirements.

2.3 Convergence and Stable Agent Ensemble

The multi-agent system is designed such that, over successive iterations, it converges to a *stable ensemble* of models and portfolio strategies. Rather than attempting a full game-theoretic treatment, we work with an operational notion of equilibrium: a fixed point of the iterative Bayesian updating, model selection, and portfolio optimization procedure.

To make this precise, let $\Theta_k^{(t)}$ denote the parameter vector of model L_k at iteration t, and let $P_k^{(t)}$ denote the corresponding posterior predictive distribution for one-period portfolio returns implied by agent A_k. We further denote by $\mathbf{w}_k^{*(t)}$ the optimal portfolio for agent A_k at iteration t, obtained by solving

$$\mathbf{w}_k^{*(t)} = \arg\max_{\mathbf{w}_k} U_k\big(\mathbf{w}_k; \mathcal{T}^{(t)}, \gamma, \lambda, \eta\big), \tag{8}$$

subject to the portfolio constraints.

We say that MAS^2 has *converged* when the following conditions hold simultaneously for all active agents over a window of recent iterations:

- **Stability of posterior predictive distributions:** for each surviving agent A_k, the predictive distribution $P_k^{(t)}$ changes only marginally between iterations, in the sense that
$$d\big(P_k^{(t)}, P_k^{(t-1)}\big) < \varepsilon_P,$$
where $d(\cdot,\cdot)$ is a chosen discrepancy measure (e.g., an L^1 distance on predictive densities or a set of moment-based diagnostics) and $\varepsilon_P > 0$ is a tolerance.

- **Stability of portfolio recommendations:** for each surviving agent A_k, the optimal portfolio vector changes only marginally between iterations:
$$\big\|\mathbf{w}_k^{*(t)} - \mathbf{w}_k^{*(t-1)}\big\|_2 < \varepsilon_w,$$
for a given tolerance $\varepsilon_w > 0$.

- **Stability of the active model set:** the set of surviving agents $\mathcal{A}^{(t)}$ does not change over a prespecified number of iterations, i.e.,
$$\mathcal{A}^{(t)} = \mathcal{A}^{(t-1)} = \cdots = \mathcal{A}^{(t-M)}$$
for some integer $M \geq 1$.

In this regime, the system has reached a fixed point in the sense that additional rounds of Bayesian updating, model evaluation, and synthetic data generation no longer lead to material changes in either the predictive distributions or the portfolio allocations. The resulting ensemble of agents can be informally viewed as an *approximate equilibrium*: each surviving model is locally optimal with respect to its own utility, given the shared data environment generated by the ensemble, and no model is eliminated by the selection criterion.

The final portfolio $\mathbf{w}^*$ implemented by MAS^2 is then obtained by aggregating the stabilized agent-specific portfolios; for example, one may take a simple or weighted average,

$$\mathbf{w}^* = \sum_{A_k \in \mathcal{A}^{(\infty)}} \pi_k \, \mathbf{w}_k^{*(\infty)},$$

where the weights π_k can be chosen based on model selection scores (e.g., normalized inverse information criteria) or set uniformly. This aggregated strategy reflects a consensus of the surviving models and incorporates model uncertainty through the ensemble structure, rather than relying on a single specification.

For clarity, the following algorithm summarizes the iterative learning and model-selection procedure implemented in the MAS^2 framework.

Algorithm: Iterative Bayesian Model Selection and Portfolio Optimization

Input: Initial real market dataset D_{real}, set of candidate models $\mathcal{L} = \{L_1, \ldots, L_K\}$, simulation environment $\mathcal{E}$, convergence tolerances $(\varepsilon_P, \varepsilon_w)$ and window size M.
Output: Final portfolio weights $\mathbf{w}^*$ from the converged agent ensemble.

1. Initialize the iteration counter $t \leftarrow 0$. Initialize all agents $A_k \in \mathcal{A}$ with data $D^{(0)} \leftarrow D_{\text{real}}$.

2. Compute empirical moment tensors

$$\mathcal{T}^{(t)} = \{\mu^{(t)}, \mathbf{\Sigma}^{(t)}, \mathbf{\Phi}^{(t)}, \mathbf{\Psi}^{(t)}\}$$

   from the dataset $D^{(t)}$.

3. For each active agent $A_k \in \mathcal{A}$, perform MCMC sampling to obtain the posterior

$$P(\Theta_k^{(t)} \mid D^{(t)}),$$

   and extract a representative parameter estimate $\bar{\Theta}_k^{(t)}$ and the posterior predictive distribution $P_k^{(t)}$.

4. For each active agent A_k, solve the optimization problem

$$\mathbf{w}_k^{*(t)} = \arg\max_{\mathbf{w}_k} U_k\big(\mathbf{w}_k; \mathcal{T}^{(t)}, \gamma, \lambda, \eta\big),$$

   subject to portfolio constraints (e.g., $\sum_i w_{ki} = 1$ and $w_{ki} \geq 0$ if short-selling is disallowed).

5. For each active agent A_k, compute the model selection criterion

$$\text{IC}_k^{(t)} \quad \text{(e.g., WAIC)},$$

   based on $D^{(t)}$ and the posterior samples. Remove any agent whose criterion value is significantly worse than a predefined threshold or dominated relative to other agents, updating the active set $\mathcal{A}^{(t)}$.

6. If only one agent remains in $\mathcal{A}^{(t)}$, optionally terminate the procedure early or continue refining this single model.

7. For each remaining agent A_k, use its current parameter estimate $\bar{\Theta}_k^{(t)}$ to simulate N_{sim} synthetic price paths in the environment $\mathcal{E}$. Store the resulting dataset as $D_{\text{synth}}^{(t+1)}$.

8. Update the dataset for the next iteration, for example via

$$D^{(t+1)} \leftarrow D_{\text{real}} \cup D_{\text{synth}}^{(t+1)},$$

   or using a weighted combination that preserves the influence of real data.

9. Increment the iteration counter: $t \leftarrow t + 1$.

10. Check the convergence criteria based on the stability of $\{P_k^{(t)}\}$, $\{\mathbf{w}_k^{*(t)}\}$, and the active set $\mathcal{A}^{(t)}$ over the last M iterations. If convergence is not satisfied, return to Step 2.

11. Once convergence is reached, construct the final robust portfolio weights $\mathbf{w}^*$ (for example, a weighted average of $\{\mathbf{w}_k^{*(t)}\}$ from the surviving agents at iteration t).

3 Conclusion

We introduced MAS^2, a multi-agent framework for dynamic portfolio optimization designed to address two persistent challenges in quantitative finance: parameter uncertainty and model uncertainty. By allowing heterogeneous model classes to coexist, learn, and compete within a unified Bayesian updating and model selection loop, the framework avoids reliance on any single probabilistic specification. Instead, MAS^2 integrates multiple perspectives on market dynamics and adaptively refines the set of candidate models through iterative posterior inference and performance-based pruning.

A key feature of the system is the use of synthetic data generation to enrich the information set available to each agent. Combined with Bayesian model comparison, this mechanism encourages the ensemble to gravitate toward a stable subset of models whose predictive distributions and implied portfolio allocations change only minimally across successive iterations. The resulting fixed point represents a data-driven consensus among the surviving agents, providing a practical notion of robustness: multiple structurally distinct models agree on the return distribution and produce consistent portfolio recommendations.

Compared with classical mean–variance or single-model approaches, MAS^2 offers greater adaptability in environments characterized by regime shifts, time-varying volatility, and heavy-tailed risks. Its iterative structure naturally accommodates new information, recalibrates model parameters as markets evolve, and allows the active model set to expand or contract when warranted by the data. This adaptability is essential when standard parametric assumptions fail or become unstable, a situation frequently encountered in real-world portfolio management.

MAS^2 brings together concepts from machine learning (multi-agent systems, synthetic data augmentation), Bayesian statistics (posterior inference, predictive validation, model selection), and financial economics (utility-based portfolio choice) into a cohesive and extensible methodology. Future work will focus on applying the framework to empirical datasets, quantifying its performance relative to existing portfolio optimization techniques, and extending it to settings that incorporate transaction costs, multi-period decision making, and dynamic constraints. Such extensions will help assess the practical viability and scalability of MAS^2 in real-world asset management contexts.

References

[1] Han Du et al. “Comparing DIC and WAIC for multilevel models with missing data”. In: *Behavior Research Methods* 56.4 (2024), pp. 2731–2750.

[2] Lorenzo Garlappi, Raman Uppal, and Tan Wang. “Portfolio selection with parameter and model uncertainty: A multi-prior approach”. In: *The Review of Financial Studies* 20.1 (2007), pp. 41–81.

[3] Abhishek Gunjan and Siddhartha Bhattacharyya. “A brief review of portfolio optimization techniques”. In: *Artificial Intelligence Review* 56.5 (2023), pp. 3847–3886.

[4] Valeriane Jokhadze and Wolfgang M Schmidt. “Measuring model risk in financial risk management and pricing”. In: *International Journal of Theoretical and Applied Finance* 23.02 (2020), p. 2050012.

[5] Harry Markowitz. “Portfolio selection”. In: *The Journal of Finance* 7.1 (1952), pp. 77–91.

[6] Adelinde M Uhrmacher and Danny Weyns. *Multi-Agent systems: Simulation and applications*. CRC press, 2018.