Learning-Augmented Online Portfolio Selection: Optimal Robustness-Consistency Tradeoffs

Ziliang Zhang
School of Computer Science
The University of Sydney
Camperdown NSW 2050, Australia
zzha0461@uni.sydney.edu.au

Tianming Zhao
School of Computer Science
The University of Sydney
Camperdown NSW 2050, Australia
tianmingzhao97@gmail.com

Albert Y. Zomaya
School of Computer Science
The University of Sydney
Camperdown NSW 2050, Australia
albert.zomaya@sydney.edu.au

Abstract

We study learning-augmented online portfolio selection (OPS), where an investor uses predictions to improve wealth while remaining protected against unreliable advice. In frictionless markets, we characterize the optimal tradeoff between robustness (wealth under arbitrary predictions) and consistency (wealth under perfect predictions) tied to the geometric mean, and show that the optimal tradeoff admits a tractable online characterization. We further quantify how performance interpolates between these endpoints under imperfect predictions by smoothness guarantees. Under transaction costs, we give exact robust-trading certificates and the first learning-augmented OPS impossibility result: for any positive cost, no GM-robust policy can be maximally consistent under exact predictions. Lastly, we show that a cost-aware greedy rule remains Pareto-optimal.

1 Introduction

Online portfolio selection (OPS) considers repeated investment decisions over a fixed universe of assets, where an investor reallocates wealth across these assets at each period before observing the upcoming returns. The objective is to maximize the multiplicative wealth over the entire horizon for arbitrary return sequences. Motivated by the increasing availability of predictive signals from Machine Learning (ML) models, recent work has begun to incorporate predictions into portfolio construction, aiming to exploit accurate forecasts while retaining robustness on portfolio wealth to incorrect or adversarial prediction. This learning-augmented viewpoint raises a fundamental question:

How much can predictions improve gain without sacrificing worst-case robustness?

Advice creates opportunity and vulnerability in the same move. Extracting value from accurate advice requires the portfolio to respond to it, yet this same sensitivity can become a liability when the advice is inaccurate. The core challenge is therefore to characterize the greatest prediction-driven improvement that can be achieved while preserving worst-case protection. Such characterization is valuable not only algorithmically, but also conceptually: the optimal frontier it reveals sharpens our understanding of which predictive signals are worth exploiting, how much they can improve wealth under explicit safety constraints and risk tolerances, and where their benefit fundamentally saturates.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

Related work. Classical OPS originates in the universal-portfolio literature, which studies sequential wealth maximization under arbitrary return sequences and without stochastic assumptions on market behavior. The Universal Portfolio of Cover [1] and the multiplicative-update approach of Helmbold et al. [2] compete with the best constant-rebalanced portfolio (BCRP) in hindsight and achieve asymptotically matching log-growth in an assumption-free setting. A complementary line of work seeks stronger empirical performance by exploiting recurring market structure, such as mean reversion [3; 4], pattern matching [5], or meta-learning [6]; see Li and Hoi [7] for a broader taxonomy. In parallel, a growing empirical literature develops predictive signals for portfolio construction, including ranking models [8], factor models [9], time-series forecasting [10], stock-movement prediction [11], and transformer architectures [12]. These developments motivate prediction-aware OPS, in which forecasts are available to the online investor but may be informative, noisy, or even adversarial.

A natural precursor to this perspective is universal portfolios with side information. Cover and Ordentlich [13] incorporate side information into universal portfolios, and Bhatt et al. [14] extend this idea to continuous signals. These works show that external information can guide sequential portfolio updates, but they do not formulate the modern learning-augmented question of explicitly balancing gains under accurate advice against worst-case guarantees when the advice is unreliable. The closest prior work to ours is Zhang et al. [15], which studies ML-assisted OPS with worst-case guarantees in frictionless markets. In contrast, our work is motivated by a broader question: rather than identifying a single robust prediction-aware strategy, we ask what robustness-consistency tradeoffs are fundamentally attainable in OPS, and how this picture changes once costs are introduced.

More broadly, our work is connected to Algorithms with Predictions, which studies algorithms that are robust to arbitrary advice yet benefit from accurate predictions. Representative examples include online caching [16], ski rental and non-clairvoyant scheduling [17], and online bipartite matching [18]; see also Mitzenmacher and Vassilvitskii [19] and the repository [20]. Within this literature, Pareto-optimal robustness-consistency tradeoffs have been studied in multi-option ski rental [21], online knapsack [22], online conversion [23], and one-max search [24], building on the general framework of best-possible tradeoffs in learning-augmented online algorithms [25]. Prediction-aware OPS is technically distinct from most learning-augmented settings because performance is evaluated through multiplicative wealth. As a result, standard competitive-ratio can become vacuous: a persistent per-round gap from the comparator compounds over time and may cause the ratio to decay exponentially with the horizon. Prediction-aware OPS therefore calls for a wealth-based Pareto formulation.

Contribution. We formulate learning-augmented OPS directly in the wealth domain. We provide theoretical analysis on robustness, consistency, and smoothness through instance-wise wealth guarantees, and adopt a strong pathwise Pareto notion tailored to OPS. This formulation reveals a residual-budget principle: a global robustness on the entire horizon can be maintained by tracking a state variable σ_i, which quantifies how aggressively predictions can be exploited at each period.

In frictionless markets, we characterize the optimal robustness-consistency tradeoff tied to the GM[1],

$$S_n^{\mathrm{GM}}(\mathbf{x}^n) := \prod_{i=1}^{n} \prod_{j=1}^{m} x_j(i)^{1/m} \tag{1}$$

We reveal that the global robustness requirement at the geometric-mean level admits a non-clairvoyant one-step characterization through a historical residual state σ: a deterministic pathwise wealth floor is preserved exactly when the next portfolio $\mathbf{b}$ lies in a reverse-KL ball around the uniform portfolio $\mathbf{u} = (1/m, \cdots, 1/m)$, i.e. $D_{\mathrm{KL}}(\mathbf{u}||\mathbf{b}) \leq \log \frac{1}{\sigma}$. This characterization lifts the single-period portfolio geometry to the sequential OPS setting, and provides a safety interface for which any one-step rule can be augmented to preserve sequential robustness. When the residual constraint is active: we solve the optimization, characterize the consistency envelope, and construct a greedy policy $\mathrm{MOP}_{\mathrm{GM}}$ that weakly dominates all GM-robust portfolios in consistency uniformly over all return sequences.

We extend the analysis to transaction costs. We derive exact state-dependent robust feasible region and prove a sharp impossibility: for any positive costs and $n \geq 3$, a maximal policy under exact predictions does not exist. We prove that the greedy template is Pareto-optimal (prediction-efficient). Lastly, we complement our theory with empirical results on NYSE and S&P500 using synthetic and ML forecasts, and show that our proposed algorithms **w/wo** transaction costs are effective in practice.

[1]We provide an analysis and justify the choice of the geometric mean in Preliminary section **Why GM?**

2 Preliminary

Why GM? In this work, we use the geometric-mean market wealth as the robustness benchmark. This choice is deliberate: GM is a pathwise safety floor, not a hindsight-regret comparator. Stronger hindsight benchmarks such as the offline optimum (OPT) and the best constant-rebalanced portfolio (BCRP) require future returns and therefore belong naturally to the consistency side, where predictions are accurate. They are not suitable as exact non-clairvoyant robustness floors.

Stronger causal benchmarks are attainable, but degenerate for our purpose. For example, uCRP is maximin-optimal under the per-round normalization in Proposition A.1. However, exact uCRP-robustness forces the policy to choose the uniform allocation at every history: otherwise an adversarial continuation with all future returns equal to $\mathbf{1}$ exposes the first deviation. Thus predictions cannot affect the decision. The same issue occurs for the uniform buy-and-hold portfolio, $S_n^{\text{uBnH}}(\mathbf{x}^n) = 1/m \cdot \sum_{j=1}^{m} \prod_{i=1}^{n} x_j(i)$, where exact domination forces the policy to follow the drifted buy-and-hold weights. Hence, these causal baselines may be useful non-predictive strategies, but they do not induce a non-degenerate robustness-consistency frontier.

GM occupies the useful middle ground. It captures average market log-growth, is deterministic and distribution-free, and is tight at unit scale: on the all-$\mathbf{1}$ return sequence, every frictionless portfolio and GM both attain wealth 1, so no policy can guarantee a uniformly strict multiplicative improvement over GM. Yet GM is non-degenerate: whenever the AM-GM inequality is strict, the resulting surplus can be converted into prediction-dependent deviations from the uniform portfolio.

Portfolio basics. Let $\mathbf{x}^n = \{\mathbf{x}(i)\}_{i=1}^n \in \mathcal{X}$ be a return sequence with m assets over n periods, where $\mathbf{x}(i) = (x_1(i), \cdots, x_m(i)) \in \mathbb{R}_{>0}^m$ is the gross return vector. Each $x_j(i) > 0$ is the return of asset j at period i, i.e., $x_j(i) = p_j(i)/p_j(i-1)$. A portfolio is a sequence of allocations $\mathbf{b}^n = \{\mathbf{b}(i)\}_{i=1}^n$ with $\mathbf{b}(i) = (b_1(i), \cdots, b_m(i)) \in \Delta^{m-1}$, where $b_j(i) \geq 0, \forall j$ and $\sum_{j=1}^m b_j(i) = 1$. Here, $b_j(i)$ is the fraction of wealth allocated to asset j at period i, and the per-period decision $\mathbf{b}(i)$ must be fixed before the return $\mathbf{x}(i)$ is revealed. We write $S_n(\mathbf{x}^n, \mathbf{b}^n) = \prod_{i=1}^n \sum_{j=1}^m b_j(i)x_j(i)$ for the wealth accumulated on sequence $\mathbf{x}^n$ with portfolio $\mathbf{b}^n$, given an initial wealth of 1. The objective is to determine $\mathbf{b}^n$ to maximize $S_n(\mathbf{x}^n, \mathbf{b}^n)$ for all return sequences in $\mathcal{X}$. For clarity, we first assume frictionless trading and defer the discussion under transaction costs in section 4.

We assume a quantitative prediction model $\mathbf{y}^n \in \mathcal{X}$. At each period i before $\mathbf{x}(i)$ is revealed, the investor observes a prediction $\mathbf{y}(i) \in \mathbb{R}_{>0}^m$. We make no assumptions about the prediction mechanism beyond positivity: $\mathbf{y}(i)$ may arise from heuristics, indicators, or any ML model that outputs next returns. We model per-period prediction by $\eta : \mathbb{R}_{>0}^m \times \mathbb{R}_{>0}^m \to (0, 1]$, where $\eta(\mathbf{x}(i), \mathbf{y}(i)) = 1$ denotes a perfect prediction. As a sufficient (but not necessary) condition, exact recovery of the return vector is a perfect prediction, i.e., $\eta(\mathbf{x}(i), \mathbf{x}(i)) = 1 \; \forall \mathbf{x}(i)$. An online policy A maps observations $(\mathbf{x}^{i-1}, \mathbf{y}^i)$ to portfolios $\mathbf{b}(i) \in \Delta^{m-1}$ and attains wealth $S_n^A(\mathbf{x}^n, \mathbf{y}^n)$ when run on sequence $(\mathbf{x}^n, \mathbf{y}^n)$.

In this work, we study the learning-augmented tradeoff tied to the geometric mean of the sequence, $S_n^{\text{GM}}(\mathbf{x}^n) := \prod_{i=1}^n \prod_{j=1}^m x_j(i)^{1/m}$. The GM captures the average logarithmic growth rate of the market universe: $\log S_n^{\text{GM}}(\mathbf{x}^n) = \frac{1}{m} \sum_{j=1}^m \sum_{i=1}^n \log x_j(i)$. Thus, a guarantee relative to GM ensures that the algorithm remains relative to the average log-growth of the assets. This gives a deterministic and distribution-free notion of broad market growth: it requires neither stochastic assumptions, stationarity, nor calibrated forecasts. These invariances are particularly important in OPS, where market regimes may shift rapidly and unpredictably. The choice is also consistent with classical and recent OPS theory: GM appears as a standard deterministic safety floor in universal-portfolio analysis [1, Prop. 5], and has also been used as a robustness measure in learning-augmented OPS [15].

Definition 2.1 (Robustness)**.** An online algorithm A is GM-robust if:

$$\forall \mathbf{x}^n \in \mathcal{X} : \inf_{\mathbf{y}^n \in \mathcal{X}} S_n^A(\mathbf{x}^n, \mathbf{y}^n) \geq S_n^{GM}(\mathbf{x}^n) = \prod_{i=1}^{n} \prod_{j=1}^{m} x_j(i)^{\frac{1}{m}} \tag{2}$$

Robustness is evaluated against arbitrary predictions. Since we impose no reliability assumptions on the prediction stream, any policy that exploits predictions must remain protected even when the advice is maximally misleading. In this sense, robustness isolates the cost of relying on predictions and ensures that any guarantee is attributable to the policy rather than to favorable advice only.

Definition 2.2 (Consistency)**.** Let $\mathcal{C} : \mathcal{X} \to \mathbb{R}_{>0}$ be some wealth functionals depending only on the return sequence. An online policy A is $\mathcal{C}$-consistent if:

$$\forall \mathbf{x}^n \in \mathcal{X} : \inf_{\mathbf{y}^n \in \mathcal{X}: \forall i, \eta(\mathbf{x}(i), \mathbf{y}(i))=1} S_n^A(\mathbf{x}^n, \mathbf{y}^n) \geq \mathcal{C}(\mathbf{x}^n) \tag{3}$$

Consistency captures the perfect-advice endpoint: it measures how well a policy can exploit ideal predictions. A higher consistency means that the policy converts accurate advice into stronger wealth guarantees. However, consistency alone is not meaningful, since it can be made arbitrarily large by policies that rely aggressively on predictions and fail badly under misleading advice. It therefore must be paired with robustness, which leads naturally to Pareto-optimality under a fixed robustness requirement. For A, define its consistency value under realized input $\mathbf{x}^n$:

$$C_A(\mathbf{x}^n) := \inf_{\mathbf{y}^n \in \mathcal{X}: \forall i, \eta(\mathbf{x}(i), \mathbf{y}(i))=1} S_n^A(\mathbf{x}^n, \mathbf{y}^n)$$

Definition 2.3 (Pareto-optimality)**.** Let $\mathcal{B}_{\text{GM}}$ denote the class of GM-robust online policies satisfying Eq. 2. An online policy $A \in \mathcal{B}_{\text{GM}}$ is pareto-optimal if there does not exist policy $A' \in \mathcal{B}_{\text{GM}}$ such that:

$$\forall \mathbf{x}^n \in \mathcal{X} : C_{A'}(\mathbf{x}^n) \geq C_A(\mathbf{x}^n) \qquad \text{and} \qquad \exists \hat{\mathbf{x}}^n \in \mathcal{X} : C_{A'}(\hat{\mathbf{x}}^n) > C_A(\hat{\mathbf{x}}^n) \tag{4}$$

Eq. 4 formalizes prediction-efficiency under GM constraint: a pareto-optimal policy cannot be uniformly improved in its perfect-advice wealth guarantee. Therefore, it identifies the efficient frontier between robust protection against misleading predictions and upside from accurate predictions.

Definition 2.4 (Maximal optimality)**.** Let $\mathcal{B}_{\text{GM}}$ denote the class of GM-robust online policies by Eq.2. An online policy $A \in \mathcal{B}_{\text{GM}}$ is maximal-optimal if it weakly dominates every policy in $\mathcal{B}_{\text{GM}}$:

$$\forall A' \in \mathcal{B}_{\text{GM}}, \forall \mathbf{x}^n \in \mathcal{X} : C_A(\mathbf{x}^n) \geq C_{A'}(\mathbf{x}^n) \tag{5}$$

Maximal optimality is a sharper, envelope-level guarantee than pareto-optimality. It attains the largest perfect-advice wealth guarantee achievable by any GM-robust policy for every realized return sequence. This rules out all sequence-dependent improvements within the GM-robust class and establishes the policy as a canonical optimal way to exploit predictions without sacrificing robustness.

In the next section, we show that such a maximal policy exists in frictionless OPS, and its construction follows from an inductive reduction. We then show that transaction costs fundamentally alter this picture: once rebalancing frictions are present, no online policy can maintain maximal optimality. This motivates the notation of Pareto-optimality which we study in the transaction-cost setting.

3 Optimal robustness-consistency tradeoff

We begin by characterizing the set $\mathcal{B}_{GM}$ of portfolio policies, the ones that are robust with respect to the geometric mean (GM-robust). Our starting point is a one-period analysis; we ask for which weight vectors $\mathbf{b} \in \Delta^{m-1}$ the following invariant holds: the one-step wealth $\sum_{j=1}^m b_j x_j$ is uniformly lower bounded by a scaled geometric mean of the returns for all $\mathbf{x} \in \mathbb{R}^m_{>0}$. This local one-step characterization serves as the building block for the robust set $\mathcal{B}_{GM}$, and will later allow us to derive a complete pathwise description of GM-robust portfolios.

Lemma 3.1 (A.2)**.** *Let $m \geq 2$, $\mathbf{u} = (\frac{1}{m}, \cdots, \frac{1}{m})$, $\mathbf{b} \in \Delta^{m-1}$, and fix $\sigma \in (0, 1]$. Then:*

$$\forall \mathbf{x} \in \mathbb{R}^m_{>0} : \sum_{j=1}^m b_j x_j \geq \sigma \prod_{j=1}^m x_j^{\frac{1}{m}} \iff \prod_{j=1}^m b_j \geq \left(\frac{\sigma}{m}\right)^m \iff D_{KL}(\mathbf{u}||\mathbf{b}) \leq \log \frac{1}{\sigma}$$

Lemma 3.1 shows that the ability to uniformly dominate a scaled geometric mean in a single period is fully captured by the geometric mean of the portfolio weights. Consequently, the feasible set:

$$\mathcal{B}(\sigma) := \left\{ \mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^m b_j \geq \left(\frac{\sigma}{m}\right)^m \right\}, \quad \sigma \in (0, 1]$$

forms the one-step admissible portfolios that are monotone in σ: as σ decreases, the constraint becomes weaker and $\mathcal{B}(\sigma)$ expands, while larger σ shrink the feasible region. The smallest residual at $\sigma = 1$ yields an admissible singleton by the application of AM-GM below.

Corollary 3.2 (A.3). *Let $m \geq 2$, $\mathbf{x} \in \mathbb{R}^m_{>0}$, and $\sum_{j=1}^m b_j = 1$ with $b_j \geq 0$. Then:*

$$\forall \mathbf{x} \in \mathbb{R}^m_{>0} : \sum_{j=1}^m b_j x_j \geq \prod_{j=1}^m x_j^{\frac{1}{m}} \iff \mathbf{b} = \left(\frac{1}{m}, \cdots, \frac{1}{m}\right)$$

The corollary identifies the boundary case: when $\sigma = 1$, GM robustness forces the uniform portfolio. Once the strategy outperforms GM, the residual drops below one, and the feasible region in Lemma 3.1 expands to allow non-uniform actions. We now characterize the entire class of GM-robust portfolios.

Lemma 3.3 (A.4). *Fix $m \geq 2$. Consider any portfolio strategy A that, at each round i, maps the past history $\mathbf{x}^{i-1}$ to a portfolio $\mathbf{b}(i) = A(\mathbf{x}^{i-1}) \in \Delta^{m-1}$. For any sequence $\mathbf{x}^n = \{\mathbf{x}(i)\}_{i=1}^n$ with $\mathbf{x}(i) \in \mathbb{R}^m_{>0}$, define $S_n^A(\mathbf{x}^n) := \prod_{i=1}^n \sum_{j=1}^m b_j(i) x_j(i)$ and $S_n^{GM}(\mathbf{x}^n) := \prod_{i=1}^n \prod_{j=1}^m x_j(i)^{\frac{1}{m}}$, with $S_0^A(\cdot) = S_0^{GM}(\cdot) = 1$. Then, the equivalence holds $\forall n \geq 1$:*

$$\forall \mathbf{x}^n : S_n^A(\mathbf{x}^n) \geq S_n^{GM}(\mathbf{x}^n) \iff \forall i \in \{1, \cdots, n\}, \forall \mathbf{x}^{i-1} : \prod_{j=1}^m b_j(i) \geq \left(\frac{S_{i-1}^{GM}(\mathbf{x}^{i-1})}{m S_{i-1}^A(\mathbf{x}^{i-1})}\right)^m$$

The condition at $n = 1$ reduces to $\prod_{j=1}^m b_j(1) \geq \frac{1}{m^m}$ which forces $\mathbf{b}(1) = (\frac{1}{m}, \cdots, \frac{1}{m})$.

By Lemma 3.3, as long as at each period i the geometric mean of A's portfolio weights stays above σ_i/m with $\sigma_i = S_{i-1}^{GM}/S_{i-1}^A$, the next-period wealth is guaranteed to dominate the GM benchmark on every return path. This guarantee is *non-clairvoyant*, since the residual factor σ_i is fully determined by the realized history and can be computed before observing the period-i returns. In addition, the condition is tight: it characterizes the largest set of portfolios that are robust to the GM benchmark, where any violation can be exploited by an adversarial choice of returns. This characterization has a broader implication than the construction of a particular algorithm. It turns global robustness into a local constraint. At the multi-period level, all decisions from the past enter the next safety requirement only through the scalar residual σ_i, encapsulated by $\mathcal{B}(\sigma_i)$. Thus Lemma 3.3 separates two questions that are usually entangled in OPS: which actions are safe, and which safe action should be preferred.

Classical portfolio rules, such as mean-variance optimization [26], choose a portfolio by optimizing a one-step statistical model of the next return. Such methods can be highly effective when their inputs are reliable, but their guarantees are inherently model-dependent: errors in predicted means, variances, or covariances can lead to large changes in weights [27; 28; 29]. Bayesian variants such as the Black-Litterman model [30], can stabilize these inputs but still do not provide a deterministic lower bound on realized wealth when the prior is wrong. Lemma 3.3 supplies exactly this missing piece: it provides a safety certificate that is independent of how the one-step objective is constructed.

More generally, let $F_i(\mathbf{b})$ be any causal one-step portfolio rule formed from the information available before period i. The certified version of the rule becomes $\mathbf{b}_F \in \arg\max_{\mathbf{b} \in \mathcal{B}(\sigma_i)} F_i(\mathbf{b})$. Whenever the optimizer exists, the enhanced rule is GM-robust by construction, regardless of the quality of the model used to define F_i. For instance, a mean-variance investor with predicted mean $\hat{\mu}_i$, covariance estimates $\mathcal{K}_i$, and risk-aversion parameter $\rho \geq 0$ may apply the constraint as an adaptive safety layer:

$$\max_{\mathbf{b} \in \mathcal{B}(\sigma_i)} \left\{ \hat{\mu}_i^T \mathbf{b} - \frac{\rho}{2} \mathbf{b}^T \mathcal{K}_i \mathbf{b} \right\}$$

We now use this safety interface to answer a sharper learning-augmented question: among all GM-robust portfolios, which local preference yields the largest consistency under accurate predictions? Since membership in $\mathcal{B}(\sigma_i)$ already certifies robustness, the remaining objective can focus entirely on predictive upside. The Maximal-optimal portfolio at GM, MOP_{GM}, is the simple greedy rule

$$S_i^{\text{MOP}}(\mathbf{x}^i, \mathbf{y}^i) = S_{i-1}^{\text{MOP}}(\mathbf{x}^{i-1}, \mathbf{y}^{i-1}) \times \hat{\mathbf{b}}(i) \cdot \mathbf{x}(i), \quad \hat{\mathbf{b}}(i) \in \arg\max_{\mathbf{b} \in \mathcal{B}(\sigma_i)} \mathbf{b} \cdot \mathbf{y}(i), \quad \sigma_i = \frac{S_{i-1}^{\text{GM}}}{S_{i-1}^{\text{MOP}}} \quad (6)$$

The consistency intuition follows naturally. At period i, a clairvoyant GM-safe action would maximize $\mathbf{b} \cdot \mathbf{x}(i)$ over $\mathcal{B}(\sigma_i)$, whereas MOP_{GM} maximizes $\mathbf{b} \cdot \mathbf{y}(i)$ over the safe set $\mathcal{B}(\sigma_i)$. These two decisions coincide whenever the prediction is oracle-consistent for the same set; in particular, this holds if $\mathbf{y}(i) = \alpha_i \mathbf{x}(i) + \beta_i \mathbf{1}$ $\forall \alpha_i > 0$ and $\beta_i \in \mathbb{R}$ by A.6. The condition is much weaker than exact return prediction. The reason this greedy step is globally optimal is that immediate gain and future flexibility are aligned. For a fixed history and realized return, the optimization simultaneously maximizes the

Algorithm 1 MOP_{GM}

Require: number of assets m, horizon n

initialize MOP_{GM} $S(0) \leftarrow 1$, GM $S^{GM}(0) \leftarrow 1$
for $i \leftarrow 1$ **to** n **do**
  Compute residual $\sigma \leftarrow S^{GM}(i-1) \div S(i-1)$
  Get $\mathcal{B}(\sigma) := \{\mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^{m} b_j \geq (\frac{\sigma}{m})^m\}$
  Receive prediction $\mathbf{y}(i) \in \mathbb{R}_+^m$
  Optimize $\mathbf{b}^* \leftarrow \arg\max_{\mathbf{b} \in \mathcal{B}(\sigma)} \sum_{j=1}^{m} b_j y_j(i)$
  Receive true return $\mathbf{x}(i) \in \mathbb{R}_+^m$
  Update MOP $S(i) \leftarrow S(i-1) \sum_{j=1}^{m} b_j^* x_j(i)$
  Update GM $S^{GM}(i) \leftarrow S^{GM}(i-1) \prod_{j=1}^{m} x_j(i)^{1/m}$
end for
return $S(n)$

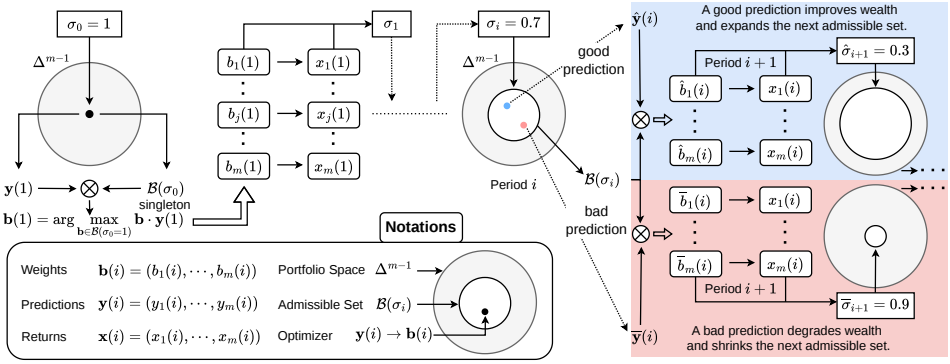

Figure 1: Visualization of MOP_{GM} algorithm.

current gain and minimizes the next residual. Since $\mathcal{B}(\sigma)$ expands as σ decreases, the best GM-safe action today also creates the largest safe action set tomorrow. Therefore, under consistent predictions, MOP_{GM} achieves the maximal consistency among all GM-robust policies. Algorithm 1 implements the residual feedback loop, and Figure 1 visualizes its dynamics. The next theorem formalizes the pathwise guarantees on robustness and consistency over all return sequences.

Theorem 3.4 (A.7). *Let S_n^{MOP} and S_n^{GM} denote the wealth of Algorithm 1 and the geometric mean with $S_0^{MOP} = S_0^{GM} = 1$, $\sigma_n := S_{n-1}^{GM}/S_{n-1}^{MOP}$, and $\mathcal{B}(\sigma_n) := \{\mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^{m} b_j \geq (\sigma_n/m)^m\}$.*

1. ***Robustness***. *$\forall \mathbf{x}^n, \mathbf{y}^n \in (\mathbb{R}_{>0}^m)^n$, $\forall m \geq 2$, $\forall n \geq 1$: $S_n^{MOP}(\mathbf{x}^n, \mathbf{y}^n) \geq S_n^{GM}(\mathbf{x}^n)$*

2. ***Consistency***[2] *Under consistent predictions $\overset{*}{\mathbf{y}}{}^n = \{\overset{*}{\mathbf{y}}(i)\}_{i=1}^n$, $\overset{*}{\mathbf{y}}(i) = \alpha_i \mathbf{x}(i) + \beta_i \mathbf{1}$, $\alpha_i > 0$, $\beta_i \in \mathbb{R}$. Let $\sigma_{1:n}^* = \{\sigma_1^*, \cdots \sigma_n^*\}$ be the realized residuals from $S_n^{MOP}(\mathbf{x}^n, \overset{*}{\mathbf{y}}{}^n)$. Then $\forall \mathbf{x}^n$:*

$$S_n^{MOP}(\mathbf{x}^n, \overset{*}{\mathbf{y}}{}^n) = S_n^*(\mathbf{x}^n | \sigma_{1:n}^*) \geq \frac{1}{m} \left[S_n^{OPT} + (m-1) \left(\frac{(S_n^{GM})^m}{S_n^{OPT}} \right)^{\frac{1}{m-1}} \right]$$

Remark 3.5. MOP_{GM} is maximal-optimal: it weakly dominates all GM-robust policies on consistency.

Theorem 3.4 makes the tradeoff transparent. If the investor demands a robustness guarantee against broad market growth, as captured by the geometric-mean, then the attainable predictive upside is still remarkably strong: under consistent predictions, MOP_{GM} achieves wealth at least a constant factor of the offline optimum. The tradeoff is also tight: it identifies the fundamental limit at which predictive upside saturates under GM robustness. Moreover, the consistency condition $\mathbf{y}(i) = \alpha_i \mathbf{x}(i) + \beta_i \mathbf{1}$ clarifies what "perfect prediction" means. What matters is not exact recovery of the next return

[2]We state only the generic case $\sigma_i \in (0, 1)$ and $x_1(i), \cdots, x_m(i)$ not all equal. The appendix gives complete result for all $\sigma_i \in [0, 1]$ and all $\mathbf{x}(i) \in \mathbb{R}_{>0}^m$; the lower bound remains unchanged. The exact consistency value requires solving a series of nonlinear optimizations, the closed-form expression is presented in the appendix. These are omitted for readability only.

vector, but preserving the cross-sectional structure that determines the best GM-robust direction. Consequently, predictions remain invariant to any common multiplicative or additive bias across assets. This perspective also suggests a different view of prediction quality. Standard forecasting losses reward point-wise accuracy of $\mathbf{y}(i)$ as an estimate of $\mathbf{x}(i)$, but the theorem shows that the portfolio decision only depends on the equivalent class of predictions up to positive scaling and common shifts. This points toward portfolio-aware training objectives that emphasize the relative cross-sectional information exploited by the optimizer, rather than absolute calibration alone.

Smoothness quantifies how prediction errors translate into loss relative to the consistency frontier. This can be expressed globally by comparing the realized wealth of MOP_{GM} with the consistency envelope $S_n^*(\mathbf{x}^n, \sigma_{1:n}^*)$. Let $\hat{\mathbf{b}}(i) \in \arg\max_{\mathcal{B}(\sigma_i)} \mathbf{b} \cdot \mathbf{y}(i)$ be the portfolio chosen by MOP_{GM} on realized σ_i, the error-dependent degradation can be presented by defining per-round error factor η_i:

$$S_n^{\text{MOP}}(\mathbf{x}^n, \mathbf{y}^n) = S_n^*(\mathbf{x}^n | \sigma_{1:n}^*) \prod_{i=1}^{n} \eta_i, \qquad \eta_i := \frac{\hat{\mathbf{b}}(i) \cdot \mathbf{x}(i)}{\max_{\mathbf{b} \in \mathcal{B}(\sigma_i^*)} \mathbf{b} \cdot \mathbf{x}(i)} \in (0, 1]$$

This provides a tractable global characterization on smoothness: η_i measures the full per-round loss relative to the envelope. It absorbs both sources of degradation simultaneously: i) the current-round mismatch between optimizing $\mathbf{x}(i)$ and $\mathbf{y}(i)$, and ii) the historical effect of earlier errors which alter the residual path. Thus, prediction error is linked directly to wealth through an efficiency factor.

To further interpret this global factor, we next localize the comparison by freezing the residual-induced safe set. Under an identical residual constraint, we upper bound the one-step wealth loss incurred when the prediction is only approximately aligned with the consistent class $\alpha_i \mathbf{x}(i) + \beta_i \mathbf{1}$.

Lemma 3.6. *[A.8] Fix $\sigma \in (0, 1]$, and denote set $\mathcal{B} = \mathcal{B}(\sigma) := \{\Delta^{m-1} : \prod_{j=1}^{m} b_j \geq (\frac{\sigma}{m})^m\}$. Let $\mathbf{b}_y \in \arg\max_{\mathbf{b} \in \mathcal{B}} \mathbf{b} \cdot \mathbf{y}$, and $\mathbf{b}_x \in \arg\max_{\mathbf{b} \in \mathcal{B}} \mathbf{b} \cdot \mathbf{x}$. Then, for $\epsilon \geq 0$:*

$$||\mathbf{y} - (\alpha \mathbf{x} + \beta \mathbf{1})||_\infty \leq \epsilon \implies \mathbf{b}_y \cdot \mathbf{x} \geq \mathbf{b}_x \cdot \mathbf{x} - \frac{2\epsilon}{\alpha}$$

The lemma shows that an ℓ_∞ deviation from this class induce an additive loss. An affine-normalized error ϵ/α in predictions causes at most $2\epsilon/\alpha$ loss in wealth compared to the clairvoyant optimizer on the same safe set. Combined with the global decomposition, this shows that prediction errors affect MOP_{GM} through a controlled local decision loss and the residual dynamic that propagates it through.

These results complete the frictionless GM picture. Lemma 3.3 identifies the exact safe certificate: the entire history needed for GM robustness is compressed into the scalar residual σ_i. The MOP_{GM} rule then uses this certificate to convert predictions into maximum upside, and Theorem 3.4 proves the tightness. The next section asks whether this reduction survives under transaction costs.

4 Portfolio with transaction costs

At the start of period i, let $\mathbf{b}(i-1)$ denote the portfolio held through period $i-1$. After returns $\mathbf{x}(i-1)$ are realized, the pre-trade (the post-return at $i-1$) weights before i are drifted by $\mathbf{x}(i-1)$:

$$\mathbf{a}(i) := \frac{\mathbf{b}(i-1) \odot \mathbf{x}(i-1)}{\mathbf{b}(i-1) \cdot \mathbf{x}(i-1)} \in \Delta^{m-1} \tag{7}$$

We then rebalance to the target $\mathbf{b}(i)$ from $\mathbf{a}(i)$. The per-asset turnover is the sell fraction $\max\{a_j - b_j, 0\}$ or buy fraction $\max\{b_j - a_j, 0\}$ respectively, and the full turnover is composed by Lemma 4.1.

Lemma 4.1 (A.9)**.** *Let $m \geq 2$, $\mathbf{a} = (a_1, \cdots, a_m) \in \Delta^{m-1}$, $\mathbf{b} = (b_1, \cdots, b_m) \in \Delta^{m-1}$. Then:*

$$\sum_{j=1}^{m} \max\{a_j - b_j, 0\} = \sum_{j=1}^{m} \max\{b_j - a_j, 0\} = \frac{1}{2}||\mathbf{a} - \mathbf{b}||_1, \quad 0 \leq ||\mathbf{a} - \mathbf{b}||_1 \leq 2$$

We assume proportional transaction costs on both sells and buys, with rates $\kappa_s + \kappa_b \leq 0.5$. Rebalancing from $\mathbf{a}(i)$ to the target portfolio $\mathbf{b}(i)$ yields the total proportional trading cost $\frac{\kappa_s + \kappa_b}{2}||\mathbf{a} - \mathbf{b}||_1$, which induces the wealth shrinkage factor Φ and the one-period net growth factor below:

$$\Phi(\mathbf{a}(i) \to \mathbf{b}(i)) = 1 - \frac{\kappa_s + \kappa_b}{2}||\mathbf{a}(i) - \mathbf{b}(i)||_1, \quad S(\mathbf{b}(i), \mathbf{x}(i)) = \Phi(\mathbf{a}(i) \to \mathbf{b}(i)) \sum_{j=1}^{m} b_j(i) x_j(i)$$

This model[3] captures several features important for OPS analysis: it charges zero cost to buy-and-hold, penalizes only active rebalancing, and allows asymmetric buy and sell rates. We assume the initial investment is frictionless: investors choose $\mathbf{b}(1)$ without paying any costs, and costs are incurred only from the first rebalancing onward. This is a normalization convention since any initial cost contributes a sequence-independent multiplicative factor and can be absorbed into the initial wealth.

Fix a per-round robustness level $\sigma \in (0, 1]$ tied to the geometric mean of the upcoming return. Under transaction cost factor Φ, the one-period dominance condition admits an exact characterization form:

$$\forall \mathbf{x} \in \mathbb{R}^m_{>0} : \Phi(\mathbf{a} \to \mathbf{b}) \sum_{j=1}^{m} b_j x_j \geq \sigma \prod_{j=1}^{m} x_j^{1/m} \iff \prod_{j=1}^{m} b_j^{\frac{1}{m}} \geq \frac{\sigma}{m\Phi(\mathbf{a} \to \mathbf{b})}$$

by Lemma A.11 (and equivalently one may include the redundant constraint $\Phi(\mathbf{a} \to \mathbf{b}) \geq \sigma$). In the sequential setting, Lemma A.13 shows that transaction-cost-aware GM-robustness is equivalent to enforcing the corresponding one-step portfolio weight constraint at every period i and history $\mathbf{x}^{i-1}$:

$$\forall \mathbf{x}^n : \prod_{i=1}^{n} \Phi_i \sum_{j=1}^{m} b_j(i) x_j(i) \geq S_n^{GM}(\mathbf{x}^n) \iff \forall i, \forall \mathbf{x}^{i-1} : \prod_{j=1}^{m} b_j(i)^{\frac{1}{m}} \geq \frac{S_{i-1}^{GM}(\mathbf{x}^{i-1})}{m\Phi_i S_{i-1}(\mathbf{x}^{i-1})}$$

We therefore define the feasible set of transaction-cost-aware GM-robust portfolios at round i as:

$$\mathcal{B}^{\Phi}_{\text{GM}}(\sigma_i, \mathbf{a}) := \left\{ \mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^{m} b_j^{\frac{1}{m}} \geq \frac{\sigma_i}{m\Phi(\mathbf{a} \to \mathbf{b})} \right\}$$

We highlight some implications of this characterization. (i) The constraint depends on observable state, thus non-clairvoyant. (ii) The set $\mathcal{B}^{\Phi}_{\text{GM}}$ is tight: if a strategy selects $\mathbf{b}(i) \notin \mathcal{B}^{\Phi}_{\text{GM}}(\sigma_i, \mathbf{a}(i))$ for some history, then $\exists \mathbf{x}(i)$ that violates the one-step inequality, implying failure of GM-robustness. (iii) The residual $\sigma_i = S_{i-1}^{GM}/S_{i-1}$ measures the strategy's accumulated "buffer" relative to GM. Smaller σ_i relaxes the current GM-robustness constraint. However, feasibility depends not only on σ_i but also on the pre-trade portfolio $\mathbf{a}(i)$, and the current choice $\mathbf{b}(i)$ moves the state via $\mathbf{a}(i+1) \propto \mathbf{b}(i) \odot \mathbf{x}(i)$. This action-dependent transition affects both future rebalancing costs $\Phi(\mathbf{a}(i+1) \to \mathbf{b}(i+1))$ and the next feasible region $\mathcal{B}^{\Phi}(\sigma_{i+1}, \mathbf{a}(i+1))$. Consequently, even under perfect predictions, a greedy policy that maximizes the one-step objective (e.g., $\max_{\mathbf{b} \in \mathcal{B}^{\Phi}} \Phi(\mathbf{a}(i) \to \mathbf{b}(i)) \sum_{j=1}^{m} b_j(i) x_j(i)$) need not maximize terminal wealth: an aggressive step can yield a large immediate gain yet steer $\mathbf{a}(i+1)$ to a region where future rebalancing becomes expensive or severely constrained, offsetting the short-term advantage. We illustrate this non-myopic effect with an intuitive example below.

Consider $m = 2$ assets with transaction cost rates $\kappa_s = \kappa_b = 0.2$ and the return sequence $\{(\frac{1}{2}, 1), (\frac{3}{2}, 1), (\frac{1}{2}, 1)\}$. Predictions are perfect and revealed sequentially. We compare the greedy GM-robust policy POP^{Φ}_{GM} (Appendix Algorithm 2) to another GM-robust policy uBnH (uniform buy-and-hold, i.e., no rebalancing after initialization, hence no future transaction costs). Both strategies start from the forced uniform allocation, yielding $S_1 = 0.75$. At round 2, POP^{Φ}_{GM} achieves $S_2^{\Phi} \approx 0.87604$, slightly exceeding uBnH's $S_2^u = 0.875$; moreover, this is the maximal two-period wealth among GM-robust strategies for this realized prefix, by construction of the greedy optimization. However, at $n = 3$ the ordering reverses: POP^{Φ}_{GM} ends at $S_3^{\Phi} \approx 0.65411$ with optimizer $\mathbf{b}^{\Phi}(3) = (0.49, 0.51)$, while uBnH attains $S_3^u = 0.6875$ with weights $\mathbf{b}^u(3) = (0.428, 0.572)$.

The example reflects a deeper impossibility rather than a limitation of the greedy implementation. Transaction costs introduce state coupling: the robust feasible set depends jointly on the residual budget and the current pre-trade portfolio, so a decision affects not only immediate wealth but also the geometry and cost of future feasibility. In fact, under positive costs, splitting a rebalance across two periods can yield strictly larger effective exposure than making the same move in one step, while preserving the GM robustness, as shown by Lemma A.15. This creates a causal conflict: two sequences with the same prefix can require incompatible period-2 actions for maximal consistency.

[3]It should be viewed as a first-order approximation to the exact self-financing model, in which the post-cost wealth fraction ρ needs to be solved implicitly by $\rho = 1 - \kappa_s \sum_{j=1}^{m} (a_j(i) - \rho b_j(i))_+ - \kappa_b \sum_{j=1}^{m} (\rho b_j(i) - a_j(i))_+$; replacing ρ by 1 inside the turnover terms yield the closed-form factor Φ. We adopt Φ to retain explicit expressions and analytical transparency. The analysis on why Φ is the first-order approximation on the exact financing model is shown in Proposition A.10, which reveals the affect of the approximation on the quality of the outcomes.

Theorem 4.2 (Impossibility result A.16). *Fix $m \geq 2$, $n \geq 3$, and any $\kappa := \kappa_s + \kappa_b \in (0, \frac{1}{2}]$. Let $\mathcal{B}^{\Phi}_{GM}(m,n)$ be the class of transaction-cost GM-robust online policies, and denote the realized consistency value $C^{\Phi}_A(\mathbf{x}^n, \mathbf{x}^n)$ under exact predictions. No policy $A \in \mathcal{B}^{\Phi}_{GM}(m,n)$ exists such that:*

$$\forall B \in \mathcal{B}^{\Phi}_{GM}(m,n), \forall \mathbf{x}^n \in (\mathbb{R}^m_{>0})^n : \quad C^{\Phi}_A(\mathbf{x}^n, \mathbf{x}^n) \geq C^{\Phi}_B(\mathbf{x}^n, \mathbf{x}^n)$$

Equivalently, maximal-optimality does not exist under any positive proportional transaction cost.

Although transaction costs preclude maximal-optimality, Algorithm 2 remains Pareto-optimal under exact-prediction consistency. Any robust policy that attempts to improve its terminal consistency must deviate from the greedy post-cost optimizer at some first history. i) If the deviation sacrifices immediate wealth, an all-one suffix exposes the loss. ii) If the deviation is an immediate tie, the two policies enter distinct states; a separating return vector (A.17) then exposes a continuation on which the greedy state is strictly better. We prove the Pareto-optimality of $\text{POP}^{\Phi}_{\text{GM}}$ in the below proposition.

Proposition 4.3 (A.18). *Fix $m \geq 2$, $n \geq 1$. Under exact predictions $\mathbf{y}(i) = \mathbf{x}(i)\ \forall i \in [n]$, Algorithm 2 POP^{Φ}_{GM} is Pareto-optimal. There does not exist a policy $A \in \mathcal{B}^{\Phi}_{GM}(m,n)$ such that:*

$$\forall \mathbf{x}^n \in (\mathbb{R}^m_{>0})^n : C^{\Phi}_A(\mathbf{x}^n) \geq C^{\Phi}_{POP}(\mathbf{x}^n) \qquad \textit{and} \qquad \exists \hat{\mathbf{x}}^n \in (\mathbb{R}^m_{>0})^n : C^{\Phi}_A(\hat{\mathbf{x}}^n) > C^{\Phi}_{POP}(\hat{\mathbf{x}}^n)$$

The proposition reveals that no robust policy can improve $\text{POP}^{\Phi}_{\text{GM}}$'s consistency uniformly over all paths. Hence, $\text{POP}^{\Phi}_{\text{GM}}$ lies on the efficient frontier of the robustness-consistency tradeoff: any policy that improves its consistency on one sequence must sacrifice consistency on another sequence. In this sense, the policy is prediction-efficient: it extracts as much one-step post-cost value from accurate forecasts as possible without admitting a uniform improvement. The result also clarifies the role of our theory: Pareto-optimality is universal over all positive return sequences, and this universality necessarily includes highly adversarial or extreme inputs. We therefore complement the theory with empirical evaluation, asking how efficiently $\text{POP}^{\Phi}_{\text{GM}}$ transfers accurate forecasts into gain on real data.

5 Numerical results

Datasets and ablations. We adopt the NYSE(O) [31] and NYSE(N) [32] datasets (spanning decades) for controlled studies under synthetic prediction streams, which are canonical OPS benchmarks used in both foundational and recent studies [1; 2; 15]. In addition to the standard stock subsets, we include further subset protocols to probe scalability with respect to the number of assets m, as reported in Table 1. We further use the S&P500 [33] dataset to evaluate performance under real-world prediction inputs. We focus on shorter horizons spanning months rather than decades, which better reflects the operational regime of learned forecasts and short-horizon portfolio deployment. These windows cover distinct market conditions, including a highly volatile COVID-period regime and a more recent regime for contemporary relevance, as summarized in Tables 2 and 3. The design is consistent with [15], ensuring transparent and reproducible comparisons to existing approaches. We ablate the universal family UPSI [13] and EGSI [2], and the recent RAM [15] which claims the same consistency lower bound as ours. To isolate the effect of the portfolio decision rule, all methods receive the same prediction stream. UPSI and EGSI can only accept discrete side information, so we encode the predicted top-return asset as the state; under perfect predictions, this permits the best side-information CRP to match the hindsight optimum. RAM consumes ranked advice only, so predicted magnitudes are converted into ranks. Costs are applied *ex post* to all cost-agnostic policies: we evaluate the realized turnover penalty without altering each algorithm's internal update rule.

Frontier realization under synthetic forecasts. Table 4 reports the optimal frontier for long-horizon NYSE experiments without costs. MOP_{GM} dominates all methods in the consistency upside. Although RAM claims the same consistency lower bound as ours[4], its exact realization deteriorates much faster with the size of the asset universe. Under i.i.d. random predictions (Table 5), MOP_{GM} remains well above its GM guarantee and achieves performance comparable to the competing methods. Table 6 evaluates the frontier under symmetric costs. For mild cost $\kappa = 0.3\%$, $\text{POP}^{\Phi}_{\text{GM}}$ retains the exponential upside and strictly dominates universal portfolio variants. When the cost increases to $\kappa = 1\%$, cost-agnostic policies become fragile. In Combination 8, all cost-agnostic policies underperform GM catastrophically, while $\text{POP}^{\Phi}_{\text{GM}}$'s consistency still reaches 2.92×10^4.

[4]The coarse corollary of $S_n \geq \frac{1}{m} S^{OPT}_n$ is shared by MOP and RAM [15], but MOP's consistency (weakly) dominates RAM's in theory. This section presents the instance-dependent gap beyond the shared $\frac{1}{m}$-OPT bound.

Real-world forecast results. The synthetic experiments isolate the frontier under ideal predictions. We next test whether the same qualitative behavior persists with learned, noisy forecasts. Table 7 and Figure 2 reports NHITS[34] generated forecasts on S&P500. In the frictionless case, $\mathrm{MOP}_{\mathrm{GM}}$ remains above GM over all six regimes and exploits useful forecast signal when present. In particular, in B-α where ML substantially outperforms uBnH, $\mathrm{MOP}_{\mathrm{GM}}$ attains 2.7041, improving over all online policies. Conversely, when the ML forecasts are misleading as in A-β and drop to 0.7071, $\mathrm{MOP}_{\mathrm{GM}}$ remains close to conservative universal portfolios and above GM. Under transaction costs: cost-agnostic policies can induce excessive turnover that destroys accurate signals. Already at $\kappa = 0.3\%$, $\mathrm{MOP}_{\mathrm{GM}}$ falls below GM in all six regimes; when $\kappa = 3\%$, all cost-agnostic policies lag behind GM, even in the B-α regime under strong forecast signals. $\mathrm{POP}^{\Phi}_{\mathrm{GM}}$ addresses this failure mode precisely. For all $\kappa \geq 0.3\%$, it improves over $\mathrm{MOP}_{\mathrm{GM}}$ in every regime. At $\kappa = 3\%$, it is the only policy that remains above GM; particularly in B-α, it still exploits favorable forecasts and reaches 3.1607. These results empirically prove $\mathrm{POP}^{\Phi}_{\mathrm{GM}}$ preserves robustness while exploiting accurate signals effectively.

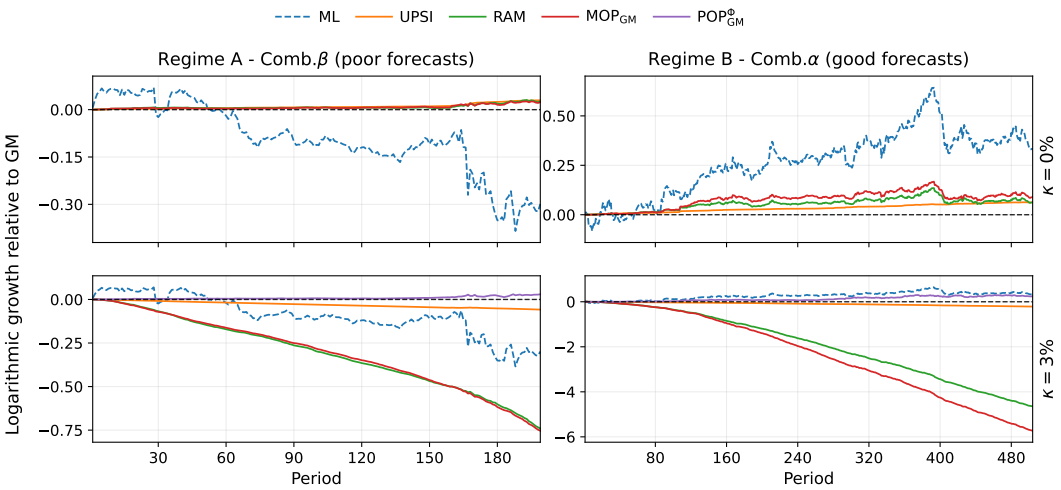

Figure 2: Log-growth ablation normalized by GM for representative trajectories from Table 7. ML is the rank-1 multiplicative log-growth signal without costs, serving as a forecast-quality proxy. Rows indicate cost settings; columns indicate market instances. Detailed settings are reported in Table 7.

6 Conclusion

In this work, we develop a pathwise theory of learning-augmented OPS in the wealth domain. We prove that $\mathrm{MOP}_{\mathrm{GM}}$ is maximal-optimal without transaction costs, and $\mathrm{POP}^{\Phi}_{\mathrm{GM}}$ is Pareto-optimal under transaction costs. Our results are universal over all return sequences; an important direction is therefore to obtain sharper frontiers under structured market models. Another direction for future work is to extend the optimal robustness-consistency frontier beyond the geometric-mean benchmark, potentially through log-sum-exp-based interpolation to achieve stronger robustness guarantees. More broadly, extending residual-style robust certificates to other learning-augmented online problems may provide a general route to prediction-efficient algorithms with explicit safety guarantees.

References

[1] T. M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991. doi: 10.1111/j.1467-9965.1991.tb00002.x.

[2] D. P. Helmbold, R. E. Schapire, Y. Singer, and M. K. Warmuth. On-line portfolio selection using multiplicative updates. *Mathematical Finance*, 8(4):325–347, 1998. doi: 10.1111/1467-9965.00058.

[3] A. Borodin, R. El-Yaniv, and V. Gogan. Can we learn to beat the best stock. *Journal of Artificial Intelligence Research*, 21:579–594, 2004. doi: 10.1613/jair.1336.

[4] B. Li, P. Zhao, S. C. H. Hoi, and V. Gopalkrishnan. Pamr: Passive aggressive mean reversion strategy for portfolio selection. *Machine Learning*, 87(2):221–258, 2012. doi: 10.1007/s10994-012-5281-z.

[5] Bin Li, Steven C. H. Hoi, and Vivekanand Gopalkrishnan. Corn: Correlation-driven nonparametric learning approach for portfolio selection. *ACM Transactions on Intelligent Systems and Technology*, 2(3):21:1–21:29, 2011. doi: 10.1145/1961189.1961193.

[6] Puja Das and Arindam Banerjee. Meta optimization and its application to portfolio selection. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'11)*, pages 1163–1171, San Diego, CA, USA, 2011. Association for Computing Machinery. ISBN 9781450308137. doi: 10.1145/2020408.2020588.

[7] B. Li and S. C. H. Hoi. Online portfolio selection. *ACM Computing Surveys*, 46(3):1–36, 2014. doi: 10.1145/2512962.

[8] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, pages 3146–3154, 2017.

[9] Lakshay Chauhan, John Alberg, and Zachary C. Lipton. Uncertainty-aware lookahead factor models for quantitative investing. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1489–1499. PMLR, 2020.

[10] Gerald Woo, Chenghao Liu, Doyen Sahoo, Akshat Kumar, and Steven C. H. Hoi. Learning deep time-index models for time series forecasting. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 37217–37237. PMLR, 2023.

[11] Shuqi Li, Yuebo Sun, Yuxin Lin, Xin Gao, Shuo Shang, and Rui Yan. Causalstock: Deep end-to-end causal discovery for news-driven multi-stock movement prediction. In *Advances in Neural Information Processing Systems 37*, pages 47432–47454. Neural Information Processing Systems Foundation, 2024.

[12] Tong Li, Zhaoyang Liu, Yanyan Shen, Xue Wang, Haokun Chen, and Sen Huang. MASTER: Market-guided stock transformer for stock price forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38(1), pages 9124–9132. AAAI Press, 2024.

[13] T. M. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363, 1996. doi: 10.1109/18.485708.

[14] Alankrita Bhatt, J. Jon Ryu, and Young-Han Kim. On universal portfolios with continuous side information. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS 2023)*, volume 206 of *Proceedings of Machine Learning Research*, pages 4147–4163, Valencia, Spain, 2023. PMLR.

[15] Ziliang Zhang, Tianming Zhao, and Albert Y. Zomaya. Online portfolio selection with ML predictions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. URL `https://openreview.net/forum?id=zOFxp98km2`. NeurIPS 2025 (poster).

[16] T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 7, pages 5201–5213, Stockholm, Sweden, July 2018. International Machine Learning Society (IMLS). ISBN 978-151086796-3. Held from July 10–15, 2018, Code 141700.

[17] Ravi Kumar, Manish Purohit, and Zoya Svitkina. Improving online algorithms via ml predictions. In H. Larochelle, N. Cesa-Bianchi, H. Wallach, K. Grauman, S. Bengio, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 2018-December, pages 9661–9670, Montreal, Canada, December 2018. Neural Information Processing Systems Foundation. Held from December 2–8, 2018, Code 147412.

[18] Davin Choo, Themistoklis Gouleakis, Chun Kai Ling, and Arnab Bhattacharyya. Online bipartite matching with imperfect advice. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 8762–8781. PMLR, 2024.

[19] Michael Mitzenmacher and Sergei Vassilvitskii. Algorithms with predictions. *Commun. ACM*, 65(7):33–35, June 2022. ISSN 0001-0782. doi: 10.1145/3528087.

[20] ALPS. ALPS: Algorithms with Predictions. `https://algorithms-with-predictions.github.io/`, 2024. Accessed May 10, 2025.

[21] Yongho Shin, Changyeol Lee, Gukryeol Lee, and Hyung-Chan An. Improved learning-augmented algorithms for the multi-option ski rental problem via best-possible competitive analysis. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 31539–31561. PMLR, 2023.

[22] Mohammadreza Daneshvaramoli, Helia Karisani, Adam Lechowicz, Bo Sun, Cameron N. Musco, and Mohammad Hajiesmaili. Near-optimal consistency-robustness trade-offs for learning-augmented online knapsack problems. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 12459–12489. PMLR, 2025.

[23] Bo Sun, Russell Lee, Mohammad Hajiesmaili, Adam Wierman, and Danny H. K. Tsang. Pareto-optimal learning-augmented algorithms for online conversion problems. In *Advances in Neural Information Processing Systems*, volume 34, pages 10339–10350, 2021.

[24] Ziyad Benomar, Lorenzo Croissant, Vianney Perchet, and Spyros Angelopoulos. Pareto-optimality, smoothness, and stochasticity in learning-augmented one-max-search. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025. URL `https://openreview.net/forum?id=ZTOVlPC5Vf`. ICML 2025 (poster).

[25] Alexander Wei and Fred Zhang. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. In *Advances in Neural Information Processing Systems*, volume 33, pages 8042–8053, 2020.

[26] Harry Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952. doi: 10.1111/j.1540-6261.1952.tb01525.x.

[27] 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.

[28] 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.

[29] Vijay K. Chopra and William T. Ziemba. The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management*, 19(2):6–11, 1993. doi: 10.3905/jpm.1993.409440.

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

[31] L. Bin. Nyse (o) dataset. `http://www.mysmu.edu.sg/faculty/chhoi/olps/datasets/NYSE_O_Dataset.html`, 2024. Accessed January 14, 2025.

[32] L. Bin. Nyse (n) dataset. `http://www.mysmu.edu.sg/faculty/chhoi/olps/datasets/NYSE_N_2_Dataset.html`, 2024. Accessed January 14, 2025.

[33] Andrewmvd. S&P 500 Stocks (daily-updated) — kaggle dataset. `https://www.kaggle.com/datasets/andrewmvd/sp-500-stocks`, 2024. Dataset accessed 10 May 2025.

[34] Cristian Challu, Kin G. Olivares, Boris N. Oreshkin, Federico Garza, Max Mergenthaler-Canseco, and Artur Dubrawski. Neural hierarchical interpolation for time series forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(6):6989–6997, 2023. doi: 10.1609/aaai.v37i6.25854.

[35] Kin G. Olivares, Cristian Challú, Azul Garza, Max Mergenthaler Canseco, and Artur Dubrawski. NeuralForecast: User friendly state-of-the-art neural forecasting models. PyCon Salt Lake City, Utah, US 2022, 2022. URL https://github.com/Nixtla/neuralforecast.

A Technical Appendix

Proposition A.1 (Maximin optimality of uCRP)**.** *Fix a horizon $n \geq 1$ and the number of assets $m \geq 2$. In each period i, an online player/policy A chooses a portfolio $\mathbf{b}(i) \in \Delta^{m-1}$, then an adaptive adversary observes $\mathbf{b}(i)$ and chooses returns $\mathbf{x}(i) \in \mathbb{R}^m_{>0}$ with $\sum_{j=1}^m x_j(i) = W(i) > 0$. The player's wealth is $S_n^A(\mathbf{x}^n) = \prod_{i=1}^n \mathbf{b}(i) \cdot \mathbf{x}(i)$. Then,*

$$\sup_A \inf_{\mathbf{x}^n} S_n^A(\mathbf{x}^n) = \prod_{i=1}^n \frac{1}{m} \sum_{j=1}^m x_j(i)$$

and this value is attained by the uniform CRP policy $b_j(i) = \frac{1}{m}\ \forall j \forall i$. Hence, no online algorithm guarantees (instance-wise, against an adaptive adversary) more wealth than uCRP, and uCRP achieves the optimal worst-case guarantee.

Proof. We decompose the adversary game into independent per-round problems and compose the bounds multiplicatively. Let $W := \sum_{j=1}^m x_j > 0$ and $\mathbf{b} \in \Delta^{m-1}$. Then, $\sum_{j=1}^m b_j x_j \geq (\min_j b_j) \sum_{j=1}^m x_j = W \cdot \min_j b_j$. To obtain the per-round maximin value, solve:

$$\max_{\mathbf{b} \in \Delta^{m-1}} \inf_{\mathbf{x} \in \mathbb{R}^m_{>0}, \sum_{j=1}^m x_j = W} \sum_{j=1}^m b_j x_j = \max_{\mathbf{b} \in \Delta^{m-1}} W \cdot \min_j b_j$$

The function is maximized over the simplex when all coordinates are equal, i.e., at $b_j = \frac{1}{m}\ \forall j$, yielding the value $\frac{1}{m}W$. Therefore the per-round maximin value is $\frac{1}{m}W$, and it is attained by the uniform portfolio. Let A be any online policy producing $\mathbf{b}(1), \cdots, \mathbf{b}(n)$. Applying the per-round inequality to obtain

$$\inf_{\mathbf{x}(i) \in \mathbb{R}^m_{>0}, \sum_{j=1}^m x_j(i) = W(i)} = W(i) \cdot \min_j b_j \leq \frac{W(i)}{m}$$

Then,

$$\inf_{\mathbf{x}^n} S_n^A(\mathbf{x}^n) = \inf_{\mathbf{x}^n} \prod_{i=1}^n \sum_{j=1}^m b_j(i) x_j(i) = \prod_{i=1}^n \inf_{\mathbf{x}(i) \in \mathbb{R}^m_{>0}, \sum_{j=1}^m x_j(i) = W(i)} \sum_{j=1}^m b_j(i) x_j(i)$$

$$= \prod_{i=1}^n W(i) \cdot \min_j b_j(i) \leq \prod_{i=1}^n \frac{W(i)}{m}$$

The equality $\inf \prod = \prod \inf$ follows because the domain factorizes as a Cartesian product and the objective decomposes as a product of nonnegative terms on disjoint coordinates: the "$\geq$" direction is immediate from each factor being bounded below by its infimum, and the reverse inequality "$\leq$" follows by ϵ-approximating each infimum simultaneously. Choosing $b_j(i) = \frac{1}{m}\ \forall j \forall i$ achieves the RHS for every feasible $\mathbf{x}(i)$: $\inf_{\mathbf{x}^n} S_n^{uCRP}(\mathbf{x}^n) = \prod_{i=1}^n \frac{1}{m} \sum_{j=1}^m x_j(i)$. □

Lemma A.2 (Residual inefficiency)**.** *Let $m \geq 2$, $\mathbf{u} = (\frac{1}{m}, \cdots, \frac{1}{m})$, $\mathbf{b} \in \Delta^{m-1}$, and fix $\sigma \in (0, 1]$. Then,*

$$\forall \mathbf{x} \in \mathbb{R}^m_{>0} : \sum_{j=1}^m b_j x_j \geq \sigma \prod_{j=1}^m x_j^{\frac{1}{m}} \iff \prod_{j=1}^m b_j \geq \left(\frac{\sigma}{m}\right)^m \iff D_{KL}(\mathbf{u}||\mathbf{b}) \leq \log \frac{1}{\sigma}$$

Proof. We first prove *sufficiency*. By AM-GM on the m nonnegative numbers $b_1 x_1, \cdots, b_m x_m$,

$$\sum_{j=1}^m b_j x_j \geq m \left(\prod_{j=1}^m b_j x_j\right)^{1/m} = m \left(\prod_{j=1}^m b_j\right)^{1/m} \left(\prod_{j=1}^m x_j\right)^{1/m}$$

If $\prod_{j=1}^m b_j \geq (\frac{\sigma}{m})^m$ holds, then

$$\sum_{j=1}^m b_j x_j \geq m \frac{\sigma}{m} \prod_{j=1}^m x_j^{\frac{1}{m}} = \sigma \prod_{j=1}^m x_j^{\frac{1}{m}}$$

Next, we prove *necessity* by first showing that $b_j > 0, \forall j$. Suppose $b_k = 0$ for some k. For $t > 0$, set $x_k = t$ and $x_j = 1$ for $j \neq k$, then

$$\sum_{j=1}^{m} b_j x_j = \sum_{j \neq k} b_j \cdot 1 + b_k \cdot t = 1 \quad \text{and} \quad \prod_{j=1}^{m} x_j{}^{\frac{1}{m}} = \sqrt[m]{t \cdot 1^{m-1}} = t^{\frac{1}{m}}$$

The asserted inequality would require $1 \geq \sigma t^{1/m}$ for all $t > 0$, which is impossible as t becomes large. Therefore $b_j > 0, \forall j$. Assume now the inequality holds for all $x_j > 0$. Set $x_j = c/b_j$ with $c > 0$. Then

$$\sum_{j=1}^{m} b_j x_j = \sum_{j=1}^{m} c = mc \quad \text{and} \quad \prod_{j=1}^{m} x_j^{\frac{1}{m}} = \left(\frac{c^m}{\prod_{j=1}^{m} b_j}\right)^{\frac{1}{m}} = \frac{c}{\prod_{j=1}^{m} b_j^{\frac{1}{m}}}$$

Substitute back gives the desired necessity,

$$\sum_{j=1}^{m} b_j x_j \geq \sigma \prod_{j=1}^{m} x_j^{\frac{1}{m}} \implies mc \geq \frac{\sigma c}{\prod_{j=1}^{m} b_j^{\frac{1}{m}}} \implies \prod_{j=1}^{m} b_j \geq \left(\frac{\sigma}{m}\right)^m$$

Lastly, the equivalence in relative entropy follows by algebra:

$$\prod_{j=1}^{m} b_j \geq \left(\frac{\sigma}{m}\right)^m \iff \sum_{j=1}^{m} \log b_j \geq m \log \frac{\sigma}{m} \iff \frac{1}{m} \sum_{j=1}^{m} \log b_j \geq \log \frac{1}{m} - \log \frac{1}{\sigma}$$

$$\iff \frac{1}{m} \sum_{j=1}^{m} \log \frac{1}{m} - \frac{1}{m} \sum_{j=1}^{m} \log b_j \leq \log \frac{1}{\sigma} \iff \frac{1}{m} \sum_{j=1}^{m} \log \frac{1}{m b_j} = D_{\mathrm{KL}}(\mathbf{u} \| \mathbf{b}) \leq \log \frac{1}{\sigma}$$

□

Corollary A.3 (Uniform initialization and maintenance)**.** *Let $m \geq 2$, $\mathbf{x} \in \mathbb{R}^m_{>0}$, and $\sum_{j=1}^m b_j = 1$ with $b_j \geq 0$. Then,*

$$\forall \mathbf{x}^n \in \mathbb{R}^m_{>0} : \sum_{j=1}^{m} b_j x_j \geq \prod_{j=1}^{m} x_j^{\frac{1}{m}} \iff \mathbf{b} = \left(\frac{1}{m}, \cdots, \frac{1}{m}\right)$$

Proof. Applying Lemma A.2 with $\sigma = 1$ gives $\prod_{j=1}^m b_j \geq m^{-m}$. On the other hand, by AM-GM applied to $b_1, \cdots, b_m$

$$\frac{1}{m} = \frac{1}{m} \sum_{j=1}^{m} b_j \geq \prod_{j=1}^{m} b_j^{1/m} \implies \prod_{j=1}^{m} b_j \leq m^{-m}$$

Combining the two inequalities $\prod_{j=1}^m b_j \leq m^{-m}$ and $\prod_{j=1}^m b_j \geq m^{-m}$ yields $\prod_{j=1}^m b_j = m^{-m}$, and equality in AM-GM implies $b_j = \frac{1}{m}\ \forall j$, which gives $\mathbf{b} = (\frac{1}{m}, \cdots \frac{1}{m})$. □

Lemma A.4 (Characterization of GM-robust portfolios)**.** *Fix $m \geq 2$. Consider any portfolio strategy A that, at each round i, maps the past history $\mathbf{x}^{i-1}$ to a portfolio $\mathbf{b}(i) = A(\mathbf{x}^{i-1}) \in \Delta^{m-1}$. For any sequence $\mathbf{x}^n = \{\mathbf{x}(i)\}_{i=1}^n$ with $\mathbf{x}(i) \in \mathbb{R}^m_{>0}$, define $S_n^A(\mathbf{x}^n) := \prod_{i=1}^n \sum_{j=1}^m b_j(i) x_j(i)$ and $S_n^{GM}(\mathbf{x}^n) := \prod_{i=1}^n \prod_{j=1}^m x_j(i)^{\frac{1}{m}}$, with $S_0^A(\cdot) = S_0^{GM}(\cdot) = 1$. Then,*

$$\forall \mathbf{x}^n : S_n^A(\mathbf{x}^n) \geq S_n^{GM}(\mathbf{x}^n) \iff \forall i \in \{1, \cdots, n\}, \forall \mathbf{x}^{i-1} : \prod_{j=1}^{m} b_j(i) \geq \left(\frac{S_{i-1}^{GM}(\mathbf{x}^{i-1})}{m S_{i-1}^A(\mathbf{x}^{i-1})}\right)^m$$

In particular, the condition at unit-horizon $n = 1$ reduces to $\prod_{j=1}^m b_j(1) \geq \frac{1}{m^m}$ which forces $\mathbf{b}(1) = (\frac{1}{m}, \cdots, \frac{1}{m})$.

Proof. Sufficiency. Assume for every period i and every history $\mathbf{x}^{i-1}$ the condition $\prod_{j=1}^m b_j(i) \geq (S_{i-1}^{GM}/mS_{i-1}^A)^m$ holds. We prove $S_i^A \geq S_i^{GM}$ by induction on i. Base $i = 1$. The condition at $i = 1$ is $\prod_{j=1}^m b_j(1) \geq \frac{1}{m^m}$ which by Corollary A.3 is equivalent to $\mathbf{b}(1) = (\frac{1}{m}, \cdots, \frac{1}{m})$, and $\frac{1}{m}\sum_{j=1}^m x_j(1) \geq \prod_{j=1}^m x_j(1)^{1/m}$. Therefore, $S_1^A \geq S_1^{GM}$. Inductive step. Suppose $S_{i-1}^A \geq S_{i-1}^{GM}$, and set $\sigma_i = S_{i-1}^{GM}/S_{i-1}^A \in (0, 1]$. By Lemma A.2 and the product condition $\prod_{j=1}^m b_j(i) \geq (\sigma_i/m)^m$, for all $\mathbf{x}(i) > 0$, we have

$$A_i = \sum_{j=1}^m b_j(i)x_j(i) \geq \sigma_i \prod_{j=1}^m x_j(i)^{1/m}$$

Multiplying by S_{i-1}^A on both sides yields

$$S_i^A = S_{i-1}^A A_i \geq S_{i-1}^A \sigma_i \prod_{j=1}^m x_j(i)^{1/m} = S_i^{GM}$$

Thus $S_n^A \geq S_n^{GM}$, $\forall \mathbf{x}^n$. Thereby proving sufficiency.

Necessity. To prove necessity, we follow the logic chain: (i) Global domination implies local prefix domination, (ii) implies per-round GM lower bound, (iii) implies product condition via Lemma A.2, and (iv) prove the special case $i = 1$. (i) We begin by showing that

$$\forall \mathbf{x}^n : S_n^A \geq S_n^{GM} \implies \forall i \in \{1, \cdots, n\}, \forall \mathbf{x}^{i-1} : S_{i-1}^A \geq S_{i-1}^{GM}$$

Fix arbitrary $i \in \{1, \cdots, n\}$ and any history $\mathbf{x}^{i-1} \in (\mathbb{R}_{>0}^m)^{i-1}$. Extend $\mathbf{x}^{i-1}$ to a full sequence $\mathbf{x}^n$ by setting $\mathbf{x}(t) := \mathbf{1}$ for all $t = i, \cdots, n$. For this particular $\mathbf{x}^n$, we have $1 = \sum_{j=1}^m b_j(t) \cdot 1$ and $1 = \prod_{j=1}^m 1^{1/m}$ for all $t \geq i$. Therefore,

$$S_n^A(\mathbf{x}^n) = S_{i-1}^A(\mathbf{x}^{i-1}) \cdot 1 = S_{i-1}^A(\mathbf{x}^{i-1}) \quad \text{and} \quad S_n^{GM}(\mathbf{x}^n) = S_{i-1}^{GM}(\mathbf{x}^n) \cdot 1 = S_{i-1}^{GM}(\mathbf{x}^n)$$

Applying the global assumption to this $\mathbf{x}^n$, we get $S_{i-1}^A(\mathbf{x}^{i-1}) \geq S_{i-1}^{GM}(\mathbf{x}^{i-1})$. Since i and $\mathbf{x}^{i-1}$ are chosen arbitrarily, this inequality holds $\forall i \in \{1, \cdots, n\}$ and $\forall \mathbf{x}^{i-1}$. Consequently $\forall i$ and $\forall \mathbf{x}^{i-1}$, the residual factor

$$\sigma_i(\mathbf{x}^{i-1}) := \frac{S_{i-1}^{GM}(\mathbf{x}^{i-1})}{S_{i-1}^A(\mathbf{x}^{i-1})} \in (0, 1]$$

is well-defined. (ii) Now fix arbitrary $i \in \{1, \cdots, n\}$ and arbitrary history $\mathbf{x}^{i-1}$. Let $\mathbf{x}(i) \in \mathbb{R}_{>0}^m$, and again extend to length n by setting $\mathbf{x}(t) := \mathbf{1}$ for all $t = i+1, \cdots, n$. For the extended sequence $\mathbf{x}^n$, we have

$$S_n^A(\mathbf{x}^n) = 1 \cdot S_{i-1}(\mathbf{x}^{i-1}) \cdot \sum_{j=1}^m b_j(i)x_j(i) \quad \text{and} \quad S_n^{GM}(\mathbf{x}^n) = 1 \cdot S_{i-1}^{GM} \cdot \prod_{j=1}^m x_j(i)^{\frac{1}{m}}$$

Now, apply the global assumption of $S_n^A \geq S_n^{GM}$ to this extended $\mathbf{x}^n$ gives

$$S_{i-1}^A(\mathbf{x}^{i-1}) \cdot \sum_{j=1}^m b_j(i)x_j(i) \geq S_{i-1}^{GM} \cdot \prod_{j=1}^m x_j(i)^{\frac{1}{m}}$$

Divide both sides by $S_{i-1}^A > 0$ and apply $\sigma_i = S_{i-1}^{GM}/S_{i-1}^A \in (0, 1]$ in (i) gives

$$\sum_{j=1}^m b_j(i)x_j(i) \geq \sigma_i \cdot \prod_{j=1}^m x_j(i)^{\frac{1}{m}}$$

for arbitrary $\mathbf{x}(i) \in \mathbb{R}_{>0}^m$ at period i. (iii) This inequality condition is exactly the necessary and sufficient condition for $\prod_{j=1}^m b_j(i) \geq (\sigma_i/m)^m$ as proven in Lemma A.2. Thus, for this fixed i and fixed $\mathbf{x}^{i-1}$, we conclude

$$\prod_{j=1}^m b_j(i) \geq \left(\frac{S_{i-1}^{GM}(\mathbf{x}^{i-1})}{mS_{i-1}^A(\mathbf{x}^{i-1})}\right)^m = \left(\frac{\sigma_i}{m}\right)^m$$

Since i and $\mathbf{x}^{i-1}$ are arbitrary, the above holds $\forall i$ and $\forall \mathbf{x}^{i-1}$. (iv) Finally, at unit-horizon $n = 1$, apply Corollary A.3 proves the $i = 1$ case, and the weights must be uniform over m. □

Lemma A.5 (Per-round optimal consistency)**.** *Fix $m \geq 2$, let $\sigma \in (0,1]$, and specify a portfolio with constraints $\mathbf{b} \in \mathcal{B}(\sigma) = \{\Delta^{m-1} : \prod_{j=1}^m b_j \geq (\frac{\sigma}{m})^m\}$. Claim*

1. ***Facts**: $\mathcal{B}(\sigma)$ is nonempty, convex and compact, and any $\mathbf{b} \in \mathcal{B}(\sigma)$ satisfies $b_j > 0\ \forall j$. At least one optimizer $\mathbf{b}^* \in \mathcal{B}(\sigma)$ exists for the below optimization problem*

$$\forall \sigma \in (0,1], \forall \mathbf{x} \in \mathbb{R}^m_{>0} : \max_{\mathbf{b} \in \mathcal{B}(\sigma)} \sum_{j=1}^m b_j x_j$$

2. ***Degenerate cases**:*
   - *If $\sigma = 1$, then $\mathcal{B}(1)$ is a singleton, and the unique optimizer is the uniform portfolio $\mathbf{b}^* = \frac{1}{m}\mathbf{1}$ with the optimal value $\frac{1}{m}\sum_{j=1}^m x_j$ for all $\mathbf{x} \in \mathbb{R}^m_{>0}$.*
   - *If all x_js equal such that $x_1 = \cdots = x_m = c > 0$, then every feasible $\mathbf{b} \in \mathcal{B}(\sigma)$ is optimal with the optimal value $\frac{1}{m}\sum_{j=1}^m x_j$ for any $\sigma \in (0,1]$.*

3. ***Generic case**: For any $\sigma \in (0,1)$ and $\mathbf{x}$ whose coordinates are not all equal. There exists $\mu^* > \max_j x_j$ solving*

$$\sigma \cdot \sum_{j=1}^m \frac{1}{\mu - x_j} \cdot \prod_{j=1}^m (\mu - x_j)^{\frac{1}{m}} = m$$

   The global optimizer $\mathbf{b}^* \in \mathcal{B}(\sigma)$ exists and solves by

$$b_j^* = \frac{\sigma \prod_{k=1}^m (\mu^* - x_k)^{\frac{1}{m}}}{m(\mu^* - x_j)},\ \forall j$$

   The optimal value is attained at $\mathbf{b}^*$ yielding maximum

$$\max_{\mathbf{b} \in \mathcal{B}(\sigma)} \sum_{j=1}^m b_j x_j = \mu^* - \sigma \prod_{j=1}^m (\mu^* - x_j)^{\frac{1}{m}}$$

Proof. We start proving claim 1. Nonemptiness: the uniform portfolio $\overline{\mathbf{b}} = (\frac{1}{m}, \cdots, \frac{1}{m})$ satisfies both constraints for all $\sigma \in (0,1]$. Positivity: $\mathbf{b} \in \mathcal{B}(\sigma) \implies \prod_{j=1}^m b_j \geq (\frac{\sigma}{m})^m \implies b_j > 0\ \forall j$. Convexity: consider two sets separately; i) the simplex $\Delta^{m-1} = \{\mathbf{b} \in \mathbb{R}^m : b_j \geq 0, \sum_{j=1}^m b_j = 1\}$ is convex; ii) $\prod_{j=1}^m b_j \geq (\frac{\sigma}{m})^m \iff \sum_{j=1}^m \log b_j \geq m \log \frac{\sigma}{m}$ where $\sum_{j=1}^m \log b_j$ is the sum of concave functions with each function being concave in $b_j > 0$. The product constraint defines a superlevel set of a concave function, so the product constraint set is convex. Subsequently, the intersection of convex sets remain convex. Compactness: the simplex $\Delta^{m-1} = \{\mathbf{b} \in \mathbb{R}^m : b_j \geq 0, \sum_{j=1}^m b_j = 1\}$ is closed and bounded, hence compact. The product map $g(\mathbf{b}) = \prod_{j=1}^m b_j$ is continuous, so the set $\{\mathbf{b} \in \mathbb{R}^m : g(\mathbf{b}) \geq (\frac{\sigma}{m})^m\}$ is the preimage of a closed interval under a continuous map, and is closed. Therefore, $\mathcal{B}(\sigma)$ is a closed subset of the compact simplex, hence compact. The objective $\sum_{j=1}^m b_j x_j$ is linear, thus continuous. Since we are maximizing a continuous function over a compact set $\mathcal{B}(\sigma)$, the extreme value theorem guarantees that there exists at least one global maximizer $\mathbf{b}^* \in \mathcal{B}(\sigma)$ attaining optimum.

Now we prove for degenerate cases. (i) Suppose $\sigma = 1$, the product constraint is $\prod_{j=1}^m b_j \geq m^{-m}$, but simplex feasibility and AM-GM demand $\prod_{j=1}^m b_j \leq ((b_1 + \cdots + b_m)/m)^m = m^{-m}$, forcing $\prod_{j=1}^m b_j = m^{-m}$. AM-GM attains equality only if all b_j are equal (noting that all $b_j > 0$), so $\mathcal{B}(1) = \{(\frac{1}{m}, \cdots, \frac{1}{m})\}$, and the optimum $\frac{1}{m}\sum_{j=1}^m x_j$ follows naturally. (ii) Suppose all x_js are equal and $x_1 = \cdots = x_m = c > 0$. Then for any $\mathbf{b}$, $\sum_{j=1}^m b_j x_j = \sum_{j=1}^m b_j c = c$. So every feasible $\mathbf{b} \in \mathcal{B}(\sigma)$ gives the same value c. Thus, every feasible $\mathbf{b} \in \mathcal{B}(\sigma)$ is optimal. Conveniently, write the optimum $\frac{1}{m}\sum_{j=1}^m x_j$.

Lastly, we solve the maximization problem in the generic case. (i) [Objective functions] Define

$$F(\mathbf{b}) := -\sum_{j=1}^m b_j x_j, \quad g(\mathbf{b}) := m \log\left(\frac{\sigma}{m}\right) - \sum_{j=1}^m \log b_j \leq 0, \quad h(\mathbf{b}) := \sum_{j=1}^m b_j - 1 = 0$$

The original optimization is equivalent to:

$$\min_{\mathbf{b}} F(\mathbf{b}) = -\sum_{j=1}^{m} b_j x_j \quad s.t. \quad g(\mathbf{b}) \leq 0, h(\mathbf{b}) = 0$$

(ii) [The optimization problem is convex] Because $F(\mathbf{b})$ is linear in $\mathbf{b}$ hence convex, $g(\mathbf{b})$ is the negative of a concave function hence convex, $h(\mathbf{b})$ is an affine function hence convex (and concave). Therefore, minimizing a convex objective under a convex inequality constraint and an affine equality constraint is a convex optimization problem. (iii) [Slater's condition holds] Take the uniform portfolio $\tilde{\mathbf{b}} = (\frac{1}{m}, \cdots, \frac{1}{m})$, we have

$$g(\tilde{\mathbf{b}}) = m \log \frac{\sigma}{m} - m \log \frac{1}{m} < 0, \quad \forall \sigma \in (0,1) \quad \text{and} \quad h(\tilde{\mathbf{b}}) = \frac{m}{m} - 1 = 0$$

There exists a strictly feasible point $\tilde{\mathbf{b}}$ with $g(\tilde{\mathbf{b}}) < 0$ and $h(\tilde{\mathbf{b}}) = 0$, hence Slater's condition holds. Since the problem is convex and Slater's condition holds (by the uniform portfolio), KKT conditions are necessary and sufficient for optimality. (iv) [KKT system]

1. Stationarity requires $\nabla F(\mathbf{b}^*) + \lambda \nabla g(\mathbf{b}^*) + \mu \nabla h(\mathbf{b}^*) = 0$.
2. Primal feasibility requires $g(\mathbf{b}^*) \leq 0$, $h(\mathbf{b}^*) = 0$, $b_j^* > 0$.
3. Dual feasibility requires $\lambda \geq 0$.
4. Complementary slackness gives $\lambda g(\mathbf{b}^*) = 0$.

By computing first derivatives, $\partial F / \partial b_j = -x_j$, $\partial g / \partial b_j = -1/b_j$, $\partial h / \partial b_j = 1$, so stationarity gives

$$-x_j + \lambda \left(-\frac{1}{b_j^*} \right) + \mu \cdot 1 = 0 \iff x_j + \frac{\lambda}{b_j^*} - \mu = 0 \iff b_j^* = \frac{\lambda}{\mu - x_j} \quad , \forall j$$

Because $b_j^* > 0$ for all j, we must have $\mu > \max_j x_j$. (v) [The product constraint is active at the optimum $\mathbf{b}^*$] Complementary slackness gives $\lambda \cdot g(\mathbf{b}^*) = 0$. We examine two possible cases: either $\lambda = 0$ or $g(\mathbf{b}^*) = 0$. If $\lambda = 0$, stationarity becomes $x_j = \mu$ for all j, which means all x_j are equal, contradicting the generic case assumption on x_js are not all equal. Hence $\lambda \neq 0$, hence $g(\mathbf{b}^*) = 0$, which means that the product constraint is active at the optimum. (vi) [Solution] We have $\sum_{j=1}^m b_j^* = 1$, $\prod_{j=1}^m b_j^* = (\sigma/m)^m$ and $b_j^* = \lambda/(\mu - x_j)\ \forall j$ where $\mu > \max_j x_j$. Now enforce $\sum_{j=1}^m b_j^* = 1$, which gives

$$\sum_{j=1}^{m} \frac{\lambda}{\mu - x_j} = 1 \iff \lambda = \left(\sum_{j=1}^{m} \frac{1}{\mu - x_j} \right)^{-1}$$

Next enforce $\prod_{j=1}^m b_j^* = (\frac{\sigma}{m})^m$, which gives

$$\prod_{j=1}^{m} \frac{\lambda}{\mu - x_j} = \left(\frac{\sigma}{m} \right)^m \iff \lambda = \frac{\sigma}{m} \prod_{j=1}^{m} (\mu - x_j)^{1/m}$$

Lastly equate over λ, which gives

$$\left(\sum_{j=1}^{m} \frac{1}{\mu - x_j} \right)^{-1} = \frac{\sigma}{m} \prod_{j=1}^{m} (\mu - x_j)^{1/m} \iff \sigma \cdot \sum_{j=1}^{m} \frac{1}{\mu - x_j} \cdot \prod_{j=1}^{m} (\mu - x_j)^{1/m} = m$$

(vii) [Existence of μ^*] We show μ^* exists by continuity argument. Define

$$S(\mu) := \sum_{j=1}^{m} \frac{1}{\mu - x_j}, \quad P(\mu) := \prod_{j=1}^{m} (\mu - x_j)^{\frac{1}{m}}, \quad f(\mu) := \sigma S(\mu) P(\mu) - m$$

We show $f(\mu) = 0$ has at least one solution. Let $\hat{x} = \max_j x_j$. First consider the case where μ approaches $\hat{x}$, we have $S(\mu) \sim \frac{1}{\mu - \hat{x}}$, $P(\mu) \sim k(\mu - \hat{x})^{\frac{1}{m}}$ for some $k > 0$. Then, $S(\mu)P(\mu) \sim$

$k(\mu-\hat{x})^{\frac{1}{m}}(\mu-\hat{x})^{-1} \to +\infty$. So $f(\mu)$ approaches $+\infty$ as μ approaches $\hat{x}$. Next consider μ approaches $+\infty$, $\mu - x_j \sim \mu$, $S(\mu) \sim \frac{m}{\mu}$, and $P(\mu) \sim \mu$. So $S(\mu)P(\mu) \to m$, and $f(\mu) \to \sigma m - m = m(\sigma-1) < 0$ since $\sigma \in (0,1)$. Given $f(\mu)$ is continuous on $(\hat{x}, +\infty)$, by the intermediate value theorem, there is at least one $\mu^* \in (\hat{x}, +\infty)$ with $f(\mu^*) = 0$. (viii) [Optimizer and optimum] Once μ^* is solved, set

$$\lambda = \frac{\sigma}{m}\prod_{j=1}^{m}(\mu^* - x_j)^{\frac{1}{m}}$$

and solve for

$$b_j^* = \frac{\sigma\prod_{j=1}^{m}(\mu^*-x_j)^{\frac{1}{m}}}{m(\mu^*-x_j)}, \quad \forall j$$

By construction, $\mathbf{b}^*$ satisfies all the constraints, and attains the global optimum. The final equation of the optimum follows by algebra. Consider the two identities

$$x_j = \mu^* - (\mu^* - x_j) \quad \text{and} \quad b_j^*(\mu^*-x_j) = \frac{\sigma\prod_{k=1}^{m}(\mu^*-x_k)^{1/m}}{m(\mu^*-x_j)}(\mu^*-x_j) = \frac{\sigma}{m}\prod_{k=1}^{m}(\mu^*-x_k)^{\frac{1}{m}}$$

then by algebra

$$\sum_{j=1}^{m} b_j^* x_j = \mu^*\sum_{j=1}^{m} b_j^* - \sum_{j=1}^{m} b_j^*(\mu^*-x_j) = \mu^* - \sum_{j=1}^{m}\frac{\sigma}{m}\prod_{k=1}^{m}(\mu^*-x_k)^{\frac{1}{m}} = \mu^* - \sigma\prod_{j=1}^{m}(\mu^*-x_j)^{\frac{1}{m}}$$

□

Lemma A.6. *Let* $\mathcal{B}(\sigma) = \{\mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^{m} b_j \geq (\frac{\sigma}{m})^m\}$, $0 < \sigma < 1$. *For all* $\mathbf{x} \in \mathbb{R}^m_{>0}$,

$$\forall \alpha > 0, \beta \in \mathbb{R} : \mathbf{y} = \alpha\mathbf{x} + \beta\mathbf{1} \implies \arg\max_{\mathbf{b}\in\mathcal{B}(\sigma)} \mathbf{b}\cdot\mathbf{x} = \arg\max_{\mathbf{b}\in\mathcal{B}(\sigma)} \mathbf{b}\cdot\mathbf{y}$$

Proof. Suppose $\mathbf{y} = \alpha\mathbf{x} + \beta\mathbf{1}$ with $\alpha > 0$. Since $\mathcal{B}(\sigma) \subseteq \Delta^{m-1}$, every feasible $\mathbf{b}$ satisfies $\mathbf{b}\cdot\mathbf{1} = 1$. Hence, $\mathbf{b}\cdot\mathbf{y} = \mathbf{b}\cdot(\alpha\cdot\mathbf{x} + \beta\cdot\mathbf{1}) = \alpha\mathbf{b}\cdot\mathbf{x} + \beta$. Therefore, for any $\mathbf{b}, \mathbf{b}' \in \mathcal{B}(\sigma)$,

$$\mathbf{b}\cdot\mathbf{y} \geq \mathbf{b}'\cdot\mathbf{y} \iff \alpha\mathbf{b}\cdot\mathbf{x} + \beta \geq \alpha\mathbf{b}'\cdot\mathbf{x} + \beta \iff \mathbf{b}\cdot\mathbf{x} \geq \mathbf{b}'\cdot\mathbf{x}$$

because $\alpha > 0$. The two objectives (maximization) include exactly the same ordering over all feasible $\mathbf{b}$, so their optimizer sets are identical. □

Theorem A.7 (Maximal-optimality at the geometric mean)**.** *Fix* $m \geq 2$. *Let* $\mathbf{x}^n = \{\mathbf{x}(i)\}_{i=1}^n$ *be an arbitrary return sequence and* $\mathbf{y}^n = \{\mathbf{y}(i)\}_{i=1}^n$ *be an arbitrary prediction sequence, where* $\mathbf{x}(i), \mathbf{y}(i) \in \mathbb{R}^m_{>0}$. *For* $\sigma \in (0,1]$, *define* $\mathcal{B}(\sigma) := \{\mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^{m} b_j \geq \left(\frac{\sigma}{m}\right)^m\}$. *Denote* $S_n^{GM}(\mathbf{x}^n) := \prod_{i=1}^{n}\prod_{j=1}^{m} x_j(i)^{1/m}$ *as the GM wealth with* $S_0^{GM} := 1$. *Denote* $S_n^{MOP}(\mathbf{x}^n, \mathbf{y}^n)$ *as the wealth of Algorithm 1 with* $S_0^{MOP} := 1$. *For each period* $i \geq 1$ *define the MOP residual* $\sigma_i^{MOP} := S_{i-1}^{GM}(\mathbf{x}^{i-1})/S_{i-1}^{MOP}(\mathbf{x}^{i-1}, \mathbf{y}^{i-1})$, *the per-period portfolio decision* $\hat{\mathbf{b}}(i) \in \arg\max_{\mathbf{b}\in\mathcal{B}(\sigma_i^{MOP})} \mathbf{b}\cdot\mathbf{y}(i)$, *and the per-period realized wealth* $S_i^{MOP}(\mathbf{x}^i, \mathbf{y}^i) = S_{i-1}^{MOP}(\mathbf{x}^{i-1}, \mathbf{y}^{i-1})\hat{\mathbf{b}}(i)\cdot\mathbf{x}(i)$. *Then, the following statements hold.*

1. ***Robustness***. *For all* $\mathbf{x}^i$, $\mathbf{y}^i$, $i = 1, \cdots, n$: $S_i^{MOP}(\mathbf{x}^i, \mathbf{y}^i) \geq S_i^{GM}(\mathbf{x}^i)$
2. ***Consistency***. *Suppose* $\overset{*}{\mathbf{y}}(i) = \alpha_i\mathbf{x}(i) + \beta_i\mathbf{1}$ *for some* $\alpha_i > 0$ *at every period* i. *Denote the optimal residual sequence* $\sigma_{1:n}^*$ *of MOP conditioned on* $\overset{*}{\mathbf{y}}(i)$ $\forall i$ *with* $\sigma_i^* = S_{i-1}^{GM}(\mathbf{x}^{i-1})/S_{i-1}^{MOP}(\mathbf{x}^{i-1}, \{\overset{*}{\mathbf{y}}(i)\}_{t=1}^{i-1})$. *Denote* $\mathbf{1}_{G_i} := 1$ *the indicator of a rare event* G_i *that* $G_i := \{\sigma_i^* = 1\} \cup \{x_1(i) = \cdots = x_m(i)\}$, *and write* $\mathbf{1}_{G_i^c} := 1 - \mathbf{1}_{G_i}$ *for its complement. Under the event* G_i^c, *let* $\mu_i^* > \max_j x_j(i)$ *be the solution of* $\sigma_i^*\sum_{j=1}^{m}(\mu_i - x_j(i))^{-1}\prod_{j=1}^{m}(\mu_i - x_j(i))^{1/m} = m$. *Then* $S_n^{MOP}(\mathbf{x}^n, \overset{*}{\mathbf{y}}^n)$ *attains the maximum* $S_n^*(\mathbf{x}^n|\sigma_{1:n}^*)$

$$S_n^{MOP} = S_n^*(\mathbf{x}^n|\sigma_{1:n}^*)$$

$$= \prod_{i=1}^{n} \mathbf{1}_{G_i} \frac{1}{m} \sum_{j=1}^{m} x_j(i) + \mathbf{1}_{G_i^c} (\mu_i^* - \sigma_i^* \prod_{j=1}^{m} (\mu_i^* - x_j(i))^{\frac{1}{m}})$$

$$\geq \frac{1}{m} \left[S_n^{OPT} + (m-1) \left(\frac{(S_n^{GM})^m}{S_n^{OPT}} \right)^{\frac{1}{m-1}} \right] \geq \frac{1}{m} S_n^{OPT}, \quad S_n^{OPT} := \prod_{i=1}^{n} \max_j x_j(i)$$

3. ***Smoothness***. *Define per-round smoothness factor* $\eta_i := \hat{\mathbf{b}}(i) \cdot \mathbf{x}(i) / \max_{\mathbf{b} \in \mathcal{B}(\sigma_i^*)} \mathbf{b} \cdot \mathbf{x}(i)$. *Then* $\eta_i \in (0, 1]$ *and the wealth of MOP evolves by*

$$S_n^{MOP}(\mathbf{x}^n, \mathbf{y}^n) = S_n^*(\mathbf{x}^n | \sigma_{1:n}^*) \prod_{i=1}^{n} \eta_i$$

4. *The strategy is maximal-optimal: no strategy that maintains the same per-round GM robustness can guarantee strictly larger consistency than* $S_n^*(\mathbf{x}^n)$ *for all* $\mathbf{x}^n$.

Proof. Robustness follows by Lemma A.4.

The exact consistency expression follows by Lemma A.5, combined with the fact that $\arg\max_{\mathbf{b}} \mathbf{b} \cdot \mathbf{x} = \arg\max_{\mathbf{b}} \mathbf{b} \cdot (\alpha \mathbf{x} + \beta \mathbf{1})$ for any $\alpha > 0$ from lemma A.6. We now prove the exact consistency lower bound $S_n^* \geq \frac{1}{m} \prod_{i=1}^{n} \max_j x_j(i)$ by proving the stronger invariant:

$$S_n^* \geq L_n := \frac{1}{m} \left[S_n^{OPT} + (m-1) \left(\frac{(S_n^{GM})^m}{S_n^{OPT}} \right)^{\frac{1}{m-1}} \right], \quad S_n^{OPT} := \prod_{i=1}^{n} \max_j x_j(i)$$

Since the second term inside the bracket is nonnegative, this implies $S_n^* \geq L_n \geq \frac{1}{m} S_n^{OPT}$.

First note $L_n \geq S_n^{GM}$, by taking one copy of S_n^{OPT} and $m-1$ copies of $((S_n^{GM})^m / S_n^{OPT})^{1/(m-1)}$ and apply the AM-GM inequality:

$$L_n = \frac{1}{m} \left[S_n^{OPT} + (m-1) \left(\frac{(S_n^{GM})^m}{S_n^{OPT}} \right)^{\frac{1}{m-1}} \right] \geq \left[S_n^{OPT} \cdot \frac{(S_n^{GM})^m}{S_n^{OPT}} \right]^{\frac{1}{m}} = S_n^{GM}$$

Now, for the base case $n = 1$, we have $\sigma_1^* = \frac{S_0^{GM}}{S_0^*} = 1$, which forms the singleton set $\mathcal{B}(1)$ containing only the uniform portfolio. Thus, $S_1^* = \frac{1}{m} \sum_{j=1}^{m} x_j(1)$. Let $j_1^* \in \arg\max_j x_j(1)$, then $S_1^{OPT} = x_{j_1^*}(1)$ and $((S_1^{GM})^m / S_1^{OPT})^{1/(m-1)} = (\prod_{j \neq j_1^*} x_j(i))^{1/(m-1)}$. By the AM-GM inequality on the $m-1$ non-winning coordinates:

$$\sum_{j \neq j_1^*} x_j(1) \geq (m-1) \left(\prod_{j \neq j_1^*} x_j(1) \right)^{\frac{1}{m-1}}$$

Therefore,

$$S_1^* = \frac{1}{m} \left[x_{j_1^*}(1) + \sum_{j \neq j_1^*} x_j(1) \right] \geq L_1$$

So the invariant holds at $n = 1$.

Inductive step. Suppose $S_{n-1}^* \geq L_{n-1}$. Then,

$$\sigma_n^* = \frac{S_{n-1}^{GM}}{S_{n-1}^*} \leq \frac{S_{n-1}^{GM}}{L_{n-1}}$$

Since $\mathcal{B}(\sigma)$ expands as σ decreases,

$$\mathcal{B}(\sigma_n^*) \supseteq \mathcal{B} \left(\frac{S_{n-1}^{GM}}{L_{n-1}} \right)$$

Therefore,

$$S_n^* = S_{n-1}^* \max_{\mathbf{b}\in\mathcal{B}(\sigma_n^*)} \mathbf{b}\cdot\mathbf{x}(n) \geq L_{n-1} \max_{\mathbf{b}\in\mathcal{B}(S_{n-1}^{GM}/L_{n-1})} \mathbf{b}\cdot\mathbf{x}(n)$$

It remains to lower bound the RHS maximum by constructing a feasible portfolio. Let $j_n^* \in \arg\max_j x_j(n)$. Define the portfolio $\overline{\mathbf{b}}$ by:

$$\text{for } j = j_n^* : \overline{b}_{j_n^*} = \frac{S_{n-1}^{OPT}}{mL_{n-1}}, \quad \text{for } j \neq j_n^* : \overline{b}_j = \frac{\left(\frac{(S_{n-1}^{GM})^m}{S_{n-1}^{OPT}}\right)^{\frac{1}{m-1}}}{mL_{n-1}}$$

We first verify that $\overline{\mathbf{b}} \in \Delta^{m-1}$. Each coordinate is strictly positive because all returns are positive. Moreover,

$$\sum_{j=1}^m \overline{b}_j = \frac{S_{n-1}^{OPT} + (m-1)\left(\frac{(S_{n-1}^{GM})^m}{S_{n-1}^{OPT}}\right)^{\frac{1}{m-1}}}{mL_{n-1}} = 1$$

Thus, $\overline{b}_j \geq 0$ for all j, $\sum_j \overline{b}_j = 1$. Hence, $\overline{b} \in \Delta^{m-1}$. Next, we verify the product constraint,

$$\prod_{j=1}^m \overline{b}_j = \frac{S_{n-1}^{OPT}\left((S_{n-1}^{GM})^m / S_{n-1}^{OPT}\right)}{(nL_{n-1})^m} = \left(\frac{S_{n-1}^{GM}}{mL_{n-1}}\right)^m$$

which satisfies the product constraint. Therefore,

$$\overline{\mathbf{b}} \in \mathcal{B}\left(\frac{S_{n-1}^{GM}}{L_{n-1}}\right)$$

By the inductive step earlier and the optimization consequence,

$$S_n^* \geq L_{n-1} \max_{\mathbf{b}\in\mathcal{B}(S_{n-1}^{GM}/L_{n-1})} \mathbf{b}\cdot\mathbf{x}(n) \geq L_{n-1}\overline{\mathbf{b}}\cdot\mathbf{x}(n)$$

Now, compute the gain of the RHS,

$$L_{n-1}\overline{\mathbf{b}}\cdot\mathbf{x}(n) = \frac{1}{m}S_{n-1}^{OPT}\max_j x_j(n) + \frac{1}{m}\left(\frac{(S_{n-1}^{GM})^m}{S_{n-1}^{OPT}}\right)^{\frac{1}{m-1}} \sum_{j\neq j_n^*} x_j(n)$$

Apply the AM-GM inequality on the $m-1$ coordinates excluding j_n^*,

$$\sum_{j\neq j_n^*} x_j(n) \geq (m-1)\left(\frac{\prod_{j=1}^m x_j(n)}{\max_j x_j(n)}\right)^{\frac{1}{m-1}}$$

Substituting this bound gives

$$L_{n-1}\overline{\mathbf{b}}\cdot\mathbf{x}(n) \geq \frac{1}{m}S_{n-1}^{OPT}\max_j x_j(n) + \frac{m-1}{m}\left[\frac{(S_{n-1}^{GM})^m}{S_{n-1}^{OPT}}\cdot\frac{\prod_{j=1}^m x_j(n)}{\max_j x_j(n)}\right]^{\frac{1}{m-1}}$$

$$= \frac{1}{m}\left[S_n^{OPT} + (m-1)\left(\frac{(S_n^{GM})^m}{S_n^{OPT}}\right)^{\frac{1}{m-1}}\right] = L_n$$

Therefore, $S_n^* \geq L_n$, completing the induction.

Thus, for every n,

$$S_n^* \geq L_n \geq \frac{1}{m}S_n^{OPT} = \frac{1}{m}\prod_{i=1}^n \max_j x_j(i)$$

This proves the claimed lower bound on consistency.

Smoothness follows by the identity $\hat{\mathbf{b}}(i)\cdot\mathbf{x}(i) = \eta_i \max_{\mathbf{b}\in\mathcal{B}(\sigma_i^*)} \mathbf{b}\cdot\mathbf{x}(i)$.

Maximal optimality follows by induction over prefixes. Fix any GM-robust policy A and any return sequence $\mathbf{x}^n$ under consistent predictions. Since A is GM-robust, Lemma A.4 implies that at every

realized history its period-i action must belong to $\mathcal{B}(\sigma_i^A)$, where $\sigma_i^A = S_{i-1}^{GM}/S_{i-1}^A$; otherwise the tightness of Lemma A.4 gives a one-step adversarial return that violates the GM guarantee.

Let S_i^* be the consistency envelope generated by MOP_{GM}, with residual $\sigma_i^* = S_{i-1}^{GM}/S_{i-1}^*$. Suppose inductively that $S_{i-1}^A \le S_{i-1}^*$. Then $\sigma_i^A \ge \sigma_i^*$. Since the feasible set $\mathcal{B}(\sigma)$ shrinks as σ increases, every action feasible for A is also feasible for the MOP oracle at residual σ_i^*. Lemma A.5 therefore implies that A's realized one-step gain at period i cannot exceed the gain chosen by the MOP oracle under consistent predictions. Multiplying by the induction hypothesis gives $S_i^A \le S_i^*$. The base case is immediate because both policies start with wealth one.

Thus, for every GM-robust policy A and every return sequence, A's consistency wealth is at most S_n^*. Since MOP_{GM} attains this value under consistent predictions, it therefore weakly dominates every GM-robust policies, and is therefore maximal-optimal. □

Lemma A.8. *Fix $\sigma \in (0,1]$, and denote set $\mathcal{B} = \mathcal{B}(\sigma) := \{\Delta^{m-1} : \prod_{j=1}^m b_j \ge (\frac{\sigma}{m})^m\}$. Let $\mathbf{b}_y \in \arg\max_{\mathbf{b}\in\mathcal{B}} \mathbf{b}\cdot\mathbf{y}$, and $\mathbf{b}_x \in \arg\max_{\mathbf{b}\in\mathcal{B}} \mathbf{b}\cdot\mathbf{x}$. Suppose there exist $\alpha > 0$, $\beta \in \mathbb{R}$, and $\epsilon \ge 0$.*

$$||\mathbf{y} - (\alpha\mathbf{x} + \beta\mathbf{1})||_\infty \le \epsilon \implies \mathbf{b}_y\cdot\mathbf{x} \ge \mathbf{b}_x\cdot\mathbf{x} - \frac{2\epsilon}{\alpha}$$

Proof. Let $\mathbf{e} := \mathbf{y} - (\alpha\mathbf{x} + \beta\mathbf{1})$, so that $||\mathbf{e}||_\infty \le \epsilon$. Since $\mathbf{b}_y$ maximizes $\mathbf{b}\cdot\mathbf{y}$ over $\mathcal{B}$, while $\mathbf{b}_x \in \mathcal{B}$, we have $\mathbf{b}_y\cdot\mathbf{y} \ge \mathbf{b}_x\cdot\mathbf{y}$. Substituting $\mathbf{y} = \alpha\mathbf{x} + \beta\mathbf{1} + \mathbf{e}$ gives

$$\mathbf{b}_y\cdot(\alpha\mathbf{x} + \beta\mathbf{1} + \mathbf{e}) \ge \mathbf{b}_x\cdot(\alpha\mathbf{x} + \beta\mathbf{1} + \mathbf{e})$$

Because every $\mathbf{b} \in \mathcal{B} \subseteq \Delta^{m-1}$ satisfies $\mathbf{b}\cdot\mathbf{1} = 1$, the β-terms cancel:

$$\alpha\mathbf{b}_y\cdot\mathbf{x} + \mathbf{b}_y\cdot\mathbf{e} \ge \alpha\mathbf{b}_x\cdot\mathbf{x} + \mathbf{b}_x\cdot\mathbf{e} \iff \alpha(\mathbf{b}_y\cdot\mathbf{x} - \mathbf{b}_x\cdot\mathbf{x}) \ge \mathbf{b}_x\cdot\mathbf{e} - \mathbf{b}_y\cdot\mathbf{e}$$

Now, for any $\mathbf{b} \in \Delta^{m-1}$,

$$|\mathbf{b}\cdot\mathbf{e}| \le \sum_{j=1}^m b_j|e_j| \le ||\mathbf{e}||_\infty \sum_{j=1}^m b_j = ||\mathbf{e}||_\infty \le \epsilon$$

Hence

$$\mathbf{b}_x\cdot\mathbf{e} - \mathbf{b}_y\cdot\mathbf{e} \ge -|\mathbf{b}_x\cdot\mathbf{e}| - |\mathbf{b}_y\cdot\mathbf{e}| \ge -2\epsilon$$

Therefore

$$\alpha(\mathbf{b}_y\cdot\mathbf{x} - \mathbf{b}_x\cdot\mathbf{x}) \ge -2\epsilon \iff \mathbf{b}_y\cdot\mathbf{x} \ge \mathbf{b}_x\cdot\mathbf{x} - \frac{2\epsilon}{\alpha}$$

since $\alpha > 0$. This proves the claim. □

Lemma A.9. *Let $m \ge 2$, $\mathbf{a} = (a_1, \cdots, a_m) \in \Delta^{m-1}$, $\mathbf{b} = (b_1, \cdots, b_m) \in \Delta^{m-1}$. Then,*

$$\sum_{j=1}^m \max\{a_j - b_j, 0\} = \sum_{j=1}^m \max\{b_j - a_j, 0\} = \frac{1}{2}||\mathbf{a} - \mathbf{b}||_1$$

Moreover, $0 \le ||\mathbf{a} - \mathbf{b}||_1 \le 2$.

Proof. Let $d_j := a_j - b_j$, $d_j^+ := \max\{d_j, 0\}$ and $d_j^- = \max\{-d_j, 0\}$. Then for each j, $d_j = d_j^+ - d_j^-$ and $|d_j| = d_j^+ + d_j^-$. Sum over j using $\sum_{j=1}^m a_j = 1 = \sum_{j=1}^m b_j$ gives $\sum_{j=1}^m d_j = \sum_{j=1}^m (a_j - b_j) = \sum_{j=1}^m a_j - \sum_{j=1}^m b_j = 0 \iff \sum_{j=1}^m d_j^+ = \sum_{j=1}^m d_j^-$. Sum over the absolutes gives $||\mathbf{a} - \mathbf{b}||_1 = \sum_{j=1}^m |d_j| = \sum_{j=1}^m (d_j^+ + d_j^-) = 2\sum_{j=1}^m d_j^+$. Thus, $\sum_{j=1}^m \max\{a_j - b_j, 0\} = \sum_{j=1}^m d_j^+ = \frac{1}{2}||\mathbf{a} - \mathbf{b}||_1$. Because $\sum_{j=1}^m d_j^+ = \sum_{j=1}^m \max\{a_j - b_j, 0\} \le 1$, we have $||\mathbf{a} - \mathbf{b}||_1 = 2\sum_{j=1}^m d_j^+ \le 2$. Lastly $||\mathbf{a} - \mathbf{b}||_1 = \sum_{j=1}^m |a_j - b_j| = \sum_{j=1}^m |b_j - a_j| = ||\mathbf{b} - \mathbf{a}||_1 \iff \sum_{j=1}^m \max\{a_j - b_j, 0\} = \sum_{j=1}^m \max\{b_j - a_j, 0\}$. $||\mathbf{a} - \mathbf{b}||_1 \ge 0$, with equality when $\mathbf{a} = \mathbf{b}$. □

Proposition A.10. *Φ is a first-order approximation of the exact self-financing model.*

Proof. Consider the exact self-financing model,

$$\rho = 1 - \kappa_s \sum_j (a_j - \rho b_j)^+ - \kappa_b \sum_j (\rho b_j - a_j)^+$$

Replace the right-hand side's ρ by 1 gives Φ. Denote the sell and buy part,

$$S(\rho) = \sum_j (a_j - \rho b_j)^+, \quad B(\rho) = \sum_j (\rho b_j - a_j)^+$$

Because the positive-part map is Lipschitz,

$$|S(\rho) - S(1)| \le 1 - \rho \quad \text{and} \quad |B(\rho) - B(1)| \le 1 - \rho$$

Therefore,

$$|\rho - \Phi| = |\kappa_s(S(\rho) - S(1)) + \kappa_b(B(\rho) - B(1))| \le (\kappa_s + \kappa_b)(1 - \rho)$$

Since,

$$1 - \rho = \kappa_s S(\rho) + \kappa_b B(\rho) \le \kappa_s + \kappa_b$$

We have,

$$|\rho - \Phi| \le (\kappa_s + \kappa_b)^2$$

Equivalently,

$$\rho = 1 - \frac{\kappa_s + \kappa_b}{2} ||a - b||_1 + O((\kappa_s + \kappa_b)^2)$$

This shows why Φ is the first-order approximation. □

Lemma A.11 (local GM-TC residual)**.** *Fix $\mathbf{a} \in \Delta^{m-1}$, $\sigma \in (0, 1]$, and $\kappa_s + \kappa_b \le \frac{1}{2}$. Let $\mathbf{b} \in \Delta^{m-1}$. Denote $||\mathbf{a} - \mathbf{b}||_1 := \sum_{j=1}^m |a_j - b_j| \in [0, 2]$, and $\Phi(\mathbf{a} \to \mathbf{b}) = 1 - \frac{\kappa_s + \kappa_b}{2} ||\mathbf{a} - \mathbf{b}||_1$. Then the following holds,*

$$\forall \mathbf{x} \in \mathbb{R}^m_{>0} : \Phi(\mathbf{a} \to \mathbf{b}) \sum_{j=1}^m b_j x_j \ge \sigma \prod_{j=1}^m x_j^{1/m} \iff \prod_{j=1}^m b_j \ge \left(\frac{\sigma}{m \Phi(\mathbf{a} \to \mathbf{b})} \right)^m$$

Remark A.12. Equivalently, We have $\forall \mathbf{x} \in \mathbb{R}^m_{>0} : \Phi(\mathbf{a} \to \mathbf{b}) \sum_{j=1}^m b_j x_j \ge \sigma \prod_{j=1}^m x_j^{1/m} \iff \mathbf{b} \in \mathcal{B}^\Phi(\sigma, \mathbf{a})$, where

$$\mathcal{B}^\Phi(\sigma, \mathbf{a}) := \left\{ \mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^m b_j \ge \left(\frac{\sigma}{m \Phi(\mathbf{a} \to \mathbf{b})} \right)^m \right\}$$

Proof. Take $\mathbf{x} = \mathbf{1}$, then $\Phi(\mathbf{a} \to \mathbf{b}) \sum_{j=1}^m b_j \ge \sigma \prod_{j=1}^m \iff \Phi(\mathbf{a} \to \mathbf{b}) \ge \sigma$, which must hold true for all $\mathbf{x} \in \mathbb{R}^m_{>0}$. Then, dividing σ by $\Phi(\mathbf{a} \to \mathbf{b})$ and applying Lemma A.2 with $\tilde{\sigma} = \sigma / \Phi(\mathbf{a} \to \mathbf{b})$ gives the claim. □

Lemma A.13 (Characterization of GM-robust portfolios under transaction costs)**.** *Fix $m \ge 2$, $\kappa_s + \kappa_b \le \frac{1}{2}$. Define, for any $\mathbf{a}, \mathbf{b} \in \Delta^{m-1}$, $\Phi(\mathbf{a} \to \mathbf{b}) := 1 - \frac{\kappa_s + \kappa_b}{2} ||\mathbf{a} - \mathbf{b}||_1$. Let A be any causal portfolio strategy that maps each history $\mathbf{x}^{i-1}$ to $\mathbf{b}(i) = A(\mathbf{x}^{i-1}) \in \Delta^{m-1}$. Initialize the pre-trade portfolio at $i = 1$ by $\mathbf{a}(1) := (1/m, \cdots, 1/m)$, $\Phi_1 := 1$. For each $i \ge 2$, define the pre-trade portfolio induced by the previous holding $\mathbf{b}(i-1)$ and returns $\mathbf{x}(i-1) \in \mathbb{R}^m_{>0}$ as*

$$\mathbf{a}(i) := \frac{\mathbf{b}(i-1) \odot \mathbf{x}(i-1)}{\mathbf{b}(i-1) \cdot \mathbf{x}(i-1)} \in \Delta^{m-1}, \qquad \Phi_i := \Phi(\mathbf{a}(i) \to \mathbf{b}(i))$$

For any return sequence $\mathbf{x}^n := \{\mathbf{x}(i)\}_{i=1}^n$ with $\mathbf{x}(i) \in \mathbb{R}^m_{>0}$, define the transaction-cost wealth of the portfolio algorithm A,

$$S_0^{TC}(\cdot, A) = 1, \qquad S_i^{TC}(\mathbf{x}^i, A) := S_{i-1}^{TC}(\mathbf{x}^{i-1}, A) \cdot \Phi_i \cdot \sum_{j=1}^m b_j(i) x_j(i)$$

and the geometric-mean benchmark $S_0^{GM}(\cdot) = 1$, $S_i^{GM}(\mathbf{x}^i) := \prod_{t=1}^{i} \prod_{j=1}^{m} x_j(t)^{1/m}$. *Then the following holds:*

$$\forall \mathbf{x}^n : S_n^{TC}(\mathbf{x}^n, A) \geq S_n^{GM}(\mathbf{x}^n)$$
$$\iff \forall i \in \{1, \cdots, n\}, \forall \mathbf{x}^{i-1} : \prod_{j=1}^{m} b_j(i) \geq \left(\frac{S_{i-1}^{GM}(\mathbf{x}^{i-1})}{m\Phi(\mathbf{a}(i) \to \mathbf{b}(i)) S_{i-1}^{TC}(\mathbf{x}^{i-1}, A)} \right)^m$$

Remark A.14. Essential properties of Lemma A.13

1. Define $\sigma_i := \frac{S_{i-1}^{GM}(\mathbf{x}^{i-1})}{S_{i-1}^{TC}(\mathbf{x}^{i-1}, A)} \in (0, 1]$, the one-step feasible set is

$$\mathcal{B}^{\Phi}(\sigma, \mathbf{a}) := \left\{ \mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^{m} b_j \geq \left(\frac{\sigma}{m\Phi(\mathbf{a} \to \mathbf{b})} \right)^m \right\}$$

   The per-round set $\mathcal{B}^{\Phi}(\sigma, \mathbf{a})$ is the largest-possible characterization of transaction-cost-aware portfolios that dominate the geometric mean benchmark. If a strategy ever picks $\mathbf{b}(i) \notin \mathcal{B}^{\Phi}(\sigma_i, \mathbf{a}(i))$ for some history, an adversary can choose a return vector $\mathbf{x}(i)$ that breaks GM-robustness immediately, and then set later returns to $\mathbf{1}$ to preserve the violation.

2. If a strategy has already outperformed GM (e.g., at period 1, if returns are not all equal) so far ($\sigma_i < 1$), then there is a slack to allow $\Phi_i < 1$ where active bets are possible. Such bets affect future feasibility by the propagation of $\mathbf{a}(i+1)$.

3. The guarantee is non-clairvoyant, and the Maximal-optimal portfolio at GM-robustness lies within feasible set $\mathcal{B}^{\Phi}(\sigma, \mathbf{a})$.

Proof. We use the one-step characterization in Lemma A.11. We first prove sufficiency. Fix an arbitrary sequence $\mathbf{x}^n$. We prove by induction on i that $S_i^{TC} \geq S_i^{GM}$ for all $i \leq n$. Base case $i = 0$ holds equality by definition. Inductive step. Assume $S_{i-1}^{TC} \geq S_{i-1}^{GM}$, then $\sigma_i = S_{i-1}^{GM} / S_{i-1}^{TC} \in (0, 1]$. By lemma A.11 and inductive hypothesis, for this history we have

$$\forall \mathbf{x} \in \mathbb{R}_{>0}^m : \Phi_i \sum_{j=1}^{m} b_j(i) x_j \geq \sigma_i \prod_{j=1}^{m} x_j^{1/m}$$

Then, apply the realized $\mathbf{x} = \mathbf{x}(i)$ and multiply both sides by S_{i-1}^{TC} gives

$$S_i^{TC} = S_{i-1}^{TC} \Phi_i \sum_{j=1}^{m} b_j(i) x_j(i) \geq S_{i-1}^{TC} \sigma_i \prod_{j=1}^{m} x_j^{1/m} = S_{i-1}^{GM} \prod_{j=1}^{m} x_j(i)^{1/m} = S_i^{GM}$$

Thus the induction closes and in particular $S_n^{TC} \geq S_n^{GM}$. Since $\mathbf{x}^n$ was arbitrary, sufficiency holds. Next, we prove necessity. First we note that global robustness implies prefix robustness:

$$\forall \mathbf{x}^n : S_n^{TC}(\mathbf{x}^n, A) \geq S_n^{GM}(\mathbf{x}^n) \implies \forall i \leq n, \forall \mathbf{x}^{i-1} : S_{i-1}^{TC}(\mathbf{x}^{i-1}, A) \geq S_{i-1}^{GM}(\mathbf{x}^{i-1})$$

hence $\sigma_i(\mathbf{x}^{i-1}) = S_{i-1}^{GM} / S_{i-1}^{TC} \in (0, 1]$ for every history (the argument can be proved in the exact same way as in Lemma A.4). Now fix any $i \in \{1, \cdots, n\}$ and any history $\mathbf{x}^{i-1}$, let $\mathbf{b}(i) = A(\mathbf{x}^{i-1})$ and let $\mathbf{a}(i)$ be the pre-trade portfolio. We claim that necessarily

$$\forall \mathbf{x} \in \mathbb{R}_{>0}^m : \Phi_i \sum_{j=1}^{m} b_j(i) x_j \geq \sigma_i \prod_{j=1}^{m} x_j^{1/m}$$

Suppose for contradiction that $\Phi_i \sum_{j=1}^{m} b_j(i) \hat{x}_j < \sigma_i \prod_{j=1}^{m} \hat{x}_j^{1/m}$ for some $\hat{\mathbf{x}} \in \mathbb{R}_{>0}^m$, then multiplying by S_{i-1}^{TC} on both sides gives the inequality

$$S_i^{TC}(\mathbf{x}^{i-1}, \hat{\mathbf{x}}, A) = S_{i-1}^{TC} \Phi_i \sum_{j=1}^{m} b_j(i) \hat{x} < S_{i-1}^{TC} \sigma_i \prod_{j=1}^{m} \hat{x}_j(i)^{1/m} = S_i^{GM}(\mathbf{x}^{i-1}, \hat{\mathbf{x}})$$

contradicting the assumption $S_n^{TC} \geq S_n^{GM}$ when extending the sequence from $i+1$ to n by setting future returns $\mathbf{x}(t) = \mathbf{1}$ for $t = \{i+1, \cdots, n\}$. Therefore $\forall \mathbf{x} \in \mathbb{R}^m_{>0} : \Phi_i \sum_{j=1}^m b_j(i) x_j \geq \sigma_i \prod_{j=1}^m x_j^{1/m}$ must hold, noting that $b_j(i) > 0 \; \forall j$. Now take $x_j := \frac{1}{b_j} \; \forall j$. Then

$$\forall \mathbf{x} \in \mathbb{R}^m_{>0} : \Phi_i \sum_{j=1}^m b_j(i) x_j \geq \sigma_i \prod_{j=1}^m x_j^{1/m} \implies \Phi_i \sum_{j=1}^m b_j \frac{1}{b_j} \geq \sigma_i \prod_{j=1}^m \frac{1}{b_j(i)^{1/m}}$$

$$\implies \prod_{j=1}^m b_j(i) \geq \left(\frac{\sigma_i}{m \Phi_i} \right)^m$$

Since i and the history were choose arbitrarily, necessity holds. Proving the full claim. □

Lemma A.15. *Fix $m \geq 2$, let $\kappa := \kappa_s + \kappa_b \in (0, \frac{1}{2}]$. For any $\mathbf{a}, \mathbf{b} \in \Delta^{m-1}$, define $\Phi(\mathbf{a}, \mathbf{b}) := 1 - \frac{\kappa}{2} ||\mathbf{a} - \mathbf{b}||_1$. Then there exists $\mathbf{q}, \mathbf{p}^*, \mathbf{b}^* \in \Delta^{m-1}$ and $\delta > 0$ with the following properties. Let $\mathcal{D}(\mathbf{q}) := \mathcal{B}^\Phi(m \prod_{j=1}^m q_j^{1/m}, \mathbf{q}) = \{\mathbf{b} \in \Delta^{m-1} : \Phi(\mathbf{q}, \mathbf{b})^m \prod_{j=1}^m b_j \geq \prod_{j=1}^m q_j\}$, choose*

$$\mathbf{b}^* \in \arg\max_{\mathbf{b} \in \mathcal{D}(\mathbf{q})} \Phi(\mathbf{q}, \mathbf{b}) b_1 \qquad \textit{and} \qquad \mathbf{p}^* := \frac{\mathbf{q} + \mathbf{b}^*}{2}$$

Then, the following holds:

1. *The first move $\mathbf{q} \to \mathbf{p}^*$ is feasible for all possible first returns: $\mathbf{p}^* \in \mathcal{D}(\mathbf{q})$*
2. *If the second return is $\mathbf{1}$, then the second move $\mathbf{p}^* \to \mathbf{b}^*$ is feasible:*

$$\mathbf{b}^* \in \mathcal{B}^\Phi \left(\frac{m \prod_{j=1}^m q_j^{1/m}}{\Phi(\mathbf{q}, \mathbf{p}^*)}, \mathbf{p}^* \right)$$

3. *The split exposure to the first asset strictly improves on the best direct feasible exposure:*

$$\Phi(\mathbf{q}, \mathbf{p}^*) \Phi(\mathbf{p}^*, \mathbf{b}^*) b_1^* = \max_{\mathbf{b} \in \mathcal{D}(\mathbf{q})} \Phi(\mathbf{q}, \mathbf{b}) b_1 + \delta$$

Proof. We first prove $\mathbf{b}^* \neq \mathbf{q}$.

Choose $a \in (0, \frac{1}{m})$ sufficiently small that $\frac{1}{a} - \frac{m-1}{1-a} - m\kappa > 0$. Then, define

$$\mathbf{q} := \left(a, \frac{1-a}{m-1}, \cdots, \frac{1-a}{m-1} \right), \qquad \sigma := m \prod_{j=1}^m q_j^{1/m} = m \left(a \left(\frac{1-a}{m-1} \right)^{m-1} \right)^{1/m}$$

Then $\mathbf{q} \in \Delta^{m-1}$ and $q_1 = a < \frac{1}{m}$. Since $\mathbf{q} \in \Delta^{m-1}$, by AM-GM and positivity:

$$0 < m \prod_{j=1}^m q_j^{1/m} \leq \sum_{j=1}^m q_j \iff 0 < \sigma \leq 1$$

Therefore, the admissible set $\mathcal{B}^\Phi$ with residual σ and pre-trade portfolio $\mathbf{q}$ is well-defined,

$$\mathcal{D}(\mathbf{q}) = \mathcal{B}(\sigma, \mathbf{q}) = \left\{ \mathbf{b} \in \Delta^{m-1} : \prod_{j=1}^m b_j^{1/m} \geq \frac{\sigma}{m \Phi(\mathbf{q}, \mathbf{b})} \right\}$$

The set $\mathcal{D}(\mathbf{q})$ is nonempty because $\mathbf{q} \in \mathcal{D}(\mathbf{q})$, compact because it is a closed subset of the simplex, and the objective $\mathbf{b} \mapsto \Phi(\mathbf{q}, \mathbf{b}) b_1$ is continuous. Therefore a maximizer $\mathbf{b}^*$ exists. For small $s > 0$, define:

$$\mathbf{z} := \left(a + s, \frac{1-a-s}{m-1}, \cdots, \frac{1-a-s}{m-1} \right)$$

Then, $||\mathbf{z} - \mathbf{q}||_1 = 2s$, and $\Phi(\mathbf{q}, \mathbf{z}) = 1 - \kappa s$. For $s = 0$

$$\Phi(\mathbf{q}, \mathbf{z})^m \prod_{j=1}^m z_j = \prod_{j=1}^m q_j$$

Moreover,

$$\frac{d}{ds}\log\left(\Phi(\mathbf{q},\mathbf{z})^m\prod_{j=1}^m z_j\right)|_{s=0}$$

$$=\frac{d}{ds}\left[m\log(1-\kappa s)+\log(a+s)+(m-1)\log\left(1-a-s\right)-(m-1)\log(m-1)\right]_{s=0}$$

$$=\left[-\frac{m\kappa}{1-\kappa s}+\frac{1}{a+\kappa s}-\frac{m-1}{1-a-s}\right]_{s=0}=\frac{1}{a}-\frac{m-1}{1-a}-m\kappa>0$$

Thus, for sufficiently small $s>0$,

$$\Phi(\mathbf{q},\mathbf{z})^m\prod_{j=1}^m z_j>\prod_{j=1}^m q_j$$

Therefore, $\mathbf{z}\in\mathcal{D}(\mathbf{q})$. Also,

$$\Phi(\mathbf{q},\mathbf{z})z_1=(1-\kappa s)(a+s)\qquad\text{and}\qquad\frac{d}{ds}[(1-\kappa s)(a+s)]\big|_{s=0}=1-\kappa a>0$$

Thus, for sufficiently small $s>0$,

$$\Phi(\mathbf{q},\mathbf{z})z_1>a=\Phi(\mathbf{q},\mathbf{q})q_1$$

Consequently,

$$\max_{\mathbf{b}\in\mathcal{D}(\mathbf{q})}\Phi(\mathbf{q},\mathbf{b})b_1>a$$

If $b_1^*\le a$, then

$$\Phi(\mathbf{q},\mathbf{b}^*)b_1^*\le b_1^*\le a$$

contradicting $\max_{\mathbf{b}\in\mathcal{D}(\mathbf{q})}\Phi(\mathbf{q},\mathbf{b})b_1>a$. Hence, $b_1^*>a$. In particular, $\mathbf{b}^*\neq\mathbf{q}$.

i) Now, we prove the first claim $\mathbf{p}^*\in\mathcal{D}(\mathbf{q})$. For $t\in[0,1]$, define

$$\mathbf{p}_t:=(1-t)\mathbf{q}+t\mathbf{b}^*,\qquad d:=\frac{\kappa}{2}||\mathbf{q}-\mathbf{b}^*||_1$$

Since $\mathbf{b}^*\neq\mathbf{q}$, $d>0$. Define

$$F(t):=m\log\Phi(\mathbf{q},\mathbf{p}_t)+\sum_{j=1}^m\log(\mathbf{p}_t)_j$$

The function is concave in t, since it is a sum of logarithms of positive affine functions. Moreover,

$$F(0)=m\log 1+\sum_{j=1}^m\log q_j\ge\log\prod_{j=1}^m q_j,\quad F(1)=m\log\Phi(\mathbf{q},\mathbf{b}^*)+\sum_{j=1}^m\log b_j^*\ge\log\prod_{j=1}^m q_j$$

where the RHS follows by $\mathbf{b}^*\in\mathcal{D}(\mathbf{q})$. By concavity,

$$F\left(\frac{1}{2}\right)\ge\frac{1}{2}F(0)+\frac{1}{2}F(1)\ge\log\prod_{j=1}^m q_j$$

Since $\mathbf{p}^*=(\mathbf{q}+\mathbf{b}^*)/2$, this gives

$$m\log\Phi(\mathbf{q},\mathbf{p}^*)+\sum_{j=1}^m\log p_j^*\ge\log\prod_{j=1}^m q_j\iff\Phi(\mathbf{q},\mathbf{p}^*)^m\prod_{j=1}^m p_j^*\ge\prod_{j=1}^m q_j$$

Therefore, $\mathbf{p}^*\in\mathcal{D}(\mathbf{q})$. This proves the first claim.

ii) Now, we prove the second claim. Suppose the intermediate return is $\mathbf{1}$. After choosing $\mathbf{p}^*$, the next pre-trade portfolio is

$$\frac{\mathbf{p}^*\odot\mathbf{1}}{\mathbf{p}^*\cdot\mathbf{1}}=\mathbf{p}^*$$

The geometric-mean is multiplied by 1, while the policy wealth is multiplied by

$$\Phi(\mathbf{q}, \mathbf{p}^*)(\mathbf{p}^* \cdot \mathbf{1}) = \Phi(\mathbf{q}, \mathbf{p}^*)$$

Therefore the next residual becomes

$$\sigma' = \frac{\sigma}{\Phi(\mathbf{q}, \mathbf{p}^*)}$$

So the second move on $\mathbf{p}^* \to \mathbf{b}^*$ is feasible at the next state if and only if

$$\mathbf{b}^* \in \mathcal{B}^\Phi(\sigma', \mathbf{p}^*) \iff \Phi(\mathbf{p}^*, \mathbf{b}^*)^m \prod_{j=1}^m b_j^* \geq \left(\frac{\sigma'}{m}\right)^m$$

$$\iff \Phi(\mathbf{q}, \mathbf{p}^*)^m \Phi(\mathbf{p}^*, \mathbf{b}^*)^m \prod_{j=1}^m b_j^* \geq \prod_{j=1}^m q_j$$

We prove the last inequality above to show that $\mathbf{p}^* \to \mathbf{b}^*$ is feasible. Because $\mathbf{p}^*$ is the midpoint between $\mathbf{q}$ and $\mathbf{b}^*$,

$$||\mathbf{q} - \mathbf{p}^*||_1 = \frac{1}{2}||\mathbf{q} - \mathbf{b}^*||_1, \qquad ||\mathbf{p}^* - \mathbf{b}^*||_1 = \frac{1}{2}||\mathbf{q} - \mathbf{b}^*||_1$$

Therefore,

$$\Phi(\mathbf{q}, \mathbf{p}^*) = 1 - \frac{d}{2}, \quad \Phi(\mathbf{p}^*, \mathbf{b}^*) = 1 - \frac{d}{2}, \quad \Phi(\mathbf{q}, \mathbf{b}^*) = 1 - d$$

So,

$$\Phi(\mathbf{q}, \mathbf{p}^*)\Phi(\mathbf{p}^*, \mathbf{b}^*) = \left(1 - \frac{d}{2}\right)^2 = 1 - d + \frac{d^2}{4} > 1 - d = \Phi(\mathbf{q}, \mathbf{b}^*)$$

Note that the above inequality is strict because $d > 0$ since $\mathbf{b}^* \neq \mathbf{q}$. Now, since $\mathbf{b}^* \in \mathcal{D}(\mathbf{q})$, we also have

$$\Phi(\mathbf{q}, \mathbf{b}^*)^m \prod_{j=1}^m b_j^* \geq \prod_{j=1}^m q_j$$

Combining the last two displays gives

$$\Phi(\mathbf{q}, \mathbf{p}^*)^m \Phi(\mathbf{p}^*, \mathbf{b}^*)^m \prod_{j=1}^m b_j^* > \Phi(\mathbf{q}, \mathbf{b}^*)^m \prod_{j=1}^m b_j^* \geq \prod_{j=1}^m q_j$$

Therefore,

$$\mathbf{b}^* \in \mathcal{B}^\Phi\left(\frac{\sigma}{\Phi(\mathbf{q}, \mathbf{p}^*)}, \mathbf{p}^*\right)$$

This proves the second claim.

iii) Now, we prove the last claim. Reusing the previous strict inequality gives

$$\Phi(\mathbf{q}, \mathbf{p}^*)\Phi(\mathbf{p}^*, \mathbf{b}^*)b_1^* > \Phi(\mathbf{q}, \mathbf{b}^*)b_1^* = \max_{\mathbf{b} \in \mathcal{D}(\mathbf{q})} \Phi(\mathbf{q}, \mathbf{b})b_1$$

Define

$$\delta := [\Phi(\mathbf{q}, \mathbf{p}^*)\Phi(\mathbf{p}^*, \mathbf{b}^*) - \Phi(\mathbf{q}, \mathbf{b}^*)]b_1^*$$

We have

$$\delta = \frac{d^2}{4} b_1^* > 0$$

since $d > 0$ and $b_1^* > 0$. Therefore

$$\Phi(\mathbf{q}, \mathbf{p}^*)\Phi(\mathbf{p}^*, \mathbf{b}^*)b_1^* = \max_{\mathbf{b} \in \mathcal{D}(\mathbf{q})} \Phi(\mathbf{q}, \mathbf{b})b_1 + \delta > \max_{\mathbf{b} \in \mathcal{D}(\mathbf{q})} \Phi(\mathbf{q}, \mathbf{b})b_1$$

This proves the lemma. □

Theorem A.16 (Non-existence of a maximal-optimality under positive transaction costs)**.** *Assume exact predictions $\mathbf{y}(i) = \mathbf{x}(i)$ are perfect predictions. Fix $m \geq 2$, $n \geq 3$, and any positive transaction-cost level $\kappa := \kappa_s + \kappa_b \in (0, \frac{1}{2}]$. Let $\mathcal{B}^{\Phi}_{GM}(m,n)$ be the class of transaction-cost GM-robust online policies, and denote the transaction-cost consistency value $C^{\Phi}_{A}(\mathbf{x}^n)$. Then there is no policy $A \in \mathcal{B}^{\Phi}_{GM}(m,n)$ such that*

$$\forall B \in \mathcal{B}^{\Phi}_{GM}(m,n), \forall \mathbf{x}^n \in (\mathbb{R}^m_{>0})^n : \quad C^{\Phi}_{A}(\mathbf{x}^n) \geq C^{\Phi}_{B}(\mathbf{x}^n)$$

Equivalently, maximal-optimality does not exist under any positive proportional transaction cost.

Proof. Maximal-optimality is a pathwise universal statement:

$$\forall B, \forall \mathbf{x}^n : C^{\Phi}_{A}(\mathbf{x}^n) \geq C^{\Phi}_{B}(\mathbf{x}^n)$$

So to refute it, it suffices to construct one robust comparison policy B and one deterministic return sequence $\mathbf{x}^n$ on which B has strictly larger consistency than A.

Suppose, for contradiction, that a maximal-optimal policy A exists.

Let $\mathbf{q} = (a, \frac{1-a}{m-1}, \cdots, \frac{1-a}{m-1}), \mathbf{p}^*, \mathbf{b}^*, \delta$ be given by Lemma A.15. Choose $M > \max\{1, \frac{1}{\delta}\}$. We will construct two return sequences $\mathbf{x}^{+,n}$ and $\mathbf{x}^{-,n}$. They agree through period 2, and differ only at period 3. Set

$$\mathbf{x}^{+}(1) = \mathbf{x}^{-}(1) := m\mathbf{q}, \qquad \mathbf{x}^{+}(2) = \mathbf{x}^{-}(2) := \mathbf{1}$$

At period 3, define

$$\mathbf{x}^{+}(3) := (M, 1, \cdots, 1), \qquad \mathbf{x}^{-}(3) := (1, M, \cdots, M)$$

For every $i = 4, \cdots, n$, set $\mathbf{x}^{+}(i) = \mathbf{x}^{-}(i) := 1$.

i) [Maximality against no-trade forces A not to trade at period 2]. Define the uniform buy-and-hold policy N, such that it chooses the uniform portfolio at period 1, and after that always chooses the current pre-trade portfolio (no rebalancing). By the application of AM-GM, N is always GM-robust under transaction costs: $N \in \mathcal{B}^{\Phi}_{GM}(m,n)$. Moreover, N ignores any prediction, so for every $\mathbf{x}^n$, $C^{\Phi}_{N}(\mathbf{x}^n) = S^{\Phi}_{N}(\mathbf{x}^n)$. On the sequence $\mathbf{x}^{-,n}$, the policy N holds $\mathbf{q}$ through period 3. Therefore,

$$\begin{aligned} C^{\Phi}_{N}(\mathbf{x}^{-,n}) = S^{\Phi}_{N}(\mathbf{x}^{-,n}) &= (\mathbf{u} \cdot m\mathbf{q})(\mathbf{q} \cdot 1)(\mathbf{q} \cdot (1, M, \cdots, M))(1) \cdots \\ &= \mathbf{q} \cdot (1, M, \cdots, M) = a + M(1-a) \end{aligned}$$

Since A is maximal-optimal and N is robust, it follows

$$S^{\Phi}_{A}(\mathbf{x}^{-,n}, \mathbf{x}^{-,n}) \geq C^{\Phi}_{A}(\mathbf{x}^{-,n}) \geq C^{\Phi}_{N}(\mathbf{x}^{-,n}) = a + M(1-a)$$

Let $\mathbf{p}$ be A's period-2 portfolio on this exact run, and let $\mathbf{b}$ be its period-3 portfolio. Since period-2 return is $\mathbf{1}$, the pre-trade portfolio at period 3 is also $\mathbf{p}$. After period 2, the residual for period 3 is

$$\sigma_3 = \frac{m \prod_{j=1}^m q_j^{1/m}}{1} \cdot \frac{1}{\Phi(\mathbf{q}, \mathbf{p})} = \frac{m \prod_{j=1}^m q_j^{1/m}}{\Phi(\mathbf{q}, \mathbf{p})}$$

Then, the local feasibility of A's period-3 action enforces

$$\prod_{j=1}^m b_j \geq \left(\frac{m \prod_{j=1}^m q_j^{1/m}}{m\Phi(\mathbf{q}, \mathbf{p})\Phi(\mathbf{p}, \mathbf{b})} \right)^m \iff \Phi(\mathbf{q}, \mathbf{p})^m \Phi(\mathbf{p}, \mathbf{b})^m \prod_{j=1}^m b_j \geq \prod_{j=1}^m q_j$$

Since $\Phi(\mathbf{q}, \mathbf{p})^m \Phi(\mathbf{p}, \mathbf{b})^m \leq 1$, this implies that

$$\prod_{j=1}^m b_j \geq \prod_{j=1}^m q_j$$

We now claim that $b_1 \geq q_1 = a$. Fix $b_1 = t$. For fixed first coordinate t, AM-GM over the remaining $m-1$ coordinates implies

$$b_1 \prod_{j \neq 1}^m b_j \leq t \left(\frac{1}{m-1} \sum_{j \neq 1}^m b_j \right)^{m-1} = t \left(\frac{1-t}{m-1} \right)^{m-1} := P(t)$$

For $t \in (0, \frac{1}{m})$,

$$\frac{P'(t)}{P(t)} = \frac{\left(\frac{1-t}{m-1}\right)^{m-1} + t(m-1)\left(\frac{1-t}{m-1}\right)^{m-2}(-\frac{1}{m-1})}{t\left(\frac{1-t}{m-1}\right)^{m-1}} = \frac{1}{t} - \frac{m-1}{1-t}$$

$$= \frac{1-mt}{t(1-t)} > 0$$

Since $P(t)$ is strictly positive, $P(t)$ is strictly increasing on $(0, \frac{1}{m})$. Since $a < \frac{1}{m}$,

$$P(a) = a\left(\frac{1-a}{m-1}\right)^{m-1} = \prod_{j=1}^{m} q_j$$

If $b_1 < a$, then

$$\prod_{j=1}^{m} b_j \le P(b_1) < P(a) = \prod_{j=1}^{m} q_j$$

contradicting feasibility. Hence $b_1 > a$. Now, upper bound the wealth of A on $\mathbf{x}^{-,n}$. At period 3, $\mathbf{x}^-(3) = (1, M, \cdots, M)$. So the period-3 gross portfolio return is $\mathbf{b} \cdot \mathbf{x}^-(3) = b_1 + M(1-b_1)$. Any later return is at most one by the sequence construction. Therefore,

$$S_A^\Phi(\mathbf{x}^{-,n}, \mathbf{x}^{-,n}) \le \Phi(\mathbf{q}, \mathbf{p})\Phi(\mathbf{p}, \mathbf{b})\,[b_1 + M(1-b_1)]$$

Now using the fact that the shrinkage factors $\Phi \le 1$, $M > 1$, and $b_1 \ge a$ gives

$$S_A^\Phi(\mathbf{x}^{-,n}, \mathbf{x}^{-,n}) \le b_1 + M(1-b_1) \le a + M(1-a)$$

Combining with the previous optimality argument on $S_A^\Phi(\mathbf{x}^{-,n}, \mathbf{x}^{-,n}) \ge a + M(1-a)$, equality must hold throughout. In particular $\Phi(\mathbf{q}, \mathbf{p})\Phi(\mathbf{p}, \mathbf{b}) = 1$. Since $\kappa > 0$,

$$\Phi(\mathbf{r}, \mathbf{s}) = 1 \iff \mathbf{r} = \mathbf{s}$$

Therefore,

$$\Phi(\mathbf{q}, \mathbf{p}) = 1 \qquad \text{and} \qquad \Phi(\mathbf{p}, \mathbf{b}) = 1$$

which forces $\mathbf{p} = \mathbf{q}$ and $\mathbf{b} = \mathbf{q}$.

Thus maximality against no-trade policy forces A to choose no trade at period 2 on the exact run for $\mathbf{x}^{-,n}$.

ii) [Causality forces the same period-2 action on $\mathbf{x}^{+,n}$]. Now consider the run for $\mathbf{x}^{+,n}$. The two sequences $\mathbf{x}^{-,n}$ and $\mathbf{x}^{+,n}$ have identical information through the second period. A causal online policy cannot distinguish these two runs at the period-2 decision. Therefore, A's period-2 portfolio on $\mathbf{x}^{+,n}$ should be exactly $\mathbf{q}$. Now, let $\mathbf{c}$ be A's period-3 decision on $\mathbf{x}^{+,n}$. We upper bound the wealth of A. By the optimality argument:

$$\Phi(\mathbf{q}, \mathbf{c})c_1 \le \max_{\mathbf{b} \in \mathcal{B}^\Phi(m\prod_{j=1}^m q_j^{1/m}, \mathbf{q})} \Phi(\mathbf{q}, \mathbf{b})b_1$$

Since period-1 and period-2 decisions are fixed, and period-4 onward has identical returns of 1, we can focus on the return after period 3,

$$C_A^\Phi(\mathbf{x}^{+,n}) \le S_A^\Phi(\mathbf{x}^{+,n}, \mathbf{x}^{+,n}) \le \Phi(\mathbf{q}, \mathbf{c})(Mc_1 + 1 - c_1) \le M \max_{\mathbf{b} \in \mathcal{B}^\Phi(m\prod_{j=1}^m q_j^{1/m}, \mathbf{q})} \Phi(\mathbf{q}, \mathbf{b})b_1 + 1$$

iii) Now we construct a robust split-trade (following Lemma A.15) that beats A on $\mathbf{x}^{+,n}$. Consider the comparison policy P: at period 1, it chooses the uniform portfolio; at period 2 if the first realized return was $m\mathbf{q}$, it chooses $\mathbf{p}^*$, otherwise it chooses no-trade; at period 3, if the first two realized returns were $(m\mathbf{q}, \mathbf{1})$, it chooses $\mathbf{b}^*$, otherwise no-trade; from period 4 onward, always choose no-trade. P is GM-robust at every period by Lemma A.15. We calculate its gain under $\mathbf{x}^{+,n}$

$$C_P^\Phi(\mathbf{x}^{+,n}) = S_P^\Phi(\mathbf{x}^{+,n}) = \left(\sum_{j=1}^{m} q_j\right)\Phi(\mathbf{q}, \mathbf{p}^*)(\mathbf{p}^* \cdot \mathbf{1})\Phi(\mathbf{p}^*, \mathbf{b}^*)[Mb_1^* + 1 - b_1^*]$$

$$= \left(\max_{\mathbf{b} \in \mathcal{B}^{\Phi}(m \prod_{j=1}^{m} q_j^{1/m}, \mathbf{q})} \Phi(\mathbf{q}, \mathbf{b}) b_1 + \delta \right) M + (1 - b_1^*) \Phi(\mathbf{q}, \mathbf{p}^*) \Phi(\mathbf{p}^*, \mathbf{b}^*)$$

$$\geq \left(\max_{\mathbf{b} \in \mathcal{B}^{\Phi}(m \prod_{j=1}^{m} q_j^{1/m}, \mathbf{q})} \Phi(\mathbf{q}, \mathbf{b}) b_1 + \delta \right) M$$

Then, by the choice of $M > \max\{1, \frac{1}{\delta}\}$

$$C_P^{\Phi}(\mathbf{x}^{+,n}) > M \max_{\mathbf{b} \in \mathcal{B}^{\Phi}(m \prod_{j=1}^{m} q_j^{1/m}, \mathbf{q})} \Phi(\mathbf{q}, \mathbf{b}) b_1 + 1 = C_A^{\Phi}(\mathbf{x}^{+,n})$$

This contradicts the assumed maximal-optimality of A, because $P \in \mathcal{B}_{GM}^{\Phi}(m, n)$. Therefore no such maximally optimal transaction-cost GM-robust policy A exists. □

Lemma A.17. *[Separation lemma] Let $\sigma \in (0, 1]$, and let $\mathbf{a}, \mathbf{a}' \in \Delta^{m-1}$ be two distinct reachable pre-trade portfolios. Then there exists a positive return vector $\mathbf{r} \in \mathbb{R}_{>0}^m$ such that*

$$\mathbf{a} \cdot \mathbf{r} > \sup_{\mathbf{b} \in \mathcal{B}_{GM}^{\Phi}(\sigma, \mathbf{a}')} \Phi(\mathbf{a}', \mathbf{b}) \mathbf{b} \cdot \mathbf{r}$$

Proof. Define the one-step achievable post-cost exposure set from state $(\sigma, \mathbf{a}')$ as

$$\mathcal{K}(\sigma, \mathbf{a}') := \left\{ \Phi(\mathbf{a}', \mathbf{b}) \mathbf{b} : \mathbf{b} \in \mathcal{B}_{GM}^{\Phi}(\sigma, \mathbf{a}') \right\} \subset \mathbb{R}_+^m$$

The feasible set $\mathcal{B}_{GM}^{\Phi}(\sigma, \mathbf{a}')$ is compact, and the map $\mathbf{b} \mapsto \Phi(\mathbf{a}', \mathbf{b}) \mathbf{b}$ is continuous, so $\mathcal{K}(\sigma, \mathbf{a}')$ is compact. For every $\mathbf{v} \in \mathcal{K}(\sigma, \mathbf{a}')$, there exists $\mathbf{b} \in \mathcal{B}_{GM}^{\Phi}(\sigma, \mathbf{a}')$ such that $\mathbf{v} = \Phi(\mathbf{a}', \mathbf{b}) \mathbf{b}$. Since $\mathbf{b} \in \Delta^{m-1}$, $\mathbf{1} \cdot \mathbf{v} = \Phi(\mathbf{a}', \mathbf{b}) \leq 1$. Moreover, because $\kappa > 0$,

$$\Phi(\mathbf{a}', \mathbf{b}) = 1 \iff \mathbf{b} = \mathbf{a}'$$

Hence, the only vector in $\mathcal{K}(\sigma, \mathbf{a}')$ with coordinate sum equal to 1 is $\mathbf{a}'$. Therefore, the only vector in conv $\mathcal{K}(\sigma, \mathbf{a}')$ with coordinate sum 1 is also $\mathbf{a}'$. Ssince $\mathbf{a} \neq \mathbf{a}'$ and $\mathbf{a} \cdot \mathbf{1} = 1$, we have $\mathbf{a} \notin$ conv $\mathcal{K}(\sigma, \mathbf{a}')$. By the strict separation theorem, there exists $\mathbf{w} \in \mathbb{R}^m$ such that

$$\mathbf{w} \cdot \mathbf{a} > \sup_{\mathbf{v}' \in \text{conv } \mathcal{K}(\sigma, \mathbf{a}')} \mathbf{w} \cdot \mathbf{v}' \geq \sup_{\mathbf{v} \in \mathcal{K}(\sigma, \mathbf{a}')} \mathbf{w} \cdot \mathbf{v}$$

Choose $\beta > 0$ large enough so that $\mathbf{r} := \mathbf{w} + \beta \mathbf{1}$ has strictly positive coordinates. For every $\mathbf{v} \in \mathcal{K}(\sigma, \mathbf{a}')$, we have $\mathbf{v} \cdot \mathbf{1} \leq 1 = \mathbf{a} \cdot \mathbf{1}$. Therefore,

$$\mathbf{r} \cdot \mathbf{a} - \mathbf{r} \cdot \mathbf{v} = \mathbf{w} \cdot \mathbf{a} - \mathbf{w} \cdot \mathbf{v} + \beta(1 - \mathbf{v} \cdot \mathbf{1}) > 0$$

Thus,

$$\mathbf{r} \cdot \mathbf{a} > \sup_{\mathbf{v} \in \mathcal{K}(\sigma, \mathbf{a}')} \mathbf{r} \cdot \mathbf{v} \iff \mathbf{a} \cdot \mathbf{r} > \sup_{\mathbf{b} \in \mathcal{B}_{GM}^{\Phi}(\sigma, \mathbf{a}')} \Phi(\mathbf{a}', \mathbf{b}) \mathbf{b} \cdot \mathbf{r}$$

□

Proposition A.18. *Fix $m \geq 2$, $n \geq 1$. Under exact predictions $\mathbf{y}(i) = \mathbf{x}(i)$ $\forall i \in [n]$, Algorithm 2 POP_{GM}^{Φ} is Pareto-optimal. There does not exist a policy $A \in \mathcal{B}_{GM}^{\Phi}(m, n)$ such that*

$$\forall \mathbf{x}^n \in (\mathbb{R}_{>0}^m)^n : C_A^{\Phi}(\mathbf{x}^n) \geq C_{POP}^{\Phi}(\mathbf{x}^n) \qquad \textit{and} \qquad \exists \hat{\mathbf{x}}^n \in (\mathbb{R}_{>0}^m)^n : C_A^{\Phi}(\hat{\mathbf{x}}^n) > C_{POP}^{\Phi}(\hat{\mathbf{x}}^n)$$

Proof. We prove by contradiction.

Assume there exists such a policy $A \in \mathcal{B}_{GM}^{\Phi}(m, n)$ that Pareto-dominates $G = POP_{GM}^{\Phi}$, then

$$\forall \mathbf{x}^n : S_A^{\Phi}(\mathbf{x}^n, \mathbf{x}^n) \geq S_G^{\Phi}(\mathbf{x}^n, \mathbf{x}^n) \qquad \text{and} \qquad S_A^{\Phi}(\hat{\mathbf{x}}^n, \hat{\mathbf{x}}^n) > S_G^{\Phi}(\hat{\mathbf{x}}^n, \hat{\mathbf{x}}^n)$$

Run both polices on the exact prediction instance $\mathbf{y}^n = \hat{\mathbf{x}}^n$. Since their terminal wealth differs, the two policies must choose different portfolios at some period. Let i be the first period at which their chosen portfolios differ. Up to period $i - 1$, they have identical wealth, residual, and pre-trade portfolio. Denote the common period-i state by $(\sigma, \mathbf{a})$. Let $\mathbf{g}$ be the period-i portfolio chosen by G, and let $\mathbf{h}$ be the period-i portfolio chosen by A. By definition on i, $\mathbf{g} \neq \mathbf{h}$.

Because A is GM-robust, its action must be locally feasible $\mathbf{h} \in \mathcal{B}^{\Phi}_{GM}(\sigma, \mathbf{a})$. Since G is Algorithm 2 and predictions are exact, it chooses a maximizer of the realized one-step post-cost gain:

$$\mathbf{g} \in \arg \max_{\mathbf{b} \in \mathcal{B}^{\Phi}_{GM}(\sigma, \mathbf{a})} \Phi(\mathbf{a}, \mathbf{b}) \mathbf{b} \cdot \hat{\mathbf{x}}(i)$$

Therefore,

$$\Phi(\mathbf{a}, \mathbf{g}) \mathbf{g} \cdot \hat{\mathbf{x}}(i) \geq \Phi(\mathbf{a}, \mathbf{h}) \mathbf{h} \cdot \hat{\mathbf{x}}(i)$$

We now split into two cases.

i) [A loses strict one-step wealth at the first deviation.] If $\Phi(\mathbf{a}, \mathbf{g}) \mathbf{g} \cdot \hat{\mathbf{x}}(i) > \Phi(\mathbf{a}, \mathbf{h}) \mathbf{h} \cdot \hat{\mathbf{x}}(i)$ where the inequality is strict, we can always pad the remaining input sequence to all $\mathbf{1}$s, so A never recovers from being dominated. This contradicts the domination assumption on A.

ii) [The first deviation is an immediate tie]. Now suppose the equality holds,

$$\Phi(\mathbf{a}, \mathbf{g}) \mathbf{g} \cdot \hat{\mathbf{x}}(i) = \Phi(\mathbf{a}, \mathbf{h}) \mathbf{h} \cdot \hat{\mathbf{x}}(i)$$

Because terminal wealth is strictly larger for A on $\hat{\mathbf{x}}^n$, this first deviation cannot occur at $i = n$; otherwise equality at the final period would imply equal terminal wealth. Hence, $i < n$.

After observing $\hat{\mathbf{x}}(i)$, the next pre-trade portfolios for A and G are respectively

$$\mathbf{a}^+_A = \frac{\mathbf{h} \odot \hat{\mathbf{x}}(i)}{\mathbf{h} \cdot \hat{\mathbf{x}}(i)}, \qquad \mathbf{a}^+_G = \frac{\mathbf{g} \odot \hat{\mathbf{x}}(i)}{\mathbf{g} \cdot \hat{\mathbf{x}}(i)}$$

These two pre-trade portfolios must be distinct. If $\mathbf{a}^+_A = \mathbf{a}^+_G$, then since $\hat{\mathbf{x}}(i) > 0$ coordinate-wise

$$\frac{h_j}{\mathbf{h} \cdot \hat{\mathbf{x}}(i)} = \frac{g_j}{\mathbf{g} \cdot \hat{\mathbf{x}}(i)}, \forall j$$

Thus $\mathbf{h}$ is a positive scalar multiple of $\mathbf{g}$. Since both vectors lie in the simplex, this implies $\mathbf{h} = \mathbf{g}$, contradicting the definition on i. Thus,

$$\mathbf{a}^+_A \neq \mathbf{a}^+_G$$

Also, because equality hold on $\Phi(\mathbf{a}, \mathbf{g}) \mathbf{g} \cdot \hat{\mathbf{x}}(i) = \Phi(\mathbf{a}, \mathbf{h}) \mathbf{h} \cdot \hat{\mathbf{x}}(i)$, the two policies must have equal wealth after period i. Therefore, they must also have the same residual at period $i + 1$. Denote this residual by σ^+.

Now, apply the separation lemma in Lemma A.17 to the two distinct reachable states

$$(\sigma^+, \mathbf{a}^+_A) \qquad \text{and} \qquad (\sigma^+, \mathbf{a}^+_G)$$

There exists a positive return vector $\mathbf{r} \in \mathbb{R}^m_{>0}$ such that

$$\mathbf{a}^+_G \cdot \mathbf{r} > \sup_{\mathbf{b} \in \mathcal{B}^{\Phi}_{GM}(\sigma^+, \mathbf{a}^+_A)} \Phi(\mathbf{a}^+_A, \mathbf{b}) \mathbf{b} \cdot \mathbf{r}$$

By padding later return sequences with all $\mathbf{1}$ vectors, this inequality remains throughout. Therefore

$$S^{\Phi}_G > S^{\Phi}_A$$

which again contradicts the domination assumption on A.

Since both cases lead to contradictions, no robust policy A can uniformly dominate $G = POP^{\Phi}_{GM}$ and strictly improve it somewhere. Hence, POP^{Φ}_{GM} is Pareto-optimal under exact-prediction consistency. □

Algorithm 2 POP^{Φ}_{GM} (GM-robust POP with transaction costs)

Require: number of assets m, horizon n, transaction cost rates κ_s, κ_b
initialize POP^{Φ}_{GM} wealth $S(0) \leftarrow 1$, GM wealth $S^{GM}(0) \leftarrow 1$
initialize pre-trade portfolio $\mathbf{a} \leftarrow (1/m, \ldots, 1/m)$
for $i \leftarrow 1$ **to** n **do**
  Compute residual $\sigma \leftarrow S^{GM}(i-1) \,/\, S(i-1)$
  Define transaction-cost factor for any $\mathbf{b} \in \Delta^{m-1}$: $\Phi(\mathbf{a} \to \mathbf{b}) \leftarrow 1 - \frac{\kappa_s + \kappa_b}{2} \|\mathbf{a} - \mathbf{b}\|_1$
  Define feasible set $\mathcal{B}^{\Phi}(\sigma, \mathbf{a}) \leftarrow \left\{ \mathbf{b} \in \Delta^{m-1} : \Phi(\mathbf{a} \to \mathbf{b}) \geq \sigma,\ \prod_{j=1}^{m} b_j \geq \left(\frac{\sigma}{m\, \Phi(\mathbf{a} \to \mathbf{b})} \right)^m \right\}$
  Receive prediction $\mathbf{y}(i) \in \mathbb{R}^m_{>0}$
  Solve $\mathbf{b}^* \leftarrow \arg\max_{\mathbf{b} \in \mathcal{B}^{\Phi}(\sigma, \mathbf{a})} \Phi(\mathbf{a} \to \mathbf{b}) \sum_{j=1}^{m} b_j y_j(i)$
  Receive true return $\mathbf{x}(i) \in \mathbb{R}^m_{>0}$
  Update POP^{Φ}_{GM} wealth $S(i) \leftarrow S(i-1)\, \Phi(\mathbf{a} \to \mathbf{b}^*) \sum_{j=1}^{m} b^*_j x_j(i)$
  Update GM wealth $S^{GM}(i) \leftarrow S^{GM}(i-1) \prod_{j=1}^{m} x_j(i)^{1/m}$
  Update next pre-trade portfolio $\mathbf{a} \leftarrow \frac{\mathbf{b}^* \odot \mathbf{x}}{\mathbf{b}^* \cdot \mathbf{x}}$
end for
return $S(n)$

B Empirical Results

Table 1: Stock combinations for NYSE(O) and NYSE(N). The suffix (O) indicates a subset drawn from NYSE(O), and (N) indicates a subset drawn from NYSE(N). NYSE(O) contains price relatives for 36 stocks from 1962–1984. NYSE(N) contains 21 stocks from 1962–2006; following Zhang et al. [15], we restrict it to 1985–2006 to avoid temporal overlap with NYSE(O). Combinations 13–18 increase the number of assets to evaluate scalability.

ID	**Combination**
1 (O)	Iro & Kin
2 (O)	Com & Kin
3 (O)	Com & Mei
4 (O)	Luk & Kin & Arc
5 (O)	Shr & Ge & Kin
6 (O)	Gul & Esp & Pil
7 (N)	Cok & Ahp & Gm & Ge & Sc
8 (N)	Dow & Mmm & Ame & Mer & Jnj
9 (N)	Ame & Kim & Alc & Ahp & Mor & Cok & Pan & Gm
10 (N)	Mmm & Dow & Mer & Ge & Ibm & Ing & Hp & Jnj
11 (N)	Ge & Mer & Gm & Ahp & Jnj & Pan & Ame & Ing & For & Dup & Dow
12 (N)	Alc & Cok & Hp & Ibm & Kim & Mmm & Mor & Sc & Gm & Ahp & Pan
13 (O)	Alc & Dow & Shr & Mmm & Kod & Gm & Sc & Ame & Mor & Hp & Esp & Cok & Tex & Fis
14 (O)	Jnj & Mer & Gul & Pil & Kin & Com & Ing & Pan & Ahp & Gte & Ibm & Kim & For & Dup
15 (O)	Arc & For & Mob & Ahp & Sc & Ibm & Dow & Kod & Gm & Tex & Esp & Pan & Jnj & Cok & Mei & Gul
16 (O)	Iro & Ge & Mor & Fis & Mer & Sea & Dup & Shr & Hp & Luk & Kim & Kin & Exx & Ame & Gte & Mmm
17 (O)	Com & Ahp & Kod & Sc & Gm & Dow & Mob & Ame & Ibm & Shr & Exx & For & Jnj & Mer & Cok & Gul & Pan & Esp
18 (O)	Ge & Alc & Sea & Gte & Mei & Kin & Arc & Dup & Fis & Ing & Pil & Luk & Kim & Mor & Hp & Iro & Mmm & Tex

Table 2: Ticker compositions of three baskets on S&P 500 (reproduced from Zhang et al. [15]).

Basket	**Constituent tickers**
α	T, MSFT, NVDA, AMZN, V
β	CSCO, MSFT, WRB, RF, T, NFLX, SBUX, BBY, ABT, BAX
γ	CMCSA, AMP, HSIC, FSLR, FCX, DE, CE, VLO, BWA, PH, ANSS, AMZN, C, EXPE, FDX, TJX, WST, EMN, PGR, FAST, PODD, HST, ADM, NVDA, PAYX, BRO, MO, ESS, DTE, WEC

Table 3: Basket specifications for the S&P 500 experiments. Regime A spans July 2019 to May 2020; regime B covers December 2022 to December 2024. (reproduced from Zhang et al. [15])

Market Regime	**Comb.** α	**Comb.** β	**Comb.** γ
A: High volatility	Popularity-weighted	Diversified sector	Random draw
B: Recent window	5 stocks	10 stocks	30 stocks

Table 4: Consistency-robustness trade-offs on NYSE without transaction costs. Each algorithm is normalized by the offline optimum, except for the GM.

ROBUSTNESS		CONSISTENCY			
ID	GM	$\frac{\text{EGSI}}{\text{OPT}}$	$\frac{\text{UPSI}}{\text{OPT}}$	$\frac{\text{RAM}}{\text{OPT}}$	$\frac{\text{MOP}_{\text{GM}}}{\text{OPT}}$
1	6.07	1.3E−12	1.7E−04	0.500	0.500
2	14.65	1.4E−12	1.8E−04	0.500	0.500
3	34.53	2.6E−12	2.9E−04	0.500	0.500
4	6.70	7.4E−26	9.9E−10	0.333	0.408
5	5.96	1.4E−24	2.1E−09	0.333	0.408
6	19.36	5.1E−25	2.7E−09	0.333	0.393
7	12.89	9.1E−35	5.2E−20	0.200	0.295
8	22.86	8.5E−31	5.6E−19	0.200	0.286
9	17.62	6.4E−46	2.0E−37	0.125	0.216
10	18.44	5.3E−45	1.1E−36	0.125	0.217
11	16.60	1.0E−51	2.7E−48	0.091	0.180
12	14.68	7.1E−57	2.4E−51	0.091	0.188
13	11.56	7.9E−65	1.2E−60	0.071	0.155
14	10.68	1.2E−70	5.3E−59	0.071	0.140
15	12.58	1.1E−72	1.1E−64	0.063	0.132
16	9.26	5.3E−79	8.4E−70	0.062	0.122
17	13.47	1.4E−73	7.8E−70	0.056	0.133
18	8.78	7.2E−89	1.2E−81	0.056	0.111

Table 5: Performance under random predictions evaluated on NYSE dataset. In each trial, the same i.i.d. random prediction stream is fed to all methods, so differences reflect the decision rules rather than prediction quality. Over 500 Monte Carlo trials, we report the mean (μ), median ($\overline{\mu}$), and standard deviation (σ).

ROBUSTNESS		SMOOTHNESS (RANDOM PREDICTIONS)								
	GM	UPSI			RAM			MOP_{GM}		
COMB.		μ	$\overline{\mu}$	σ	μ	$\overline{\mu}$	σ	μ	$\overline{\mu}$	σ
1	6.07	53.52	40.75	43.62	71.02	13.46	275.70	71.02	13.46	275.70
2	14.65	95.02	80.43	41.52	106.46	35.93	270.13	106.46	35.93	270.12
3	34.53	84.95	77.86	20.59	95.24	51.05	144.80	95.24	51.05	144.80
4	6.70	40.83	37.23	11.29	39.61	18.08	93.77	37.21	16.83	82.44
5	5.96	30.79	27.95	9.53	36.99	15.15	164.28	35.34	13.58	151.89
6	19.36	55.65	53.39	8.92	50.05	35.19	42.26	51.41	34.68	49.34
7	12.89	21.88	21.82	0.58	22.36	19.62	11.20	22.46	19.09	12.54
8	22.86	33.56	33.50	0.61	33.95	31.42	9.08	34.19	30.46	10.80
9	17.62	30.47	30.43	0.41	31.00	28.12	13.59	31.12	27.50	14.98
10	18.44	31.90	31.86	0.38	31.24	28.95	10.02	31.13	27.98	10.64
11	16.60	28.18	28.17	0.21	28.51	26.29	8.29	28.67	25.70	9.74
12	14.68	28.29	28.28	0.23	28.68	26.02	10.85	28.65	24.77	14.39
13	11.56	22.95	22.93	0.18	23.01	21.00	7.72	22.86	20.42	9.01
14	10.68	27.65	27.62	0.34	27.82	24.49	12.48	28.04	23.28	17.17
15	12.58	26.63	26.60	0.18	26.72	24.58	9.39	27.00	23.73	13.02
16	9.26	25.82	25.79	0.27	26.67	23.19	15.96	26.72	20.83	18.94
17	13.47	26.37	26.37	0.13	26.52	24.55	7.78	26.40	24.28	8.95
18	8.78	26.43	26.42	0.23	26.39	22.20	19.73	26.20	19.95	24.33

Table 6: Consistency under transaction costs on NYSE. GM reports the robustness baseline; all other entries are terminal wealth under perfect predictions after proportional transaction costs.

Robustness		Consistency ($\kappa = 0.003$)					Consistency ($\kappa = 0.01$)				
ID	GM	EGSI	UPSI	RAM	MOP_{GM}	$\text{POP}^{\Phi}_{\text{GM}}$	EGSI	UPSI	RAM	MOP_{GM}	$\text{POP}^{\Phi}_{\text{GM}}$
1	6.07	3.32E35	3.08E42	2.83E45	2.71E45	1.15E45	2.32E20	4.14E24	2.52E26	2.26E26	2.56E29
2	14.65	1.24E31	8.83E37	7.22E40	7.88E40	4.29E40	1.63E16	1.22E20	5.71E21	7.94E21	6.58E25
3	34.53	5.55E18	9.18E24	2.27E27	2.32E27	3.80E27	1.02E7	7.68E8	1.91E9	2.08E9	1.41E15
4	6.70	1.00E28	2.23E41	2.49E48	3.14E48	3.14E47	5.58E13	3.27E20	1.15E24	1.58E24	1.79E29
5	5.96	8.04E24	1.53E37	5.75E43	6.47E43	5.82E42	3.14E11	8.28E16	4.18E19	3.91E19	1.56E25
6	19.36	2.96E18	1.29E31	3.83E37	4.63E37	6.83E36	4.19E6	9.25E11	3.99E14	5.18E14	2.23E20
7	12.89	6.76E5	3.89E16	7.28E29	9.50E29	1.29E27	3.00E1	6.86E2	3.41E3	3.34E3	5.55E7
8	22.86	1.85E4	1.67E12	1.07E24	1.39E24	2.07E21	4.55	0.341	0.005	0.005	2.92E4
9	17.62	7.65E3	1.51E10	1.26E37	1.95E37	4.50E30	4.74E1	6.72E2	5.94E8	7.02E8	7.05E7
10	18.44	1.10E4	1.06E10	1.89E36	2.92E36	5.04E29	2.61E1	9.37E1	2.39E7	2.80E7	2.33E6
11	16.60	7.23E2	2.24E5	6.96E40	1.21E41	4.28E30	2.73E1	5.88E1	1.67E11	2.17E11	2.42E5
12	14.68	3.37E3	4.88E7	7.75E46	1.43E47	2.85E36	6.39E1	5.57E2	2.92E17	4.11E17	1.51E9
13	11.56	1.23E6	2.68E9	3.25E57	6.69E57	1.78E40	1.45E3	3.20E4	2.38E27	4.34E27	1.00E7
14	10.68	1.99E13	1.39E23	4.48E71	8.20E71	5.45E51	2.34E7	8.84E12	1.28E41	2.00E41	4.79E14
15	12.58	4.62E6	1.24E13	1.15E66	2.25E66	8.23E42	3.53E3	1.64E6	1.28E36	2.14E36	1.37E7
16	9.26	5.24E13	3.38E21	1.95E80	3.58E80	3.16E55	8.43E7	2.97E12	5.23E49	8.27E49	1.93E14
17	13.47	3.14E5	2.68E8	2.19E65	4.89E65	3.02E40	1.06E3	1.13E4	1.42E35	2.71E35	1.36E6
18	8.78	3.62E12	4.85E18	3.79E88	7.31E88	1.42E60	2.99E7	1.10E11	7.08E57	1.25E58	3.36E16

Table 7: Transaction-cost ablations on S&P500 using NHITS forecasts (NeuralForecast). To generate predictions on price relatives, we forecast next-day price relatives using a one-step-ahead time-series model trained on historical price relatives. Specifically, we use the NHITS model [34] as implemented in Nixtla's NeuralForecast open library [35]. For the price relative $\mathbf{x}(i)$ on day i, define its logarithmic form $\mathbf{r}(i) = \log \mathbf{x}(i)$. For each test day i, we refit the model using only the most recent 150-days historical relatives $\{\mathbf{r}(i)\}_{i-150}^{i-1}$, reserving the last 20 days within this window for validation, and then produce a one-step-ahead forecast $\hat{\mathbf{r}}(i)$. We exponentiate the forecast to obtain predicted price relatives $\mathbf{y}(i) = \exp(\hat{\mathbf{r}}(i))$. This rolling refit protocol ensures strict online deployment without look-ahead bias. Full hyper-parameter settings, training code, and random seeds are provided in the supplementary materials to ensure transparency and reproducibility. We report terminal wealth factors (initial wealth of 1) under four friction settings; κ denotes symmetric proportional buy/sell costs (i.e., $\kappa_s = \kappa_b = \kappa$). Regimes and IDs follow Table 2 and 3. GM is the geometric-mean benchmark; uBnH is uniform buy-and-hold (no rebalancing costs); Best is the clairvoyant single best stock in hindsight; ML invests fully in the model's rank-1 prediction each period and incurs cost only upon switches (a multiplicative factor $1 - \kappa$) per switch under our turnover model); in Figure 2, ML costs are not applied for readability, so the curves better reflect signal quality rather than turnover penalties. For EGSI, UPSI, RAM and MOP_{GM}, costs are applied *ex post*: we evaluate the realized turnover penalty without altering each algorithm's internal update rule, thereby isolating the impact of transaction costs on their frictionless decisions. $\text{POP}_{\text{GM}}^{\Phi}$ (Algorithm 2) is the cost-aware variant of MOP_{GM}.

Friction	**Reg**	**ID**	**Benchmarks**				**Portfolios**				
			GM	**uBnH**	**Best**	**ML**	**EGSI**	**UPSI**	**RAM**	**MOP_{GM}**	**$\text{POP}_{\text{GM}}^{\Phi}$**
$\kappa = 0$	A	α	1.2062	1.2308	1.7334	1.2763	1.2307	1.2312	1.2385	1.2344	1.2344
		β	0.9534	0.9828	1.3588	0.7071	0.9826	0.9821	0.9788	0.9762	0.9762
		γ	0.8853	0.9248	1.7334	1.1943	0.9248	0.9255	0.9215	0.9224	0.9224
	B	α	2.4704	2.6290	8.2932	3.4376	2.6299	2.6315	2.6288	2.7041	2.7041
		β	1.3004	1.3725	3.1531	1.4163	1.3729	1.3715	1.3436	1.3485	1.3485
		γ	1.4015	1.5009	8.2932	1.2589	1.5019	1.5018	1.4864	1.4831	1.4831
$\kappa = 0.0003$	A	α	1.2062	1.2308	1.7334	1.1790	1.2300	1.2302	1.2334	1.2284	1.2230
		β	0.9534	0.9828	1.3588	0.6363	0.9819	0.9812	0.9714	0.9687	0.9790
		γ	0.8853	0.9248	1.7334	1.0772	0.9241	0.9246	0.9127	0.9135	0.9301
	B	α	2.4704	2.6290	8.2932	2.7647	2.6253	2.6242	2.5087	2.5524	2.5892
		β	1.3004	1.3725	3.1531	1.0948	1.3707	1.3685	1.3000	1.3035	1.3280
		γ	1.4015	1.5009	8.2932	0.9743	1.4993	1.4984	1.4358	1.4339	1.4906
$\kappa = 0.003$	A	α	1.2062	1.2308	1.7334	0.5767	1.2236	1.2210	1.1880	1.1761	1.2193
		β	0.9534	0.9828	1.3588	0.2452	0.9760	0.9735	0.9067	0.9034	0.9679
		γ	0.8853	0.9248	1.7334	0.4242	0.9176	0.9164	0.8375	0.8365	0.9092
	B	α	2.4704	2.6290	8.2932	0.3868	2.5846	2.5595	1.6465	1.5173	2.4744
		β	1.3004	1.3725	3.1531	0.1071	1.3512	1.3419	0.9654	0.9601	1.3114
		γ	1.4015	1.5009	8.2932	0.0964	1.4761	1.4685	1.0516	1.0580	1.5272
$\kappa = 0.03$	A	α	1.2062	1.2308	1.7334	3.62E-04	1.1614	1.1327	0.8158	0.7603	1.2357
		β	0.9534	0.9828	1.3588	1.31E-05	0.9186	0.8996	0.4545	0.4488	0.9820
		γ	0.8853	0.9248	1.7334	2.85E-05	0.8554	0.8381	0.3533	0.3462	0.9186
	B	α	2.4704	2.6290	8.2932	6.04E-10	2.2104	1.9942	0.0238	0.0080	3.1607
		β	1.3004	1.3725	3.1531	4.19E-12	1.1711	1.1029	0.0487	0.0446	1.4165
		γ	1.4015	1.5009	8.2932	4.22E-12	1.2628	1.2002	0.0461	0.0500	1.6275

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims are rigorously grounded with detailed mathematical proofs. i) We characterize the feasible set of GM-robust portfolios in Lemma 3.3, characterize the consistency envelope under GM-constraints in Lemma A.5, lower bound the exact consistency value in Theorem 3.4, and show that the proposed MOP_{GM} is maximal-optimal (attaining the optimal tradeoff) in Theorem 3.4. ii) Smoothness is provided in both Theorem A.7 and Lemma A.8. iii) Under transaction costs, we retain robustness characterization in Lemma A.13. iv) We prove that no maximal policy exists under positive costs in Theorem 4.2. v) We prove that $\text{POP}^{\Phi}_{\text{GM}}$ remains Pareto-optimal in Proposition 4.3. The empirical section corroborates with these results, and shows how the theory realizes in practice.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: This work studies the strongest useful exact pathwise frontier that is non-degenerate under arbitrary predictions, anchored at the geometric mean. Preliminary section **Why GM?** justifies this benchmark choice: exact BCRP/OPT robustness is infeasible, while exact uCRP and uniform buy-and-hold robustness are degenerate; GM is the canonical symmetric floor that preserves deterministic safety while allowing prediction-driven upside.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: Our main results are universal over all positive return sequences. Specific assumptions on consistency are fully provided in the statements. All proofs are complete, correct, and rigorously verified by the authors.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: All sources of randomness are explicitly documented and reported in the supplementary materials. The main paper, appendix, and supplementary materials provide every implementation details required to replicate the results in the empirical section.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.

(b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
(c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The datasets used in this work are open-sourced. The supplementary materials contain a Colab notebook that can be run on a single-click to reproduce our results.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: All details are summarized in the Appendix and supplementary materials.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Under synthetic predictions, we report error bars including mean, median and standard deviations to assess the robustness of our algorithm under noisy predictive signals.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: All experiments compute under 24 hours on a Tesla T4 GPU, as indicated in the supplementary materials.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: All data and model used are open and publicly available.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All third-party assets are fully credited and their licenses respected.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

Justification:

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification:

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.