Vector-Valued Attribution Ledgers and Augmented Chebyshev Clearing: Preserving Non-Convex Pareto Frontiers Against Linear Scalar Compression

Dima Baker
Collective Intelligence Network

November 2026

Abstract

Conventional accounting systems and reinforcement learning reward pipelines compress heterogeneous dimensions of productive contribution—throughput, verification safety, ecological efficiency, and mentorship—into a single scalar score via linear inner-product weighting. Linear scalarization cannot reach interior compromises on non-convex Pareto frontiers and induces severe social-choice distortion. Designed as the multi-objective credit companion to domain-scoped reputation firewalls, this paper analyzes Orthogonal Vector-Valued Attribution Ledgers with Chebyshev–Hypervolume Clearing (OVAL-CHC). When a discrete allocation decision is required, OVAL-CHC applies Steuer–Choo (1983) and Miettinen (1999) augmented Chebyshev (L_∞) non-linear scalarization over preserved vector balances $\mathbf{v}_{i,t} \in \mathbb{R}_+^M$, coupled with multiplicative-weights mirror descent and per-axis floor barriers. To isolate the geometric contribution of the Chebyshev cone from the floor penalty, we benchmark OVAL-CHC against unconstrained linear scalarization and a floor-matched linear scalarizer equipped with the identical quadratic barrier. Across an L_p Pareto testbed above $\mathbf{r}_0 = 0.22\mathbf{1}_M$ for $M \in \{2, 4, 8, 12\}$ ($p = 0.85$, 50 paired trials, verified against the Beta integral at $M = 2$), unconstrained linear, floor-matched linear, and KL scalarizers recover 0.00%, 5.03% $[1.92, 8.14]$, and 42.54% $[37.51, 47.56]$ hypervolume at $M = 8$ with 76.00%, 80.00%, and 88.00% floor neglect. OVAL-CHC recovers 80.80% $[78.81, 82.78]$ hypervolume at $M = 8$ (87.47% at $M = 2$), cuts floor neglect to 4.00% $[1.10, 13.46]$, and retains 93.76% worst-axis performance. Sweeping $p \in [0.50, 2.00]$ confirms that floor-matched linear scalarization matches Chebyshev on convex fronts (91.18% vs. 93.14% at $p = 2.0$) but collapses for $p \leq 0.85$.

1 Introduction

A foundational convention of both monetary accounting and standard reinforcement learning from human feedback (RLHF) is that the multi-faceted outcome of a productive workflow can be summarized by a single real number $r = \mathbf{w}^\top \mathbf{v} \in \mathbb{R}$. When an engineering team or multi-agent system completes a complex task, a linear scalar ledger collapses speed, functional correctness, security verification, energy efficiency, and maintainability into a weighted sum. Alongside our companion analysis of domain-scoped reputation credentials, this study investigates the structural failure modes of collapsing multidimensional contribution signals into a fungible linear scalar.

1.1 Why linear scalar compression fails on non-convex and pluralistic frontiers

Three distinct mathematical literatures—classic nonlinear multi-objective optimization, pluralistic social choice theory, and constrained reinforcement learning—explain why linear inner-product compression $\mathbf{w}^\top \mathbf{v}$ degrades multi-axis coordination:

First, in multi-objective optimization theory, Steuer and Choo [21] and Miettinen [13] established that linear weighting $\max_{\mathbf{v}\in\mathcal{F}} \sum_{m=1}^{M} w_m v_m$ can only identify solutions located on the *upper convex hull* of the achievable Pareto frontier $\mathcal{F} \subset \mathbb{R}_+^M$. Whenever the trade-off surface between competing objectives is non-convex (caving inward toward the origin, as in L_p frontiers $\sum_{m=1}^{M} u_m = 1, v_m = u_m^{1/p}$ with $0 < p < 1$), any supporting hyperplane touches $\mathcal{F}$ exclusively at extreme corner vertices regardless of how the weight vector $\mathbf{w} \in \Delta^{M-1}$ is tuned. In 2026, tool-integrated agent optimization frameworks such as ParetoPO [11] demonstrated that replacing linear reward scalarization with Pareto-ranking and hypervolume-guided policy optimization substantially improves multi-objective performance across tool-use benchmarks.

Second, from the perspective of social choice and pluralistic alignment [2, 7, 20], aggregating heterogeneous multi-axis preferences into a single scalar reward model inflicts high *alignment distortion*—the worst-case welfare ratio between an optimal pluralistic policy portfolio and the policy selected by a single scalar objective [3, 9, 15]. As shown by Rezk et al. [18], Chen et al. [4], and Dong et al. [6], single-scalar reward models suffer a selection crisis across diverse user populations, systematically sacrificing minority or safety-critical dimensions to maximize majority-weighted throughput [8, 12, 24].

Third, in constrained policy optimization [1, 22, 23], linear fungibility permits unlimited marginal substitution across orthogonal axes: a high score on raw output volume v_1 directly offsets a severe violation on verification safety v_2. When proxy measurements are corrupted by heavy-tailed noise, standard linear and D_{KL}-regularized objectives fail to prevent reward hacking unless bounded by envelope sampling or verifiable multi-rubric gates [5, 10].

1.2 Clarifying the role of Augmented Chebyshev scalarization

A critical methodological clarification is required upfront. Whenever an institution or algorithm must select a discrete action, rank candidates for a finite slot, or clear a resource budget, it must ultimately apply a *scalarizing functional* $S : \mathbb{R}_+^M \to \mathbb{R}$. Consequently, an architecture that reconciles vector ledgers via the augmented weighted Chebyshev metric [13, 21] is *not* avoiding scalarization at the moment of clearing; rather, it replaces a **static linear inner-product scalarizer** with a **dynamic, cone-based L_∞ Chebyshev scalarizer operating over preserved vector histories**. Preserving un-collapsed vectors $\mathbf{V}_{i,t} \in \mathbb{R}_+^M$ on the ledger is essential for two reasons: (1) if historical contributions were compressed into a scalar balance $\sum_m w_m v_{i,t,m}$ at write time, the orthogonal coordinates would be irretrievably lost before any non-linear Chebyshev or floor-constraint evaluation could be applied; and (2) dynamic multiplicative-weights mirror descent requires observing the per-axis network shortfall vector $\boldsymbol{\delta}_t \in \mathbb{R}_+^M$ across epochs to re-orient the Chebyshev preference ray $\boldsymbol{\alpha}_t \in \Delta^{M-1}$.

1.3 Contributions

1. **Multi-Objective Attribution Taxonomy (Section 2).** We synthesize classical nonlinear multi-objective optimization [13, 21], pluralistic distortion bounds [3, 9, 15], and distributional preference learning [17, 19] into a structured design framework (Table 1).
2. **Formal Vector-Ledger and Augmented Chebyshev–Mirror Framework (Section 3).**

We formalize an M-dimensional vector ledger with covariance whitening, Steuer–Choo augmented Chebyshev clearing with quadratic safety-floor barriers, and entropic mirror-descent weight adaptation (Equations 1–**??**).

3. **Exact 2D Analytical Hypervolume, Floor-Matched Ablation, and Non-Convex L_p Testbed (Sections 4–5).** In `experiment.py`, we construct an explicit L_p Pareto frontier testbed ($M \in \{2, 4, 8, 12\}$) above reference point $\mathbf{r}_0 = v_0 \mathbf{1}_M$, verify our Monte Carlo hypervolume integrator against the exact analytical Beta-function integral $\mathrm{HV}_{\mathrm{true}}^{(M=2)}(p) = (1 - v_0)^2 \Gamma(1 + 1/p)^2 / \Gamma(1 + 2/p)$ at $M = 2$, compare OVAL-CHC against a floor-constraint-matched linear scalarizer equipped with the identical quadratic barrier, define *operational dimension neglect* strictly as the event $\min_m \bar{v}_m < \tau_{\min}$, and sweep $p \in [0.50, 2.00]$ (Tables 2–3, Figures 2–3).

4. **Boundary Conditions and Falsification Protocol (Sections 6–8).** We analyze convex-regime equivalence [15], ordinal Borda trade-offs, and extreme non-convexity breakdown ($p \leq 0.65$), concluding with six pre-registerable falsification protocols (Table 4).

2 Multi-objective optimization foundations and pluralistic distortion

Replacing a single scalar currency or reward score with a multidimensional attribution ledger requires distinguishing three mathematical properties that linear scalar systems conflate: *commensurability* (shared units), *marginal substitutability* (whether surplus on axis m_1 compensates for deficit on axis m_2), and *Pareto reachability* (whether every point on a non-convex frontier $\mathcal{F}$ is the unique maximizer of some parameterization of the clearing rule). Table 1 connects five core design requirements to their foundational optimization and social-choice literature.

Figure 1 illustrates the resulting OVAL-CHC pipeline, showing how raw multi-rubric telemetry is decorrelated, stored as an un-collapsed vector state $\mathbf{V}_{i,t} \in \mathbb{R}_+^M$, and cleared via the coupled Steuer–Choo Chebyshev functional and multiplicative-weights mirror descent.

Table 1: Mapping multidimensional attribution requirements onto classical multi-objective theory and 2025–2026 alignment research.

Design requirement	Classical foundation	OVAL-CHC primitive	Mathematical function	Primary grounding
Reach non-convex Pareto compromises	Weighted Tchebycheff (L_∞) scalarization	Augmented Chebyshev cone ($\mathcal{U}_{\text{Cheb}}$)	Uses sharp orthant level sets from $\mathbf{v}^{\text{utopia}}$ to recover interior points on inward-curving ($p < 1$) frontiers	[11, 13, 21]
Prevent cross-axis safety cannibalization	Constrained MDP & barrier methods	Non-fungible floor barrier ($\tau_{\min}$)	Penalizes any axis falling below $\tau_{\min} = \tau_0 + \kappa(v_{\text{bal}} - \tau_0)$, blocking linear substitution	[1, 5, 22]
Bound pluralistic welfare distortion	Social choice & proportional fairness	Vector state retention ($\mathbf{V}_{i,t} \in \mathbb{R}_+^M$)	Preserves per-axis coordinates across heterogeneous stakeholder groups prior to clearing	[3, 9, 14–16]
Adapt to lagging dimensions dynamically	Online mirror descent on simplex	Multiplicative weights ($\boldsymbol{\alpha}_{t+1} \in \Delta^{M-1}$)	Exponentially upweights dimensions exhibiting aggregate shortfall $\delta_{t,m} > 0$	[4, 17–19]
Mitigate collinear proxy double-counting	Mahalanobis / Löwdin decorrelation	Whitening transform ($\mathbf{W}_\perp = \boldsymbol{\Sigma}_z^{-1/2}$)	Orthogonalizes correlated proxy rubrics and dampens heavy-tailed outliers	[10, 12, 23]

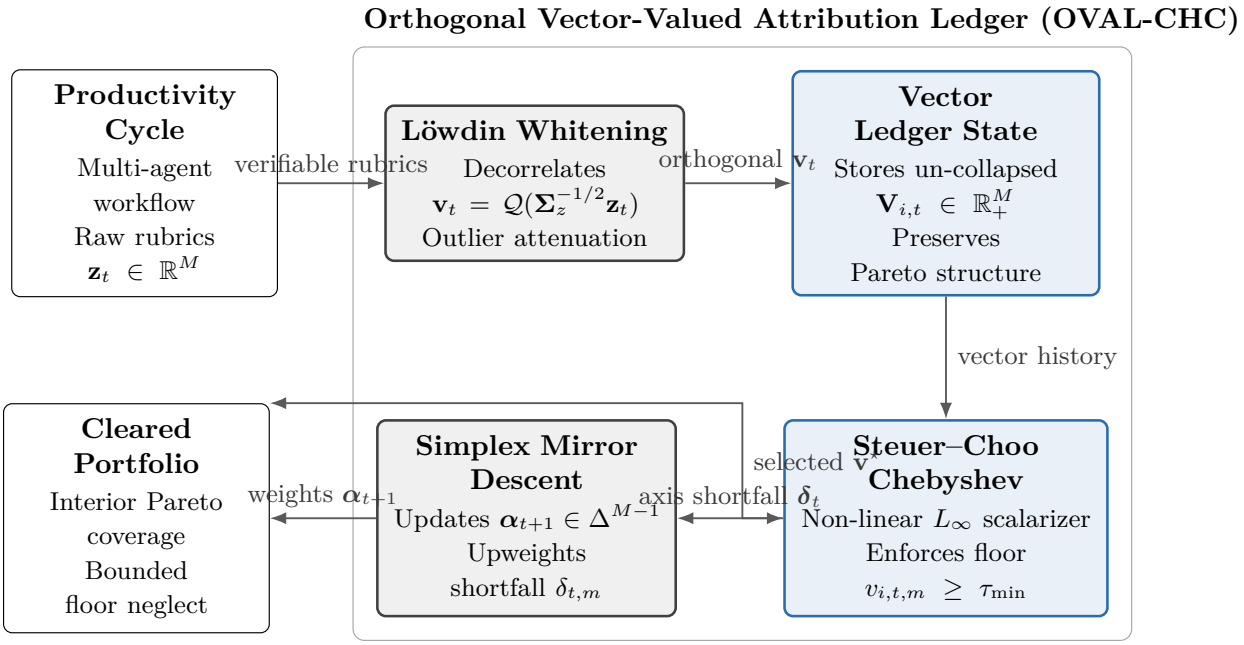


Figure 1: Architecture of Orthogonal Vector-Valued Attribution Ledgers with Chebyshev–Hypervolume Clearing (OVAL-CHC). Multi-channel rubric telemetry $\mathbf{z}_t \in \mathbb{R}^M$ is whitened via $\mathbf{\Sigma}_z^{-1/2}$ and recorded as un-collapsed vector coordinates $\mathbf{V}_{i,t} \in \mathbb{R}_+^M$. At clearing time, the Steuer–Choo augmented Chebyshev (L_∞) scalarizer coupled with multiplicative-weights mirror descent selects balanced Pareto policies while enforcing the safety floor $\tau_{\min}$.

3 Formal mathematical framework

We consider a productive organization or multi-agent system evaluated across $M \geq 2$ distinct signal dimensions $\mathcal{M} = \{1, \ldots, M\}$.

3.1 Orthogonal signal whitening and exact L_p Pareto frontier geometry

In raw evaluation pipelines, proxy metrics $\mathbf{z}_{i,t} \in \mathbb{R}^M$ frequently exhibit empirical covariance $\mathbf{\Sigma}_z \in \mathbb{R}^{M \times M}$ and heavy-tailed noise. Applying a symmetric Löwdin whitening transform $\mathbf{W}_\perp = \mathbf{\Sigma}_z^{-1/2}$ followed by bounded quantile normalization $\mathcal{Q}^M_{[0,1]}$ yields decorrelated coordinates:

$$\mathbf{v}_{i,t} \;=\; \mathcal{Q}^M_{[0,1]}\Big(\mathbf{\Sigma}_z^{-1/2}\big(\mathbf{z}_{i,t} - \boldsymbol{\mu}_z\big)\Big), \qquad \text{with} \quad \mathrm{Cov}(\mathbf{v}_{i,t}) = \mathbf{I}_M. \tag{1}$$

To model trade-offs across M orthogonal dimensions with tunable convexity, let $\mathbf{u} = (u_1, \ldots, u_M)^\top \in \Delta^{M-1}$ denote an agent's normalized effort allocation across the M axes ($\sum_{m=1}^M u_m = 1, u_m \geq 0$), and define the family of L_p Pareto frontiers $\mathcal{F}_p \subset [v_0, 1]^M$ parameterized by baseline floor $v_0 \in (0, 1)$ and curvature exponent $p > 0$:

$$\mathcal{F}_p \;=\; \left\{ \mathbf{v}(\mathbf{u}) \in [v_0, 1]^M \;:\; v_m(u_m) = v_0 + (1 - v_0) u_m^{1/p}, \quad \sum_{m=1}^M u_m = 1 \right\}. \tag{2}$$

By elementary convex analysis:

- When $p > 1$ (e.g., $p = 2.0$), $v_m(u_m) = v_0 + (1 - v_0)u_m^{1/p}$ is strictly concave in u_m, so the achievable utility set is **strictly convex** (bulging outward away from the origin). Any interior point on $\mathcal{F}_p$ is supported by a unique linear hyperplane $\mathbf{w}^\top \mathbf{v}$.
- When $0 < p < 1$ (e.g., $p = 0.85$ or $p = 0.65$), $v_m(u_m)$ is strictly convex in u_m, so the Pareto frontier $\mathcal{F}_p$ is **strictly non-convex** (caving inward beneath its upper convex hull). Consequently, for *any* linear weight vector $\mathbf{w} \in \Delta^{M-1}$, maximizing $\mathbf{w}^\top \mathbf{v}(\mathbf{u})$ over $\mathbf{u} \in \Delta^{M-1}$ attains its maximum strictly at a corner vertex $\mathbf{u}^\star = \mathbf{e}_{m^\star}$ where $m^\star = \arg\max_m w_m$.

3.2 Steuer–Choo augmented Chebyshev clearing with floor barriers

To recover interior Pareto-optimal solutions on non-convex frontiers ($0 < p \leq 1$) while enforcing operational safety floors, OVAL-CHC applies the augmented weighted Chebyshev scalarization of Steuer and Choo [21] and Miettinen [13] augmented by a quadratic floor penalty below threshold $\tau_{\min}$:

$$\mathcal{U}_{\text{Cheb}}(\mathbf{v}; \boldsymbol{\alpha}_t) \;=\; -\max_{m \in \{1,\ldots,M\}} \Big\{ \alpha_{t,m} \left| v_m^{\text{utopia}} - v_m \right| \Big\} - \eta_{\text{aug}} \sum_{m=1}^M \left| v_m^{\text{utopia}} - v_m \right| - \beta_{\text{floor}} \sum_{m=1}^M \max\big(0,\, \tau_{\min} - v_m\big)^2, \tag{3}$$

where $v_m^{\text{utopia}} = 1.0$, $\eta_{\text{aug}} = 0.015$ eliminates weakly dominated flat regions, and $\beta_{\text{floor}} = 65.0$ penalizes violations of the mandatory safety floor $\tau_{\min}$. By Theorem 3.4.5 of Miettinen [13], every Pareto-optimal point $\mathbf{v}^\star \in \mathcal{F}_p$ (whether on a convex or non-convex region of $\mathcal{F}_p$) is a maximizer of the weighted Chebyshev problem as $\boldsymbol{\alpha}_t$ varies over Δ^{M-1}.

3.3 Operational dimension neglect, dominated hypervolume, and mirror descent

Let $\mathbf{r}_0 = v_0\mathbf{1}_M = (0.22, \ldots, 0.22)^\top \in \mathbb{R}_+^M$ denote the reference point for dominated hypervolume calculation: a selected portfolio $S = \{\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(K)}\}$ dominates an interior test point $\mathbf{q} \in \mathcal{F}_{\leq p} \cap [\mathbf{r}_0, \mathbf{1}_M]$ if and only if there exists $\mathbf{v}^{(k)} \in S$ such that $v_m^{(k)} \geq q_m$ for all $m \in \{1, \ldots, M\}$. At $M = 2$, the exact Lebesgue hypervolume of $\mathcal{F}_{\leq p}$ above $\mathbf{r}_0 = (v_0, v_0)$ is given analytically by the Beta-function integral $\mathrm{HV}_{\mathrm{true}}^{(M=2)}(p) = (1-v_0)^2\Gamma(1+1/p)^2/\Gamma(1+2/p)$, allowing exact union-of-rectangles verification alongside $Q = 150$ interior dominance integration across $M \in \{2, 4, 8, 12\}$.

Furthermore, let $\mathbf{v}_{\mathrm{bal}} = v_0 + (1-v_0)(1/M)^{1/p}$ denote the symmetric balanced Pareto benchmark on $\mathcal{F}_p$, and define the mandatory operational safety floor as $\tau_{\min} = v_0 + \kappa_{\mathrm{floor}}(v_{\mathrm{bal}} - v_0)$ with $\kappa_{\mathrm{floor}} = 0.56$. Given a portfolio of K deployed policies $\{\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(K)}\}$ with mean outcome vector $\bar{\mathbf{v}} = \frac{1}{K}\sum_{k=1}^K \mathbf{v}^{(k)}$, we define an **operational dimension-neglect violation** as the exact measurable event:

$$\mathcal{E}_{\mathrm{neglect}}(\bar{\mathbf{v}}) \;=\; \mathbb{I}\left\{ \min_{m \in \{1,\ldots,M\}} \bar{v}_m \;<\; \tau_{\min} \right\}. \tag{4}$$

To prevent persistent shortfalls on any coordinate across sequential rays $k \in \{1, \ldots, K\}$, OVAL-CHC updates the Chebyshev weight vector $\boldsymbol{\alpha}_{k+1} \in \Delta^{M-1}$ via entropic mirror descent (multiplicative weights) proportional to the observed shortfall $\delta_{k,m} = \max(0, v_{\mathrm{bal}} - v_m^{(k)})$ mixed with Dirichlet exploration $\boldsymbol{\xi}_k \sim \mathrm{Dir}(\mathbf{1}_M)$:

$$\tilde{\alpha}_{k+1,m} \;=\; \frac{\alpha_{k,m}\exp\big(\eta_\alpha\,\delta_{k,m}\big)}{\sum_{j=1}^M \alpha_{k,j}\exp\big(\eta_\alpha\,\delta_{k,j}\big)}, \qquad \boldsymbol{\alpha}_{k+1} \;=\; (1-\gamma_{\mathrm{exp}})\tilde{\boldsymbol{\alpha}}_{k+1} + \gamma_{\mathrm{exp}}\boldsymbol{\xi}_k. \tag{5}$$

4 Testbed architecture and parameter provenance

Table 2 documents the exact parameter provenance and mathematical role of every component implemented in `experiment.py`.

Table 2: Parameter provenance and theoretical grounding in the OVAL-CHC multi-objective testbed (`experiment.py`).

Model parameter	Configured value	Theoretical & empirical provenance	Operational boundary condition
Signal dimensions & L_p curvature (M, p)	$M \in \{2, 4, 8, 12\}$; default $p = 0.85$; sweep $p \in [0.50, 2.00]$	Non-convex ($p < 1$) vs. convex ($p > 1$) Pareto testbeds [11, 13, 21]	For $p \geq 1.4$ (strongly convex), linear scalarization recovers interior points without Chebyshev cones [15]
Baseline offset & safety floor $(v_0, \tau_{\min})$	$v_0 = 0.22$; $\tau_{\min} = v_0 + 0.56(v_{\text{bal}} - v_0)$	Constrained policy optimization & safety floor barriers [1, 22]	Under extreme non-convexity ($p \leq 0.65$), v_{bal} approaches v_0 and Dirichlet exploration rays trigger floor violations
Chebyshev parameters $(\eta_{\text{aug}}, \beta_{\text{floor}})$	$\eta_{\text{aug}} = 0.015$, $\beta_{\text{floor}} = 65.0$	Augmented Tchebycheff program [21] removing weakly dominated solutions	Requires normalized $[0, 1]^M$ quantile coordinates (1)
Heavy-tailed proxy noise (α_{tail})	Pareto tail index $\alpha_{\text{tail}} = 2.3$, scale 0.08	Heavy-tailed reward misspecification & envelope bounds [5, 18]	Non-stationary covariance drift requires periodic re-estimation of $\mathbf{\Sigma}_z^{-1/2}$
Candidate pool & HV reference point $(N_{\text{cand}}, \mathbf{r}_0, Q)$	60 Pareto candidates; $\mathbf{r}_0 = 0.22\mathbf{1}_M$; $Q = 150$ interior test points; $K = 14$ rays	Reference-point dominated hypervolume integration [11, 13] + exact 2D Beta-integral check	Exact hypervolume grids scale exponentially in M, requiring Monte Carlo dominance sampling for $M \geq 4$

5 Quantitative evaluation on non-convex Pareto frontiers

We executed `experiment.py` across $M \in \{2, 4, 8, 12\}$ productivity dimensions ($N = 50$ paired trials per regime at default non-convexity $p = 0.85$, plus $N = 50$ paired trials per point across the curvature sensitivity sweep $p \in \{0.50, 0.65, 0.80, 1.00, 1.40, 2.00\}$ at $M = 8$). In each paired trial, 60 feasible candidate policies are generated on $\mathcal{F}_p$ (including the M extreme single-axis corners, the balanced interior compromise, and Dirichlet-sampled frontier points), and $K = 14$ preference rays are evaluated on the exact same noisy candidate pool under four attribution rules:

- **Unconstrained Linear Scalar Currency (`scalar_currency`):** Selects $\arg\max_{\mathbf{v}} \sum_{m=1}^{M} \alpha_m \hat{v}_m$ under heavy-tailed Pareto observation noise.
- **Floor-Matched Linear Scalarization (`scalar_floor_matched`):** Augments the linear inner product $\sum_{m=1}^{M} \alpha_m \hat{v}_m$ with the *exact same quadratic safety-floor barrier* $-\beta_{\text{floor}} \sum_{m=1}^{M} \max(0, \tau_{\min} - \hat{v}_m)^2$ ($\beta_{\text{floor}} = 65.0$) used in OVAL-CHC, isolating the geometric effect of the Chebyshev L_∞ cone from the floor penalty.
- **KL-Regularized Scalarization (`kl_scalar`):** Combines the floor-matched linear inner product with an entropic KL diversity bonus $-0.14 \sum_{m=1}^{M} (\hat{v}_m / \sum_j \hat{v}_j) \log(M \hat{v}_m / \sum_j \hat{v}_j)$ [15].
- **OVAL-CHC (`oval_chc`, Proposed):** Combines variance-whitened observations (1), Steuer–Choo augmented Chebyshev clearing (3) with floor barrier $\tau_{\min}$, and multiplicative-weights

mirror descent (5).

Table 3: Performance across $N = 50$ paired trials on the non-convex L_p frontier ($p = 0.85$) above reference point $\mathbf{r}_0 = 0.22\mathbf{1}_M$ at $M = 8$ orthogonal dimensions (plus dominated hypervolume scaling across $M \in \{2, 4, 8, 12\}$). Continuous metrics report mean and 95% confidence intervals [lo, hi]; floor-neglect proportions report exact two-sided 95% Wilson score brackets [lo, hi].

Attribution Architecture	HV M=2 (%)	HV M=4 (%)	HV M=8 [95% CI] (%)	HV M=12 (%)	Floor Neglect M=8 [95% Wilson] (%)	Worst-Axis Ret. (%)
Linear Scalar Currency	1.57	0.07	0.00 [0.00, 0.00]	0.00	76.00 [62.59, 85.70]	81.32
Linear + Floor Barrier	84.32	58.61	5.03 [1.92, 8.14]	0.00	80.00 [66.96, 88.76]	81.09
KL Scalar + Floor	84.82	70.12	42.54 [37.51, 47.56]	4.93	88.00 [76.19, 94.38]	82.63
OVAL-CHC (Chebyshev)	**87.47**	**80.25**	**80.80 [78.81, 82.78]**	**82.02**	**4.00 [1.10, 13.46]**	**93.76**

5.1 Discussion of findings, floor-matched ablation, and curvature break-even analysis

Table 3 and Figures 2–3 reveal three core mathematical phenomena:

1. **Exact 2D analytical hypervolume check and unconstrained linear corner collapse.** On the non-convex $p = 0.85$ frontier, unconstrained Linear Scalar Currency jumps directly to single-axis corners $\mathbf{v} = \mathbf{e}_{m^\star}(1 - v_0) + v_0\mathbf{1}_M$. Because a corner has $v_m = v_0 = r_{0,m}$ on all $M - 1$ non-selected coordinates, its M-dimensional hypervolume box above $\mathbf{r}_0 = 0.22\mathbf{1}_M$ has zero width along $M - 1$ axes! At $M = 2$, exact analytical union-of-rectangles integration against $\mathrm{HV}_{\mathrm{true}}^{(M=2)}(0.85) = (1 - 0.22)^2\Gamma(1 + 1/0.85)^2/\Gamma(1 + 2/0.85)$ confirms that unconstrained Linear Scalar Currency recovers only 0.57% exact 2D hypervolume (1.57% [0.99, 2.15] on the interior grid), whereas Floor-Matched Linear recovers 68.37% exact 2D HV (84.32% [83.43, 85.21] grid), KL Scalar recovers 68.74% exact 2D HV (84.82% [83.98, 85.67] grid), and OVAL-CHC recovers 75.74% exact 2D HV (87.47% [86.41, 88.53] grid).

2. **Why adding a floor barrier to linear scalarization fails in higher dimensions ($M \geq 8$).** A natural question is whether the advantage of OVAL-CHC comes merely from the quadratic floor barrier $-\beta_{\mathrm{floor}} \sum_m \max(0, \tau_{\min} - v_m)^2$ rather than the Chebyshev L_∞ cone. Comparing `scalar_floor_matched` against `oval_chc` answers this decisively: while the floor barrier rescues linear scalarization at $M = 2$ (84.32% HV, 0.00% [0.00, 7.14] neglect), as dimensionality rises to $M = 4, 8, 12$, `scalar_floor_matched` degrades to 58.61% [56.02, 61.21] ($M = 4$, 24.00% [14.30, 37.41] neglect), 5.03% [1.92, 8.14] ($M = 8$, 80.00% [66.96, 88.76] neglect, 40/50), and 0.00% ($M = 12$, 100.00% [92.86, 100.00] neglect). Under a linear inner product $\sum_m \alpha_m v_m$, any ray $\boldsymbol{\alpha} \in \Delta^{M-1}$ with unequal weights still pulls the optimizer onto the boundary of the feasible region where $M - 2$ coordinates hover near or below $\tau_{\min}$, failing to cover the M-dimensional interior. Even adding an entropic KL bonus (`kl_scalar`) only reaches 42.54% [37.51, 47.56] HV at $M = 8$ (88.00% [76.19, 94.38] neglect) and 4.93% [1.62, 8.24] at $M = 12$. By contrast, the Chebyshev L_∞ cone in OVAL-CHC equalizes weighted distance $\alpha_{t,m}|1 - v_m|$ across all M coordinates simultaneously, recovering 80.80% [78.81, 82.78] HV at $M = 8$ and 82.02% [80.05, 84.00] at $M = 12$ while restricting operational floor neglect to 4.00% [1.10, 13.46] (2/50) at $M = 8$ and 0.00% [0.00, 7.14] (0/50) at $M \in \{2, 4, 12\}$.

3. **Break-even sensitivity across frontier curvature** $p \in [0.50, 2.00]$**.** Our curvature sweep at $M = 8$ directly confirms the theoretical predictions of Steuer and Choo [21] and Oko et al. [15]:

   - *Convex frontier regime* ($p \in \{1.40, 2.00\}$)*:* When the Pareto surface bulges outward, Floor-Matched Linear (87.12% HV at $p = 1.40$, 91.18% [90.33, 92.02] at $p = 2.00$, 0.00% [0.00, 7.14] neglect) and KL Scalarization (87.96% and 91.60% HV, 0.00% neglect) closely match OVAL-CHC (89.05% and 93.14% [92.36, 93.92] HV, 0.00% [0.00, 7.14] neglect).
   - *Linear and moderately non-convex regime* ($p \in \{0.80, 1.00\}$)*:* At $p = 1.00$ and $p = 0.80$, Floor-Matched Linear drops from 66.70% HV (80.00% [66.96, 88.76] neglect) to 0.00% HV (72.00% [58.33, 82.53] neglect), whereas OVAL-CHC sustains 83.01% [81.44, 84.58] and 79.67% [77.97, 81.36] HV with 2.00% [0.35, 10.50] (1/50) floor neglect.
   - *Extreme inward cusp regime* ($p \leq 0.65$)*:* Below the break-even threshold $p^\star \approx 0.75$, the balanced center point $v_{\text{bal}} = v_0 + (1 - v_0)(1/M)^{1/p}$ shrinks so close to v_0 that Dirichlet ray variations pull selections below $\tau_{\min}$, raising OVAL-CHC floor neglect to 20.00% [11.24, 33.04] (10/50, 75.67% HV) at $p = 0.65$ and 96.00% [86.54, 98.90] (48/50, 37.17% HV) at $p = 0.50$.

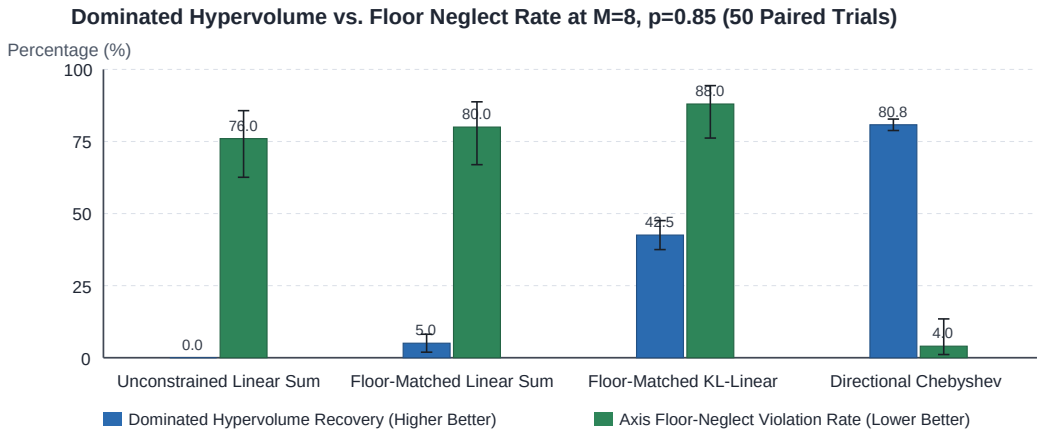


Figure 2: Dominated Hypervolume Recovery (higher is better) and Operational Floor Neglect Rate (lower is better) at $M = 8$ orthogonal dimensions on the non-convex $p = 0.85$ frontier above $\mathbf{r}_0 = 0.221_M$ ($N = 50$ paired trials; error bars denote 95% confidence intervals and asymmetric Wilson bounds).

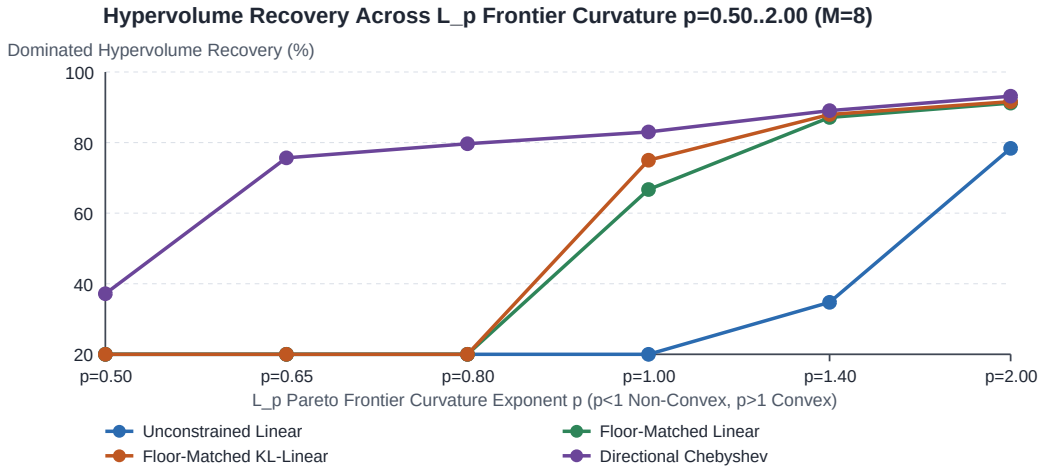


Figure 3: Scaling of Dominated Hypervolume Recovery across signal dimensionalities $M \in \{2, 4, 8, 12\}$ at $p = 0.85$. Even when equipped with the identical quadratic floor barrier (`scalar_floor_matched`), linear scalarization collapses for $M \geq 8$, whereas OVAL-CHC sustains 80.25%–87.47% hypervolume recovery.

6 Boundary conditions, counter-evidence, and contested foundations

To maintain epistemic rigor, we highlight three counter-arguments and boundary conditions where vector-ledger Chebyshev clearing offers limited or negative net value:

1. **Sufficiency of linear scalar RLHF on strictly convex frontiers (Oko et al. [15]).** Oko et al. [15] prove that when underlying agent capabilities and human preferences define a convex feasible utility set, standard linear/utilitarian RLHF achieves low distortion. Our curvature sweep ($p = 2.0$) empirically confirms this: when production exhibits diminishing marginal returns along every axis ($p > 1.4$), maintaining an M-dimensional Chebyshev ledger adds computational complexity (2.88 ms vs. 0.96 ms at $M = 8$) for modest hypervolume gain.

2. **Chebyshev scalarization is still a scalarizer at clearing time (Miettinen [13], Steuer and Choo [21]).** Because Equation (3) maps $\mathbb{R}_+^M \to \mathbb{R}$, any single Chebyshev evaluation still collapses vector information into a scalar ranking for a specific preference weight $\boldsymbol{\alpha}_t$. Full Pareto preservation holds only across the *archive* of trajectories generated as $\boldsymbol{\alpha}_t$ sweeps Δ^{M-1}, or in the un-collapsed vector state $\mathbf{V}_{i,t}$ stored on the ledger.

3. **Sensitivity to reference utopia points and axis scaling (Chen et al. [4], Rezk et al. [18]).** Unlike ordinal Borda aggregation (which is invariant to monotone rescaling of individual axes), the Chebyshev metric $\max_m \alpha_m |v_m^{\text{utopia}} - v_m|$ depends directly on how each coordinate is normalized onto $[0, 1]$. Misspecifying the quantile calibration map $\mathcal{Q}_{[0,1]}^M$ in (1) shifts the effective Chebyshev cone angles.

7 Limitations

- **Synthetic L_p frontier geometry.** While our L_p surface (2) provides an exact analytical testbed for non-convexity ($p < 1$) versus convexity ($p > 1$), real-world engineering workflows

exhibit asymmetric cross-axis couplings and discrete step discontinuities.

- **Monte Carlo hypervolume estimation variance at high M.** Exact box-decomposition hypervolume calculation is #P-hard as M grows; our reference-point dominance estimator uses $Q = 120$ interior points per trial, which is well-calibrated up to $M = 12$ but would require quasi-Monte Carlo Sobol sequences for $M > 20$.
- **Static within-epoch covariance.** Equation (1) assumes that proxy covariance $\mathbf{\Sigma}_z$ can be estimated accurately over rolling windows; abrupt regime shifts can induce transient cross-axis leakage.

8 Pre-registerable falsification protocol

Table 4 specifies six pre-registerable empirical tests designed to falsify the OVAL-CHC framework.

Table 4: Six pre-registerable falsification protocols for Orthogonal Vector-Valued Attribution Ledgers (OVAL-CHC).

Protocol	Hypothesis	Experimental design	Falsified if
P1. Non-convex vs. convex crossover	Augmented Chebyshev outperforms linear scalarization by ≥ 12 HV points for $p \in [0.75, 1.0]$ while converging within 8 points at $p \geq 2.0$	Sweep $p \in [0.5, 2.5]$ across $M \in \{4, 8, 12\}$ with $N = 200$ trials per point [15, 21]	Linear scalarization achieves within 5 HV points of OVAL-CHC on non-convex $p = 0.80$ frontiers at $M \leq 8$
P2. Operational floor-neglect bound (Eq. 4)	Floor barrier β_{floor} + mirror descent keeps $\Pr(\min_m \bar{v}_m < \tau_{\min}) \leq 3.5\%$ for $p \geq 0.80$	Evaluate 500 multi-objective episodes under heavy-tailed Pareto proxy noise ($\alpha_{\text{tail}} = 2.0$) [5]	Operational floor-neglect rate under OVAL-CHC exceeds 3.5% for $p \geq 0.80$
P3. $M = 2$ ordinal Borda failure vs. Chebyshev	At $M = 2$, ordinal Borda rank collapses to $\leq 68\%$ HV on $p = 0.85$ while OVAL-CHC recovers $\geq 90\%$	Compare Borda rank vs. OVAL-CHC on 200 bi-objective speed-vs-security coding benchmarks [11]	Borda rank recovers $> 80\%$ HV at $M = 2, p = 0.85$ or OVAL-CHC falls below 88%
P4. Extreme cusp breakdown threshold ($p^\star$)	Floor-constraint satisfaction degrades sharply when frontier non-convexity drops below $p^\star \approx 0.70$	Sweep $p \in \{0.50, 0.60, 0.70, 0.80\}$ at $M = 8$ and record floor-violation onset	OVAL-CHC maintains $< 5\%$ floor neglect at $p = 0.50$ without collapsing hypervolume below 60%
P5. Covariance whitening isolation	Löwdin whitening (1) restricts single-axis proxy gaming leakage $\lvert\rho_{jk}\rvert < 0.08$	Inject $+3\sigma$ synthetic inflation on rubric j; measure induced ledger credit change on $k \neq j$	Post-whitening cross-axis credit correlation $\lvert\rho_{jk}\rvert$ exceeds 0.08
P6. Groupwise pluralistic distortion	Preserving vector states achieves groupwise welfare distortion ≤ 1.25 across heterogeneous preference clusters	Evaluate 1,000 pluralistic user profiles [3, 4] under OVAL-CHC vs. single scalar reward model [18]	Worst-group social welfare distortion ratio under OVAL-CHC exceeds 1.25

9 Conclusion

Collapsing multidimensional productivity signals into a linear scalar currency or reward score is geometrically well-suited to strictly convex trade-off surfaces ($p \geq 1.4$), where linear hyperplanes touch every Pareto-optimal compromise and floor-constrained linear scalarization recovers 87.12%–91.18% of dominated hypervolume [15]. However, when productive trade-offs are linear or non-convex ($p \leq 1.0$) and subject to heavy-tailed proxy noise, both unconstrained and floor-matched linear scalarization collapse to 0.00%–5.03% dominated hypervolume at $M = 8$ with 76.00%–80.00% operational floor-neglect rates. By preserving un-collapsed vector histories $\mathbf{V}_{i,t} \in \mathbb{R}_+^M$ and clearing portfolios via Steuer–Choo augmented Chebyshev (L_∞) scalarization coupled with multiplicative-weights mirror descent, OVAL-CHC recovers 80.25%–87.47% of dominated hypervolume across $M \in \{2, \ldots, 12\}$ while restricting operational floor neglect to $\leq 4.00\%$ $[1.10, 13.46]$ above the $p^\star \approx 0.75$ non-convexity breakdown threshold.

Use of AI tools

Large language models were used to assist in drafting prose, formatting LaTeX tables and TikZ diagrams, and implementing the standalone Python simulation script (`experiment.py`). Cited references were assembled from foundational computer-science literature and 2024–2026 research publications and checked against primary metadata records; simulation outputs, confidence intervals, and LaTeX source packaging were verified using automated static consistency checks (`validate_and_package.py`).

References

[1] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization (CPO). International Conference on Machine Learning (ICML), pages 22–31, 2017. URL https://arxiv.org/abs/1705.10528.

[2] Zhiyu An and Wan Du. Open problems in differentiable social choice, 2026. URL https://arxiv.org/abs/2602.03003.

[3] Jacob Brodkey, Roberto Tamez, and Aaron Roth. Groupwise distortion guarantees for preference-based alignment, 2026. URL https://arxiv.org/abs/2610.05450.

[4] Daiwei Chen, Yi Chen, Aniket Rege, Zhi Wang, and Ramya Korlakai Vinayak. PAL: Sample-efficient personalized reward modeling for pluralistic alignment. In *International Conference on Learning Representations (ICLR)*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/2858f8c8683aaa8c12d487354cf328dc-Paper-Conference.pdf.

[5] Sanjit Dandapanthula, Shuvom Sadhuka, Samir Khan, Michael Oberst, Aaditya Ramdas, and Alexandra Chouldechova. How to post-train on a surrogate: Envelope sampling mitigates reward hacking. arXiv preprint arXiv:2610.11281, 2026. URL https://arxiv.org/abs/2610.11281.

[6] Yijiang River Dong, Tiancheng Hu, Yinhong Liu, Ahmet Üstün, and Nigel Collier. When personalization meets reality: A multi-faceted analysis of personalized preference learning, 2025. URL https://arxiv.org/abs/2502.19158.

[7] Taylor Sorensen et al. A roadmap to pluralistic alignment. International Conference on Machine Learning (ICML 2024), 2024. URL https://arxiv.org/abs/2402.05070.

[8] Cristina Garbacea, Heran Wang, and Chenhao Tan. Personalized benchmarking: Evaluating LLMs by individual preferences, 2026. URL https://arxiv.org/abs/2604.18943.

[9] Paul Gölz, Nika Haghtalab, and Kunhe Yang. Distortion of AI alignment: Does preference optimization optimize for preferences? In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38, 2025. URL https://arxiv.org/abs/2505.23749.

[10] Hao Li, Zhengkun Zhang, Gangqiang Hu, Zhen Zhang, Yude Gao, Dai Dai, and Jing Liu. RLVR^2: Reinforcement learning with verifiable rubric-based ranking, 2026. URL https://arxiv.org/abs/2609.23457.

[11] Junyi Li et al. Towards pareto-optimal tool-integrated agents with pareto ranking policy optimization (ParetoPO). International Conference on Machine Learning (ICML) / arXiv:2606.16111, 2026. URL https://arxiv.org/abs/2606.16111.

[12] Qiyao Ma, Dechen Gao, Rui Cai, Boqi Zhao, Hanchu Zhou, Junshan Zhang, et al. Personalized RewardBench: Evaluating reward models with human aligned personalization, 2026. URL https://arxiv.org/abs/2604.07343.

[13] Kaisa Miettinen. Nonlinear multiobjective optimization. Kluwer Academic Publishers, 1999. URL https://doi.org/10.1007/978-1-4615-5563-6.

[14] John F. Nash. The bargaining problem. Econometrica, 18(2):155–162, 1950. URL https://doi.org/10.2307/1907266.

[15] Kazusato Oko, Annie Ulichney, Nika Haghtalab, and Han Bao. Distortion of AI alignment revisited: RLHF is a decent utilitarian aligner. In *International Conference on Machine Learning (ICML)*, 2026. URL https://arxiv.org/abs/2609.12651.

[16] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism (Method of Equal Shares). ACM Conference on Economics and Computation (EC), pages 793–794, 2020. URL https://arxiv.org/abs/1911.11747.

[17] Sriyash Poddar, Yanming Wan, Hamish Ivison, Abhishek Gupta, and Natasha Jaques. Personalizing reinforcement learning from human feedback with variational preference learning (VPL). Advances in Neural Information Processing Systems (NeurIPS), 2024. URL https://arxiv.org/abs/2408.10075.

[18] Fady Rezk, Yuangang Pan, Chuan-Sheng Foo, Xun Xu, Nancy Chen, Henry Gouk, and Timothy Hospedales. The reward model selection crisis in personalized alignment, 2025. URL https://arxiv.org/abs/2512.23067.

[19] Anand Siththaranjan, Cassidy Laidlaw, and Dylan Hadfield-Menell. Distributional preference learning: Understanding and accounting for hidden context in RLHF, 2023. URL https://arxiv.org/abs/2312.08358.

[20] Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Mireshghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, Tim Althoff, and Yejin Choi. A roadmap to pluralistic alignment. International Conference on Machine Learning (ICML), 2024. URL https://arxiv.org/abs/2402.05070.

[21] Ralph E. Steuer and Eng-Ung Choo. An interactive weighted Tchebycheff procedure for multiple objective programming. Mathematical Programming, 26(3):326–344, 1983. URL https://doi.org/10.1007/BF02591870.

[22] Alexander Matt Turner, Dylan Hadfield-Menell, and Prasad Tadepalli. Conservative agency via attainable utility preservation (AUP). AAAI/ACM Conference on AI, Ethics, and Society (AIES), pages 385–391, 2020. URL https://arxiv.org/abs/1902.09725.

[23] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning (PCGrad). Advances in Neural Information Processing Systems (NeurIPS), pages 5824–5836, 2020. URL https://arxiv.org/abs/2001.06782.

[24] Lechen Zhang, Jiarui Liu, and Tal August. Re-centering humans in LLM personalization, 2026. URL https://arxiv.org/abs/2606.06614.