

01 Motivation & Contribution

Multivariate financial returns exhibit volatility clustering, asymmetric tails, and cross-asset dependence that **evolve over time**. A useful scenario generator for portfolio decisions should:

- **Capture nonlinear, asymmetric** cross-sectional dependence.
- Let the **scale of returns change over time**, as in volatility clustering.
- Support **efficient conditional sampling** for downstream decisions.

Riemannian Flow Matching (RFM; Chen & Lipman, 2024) offers a flexible nonparametric generative model, but in its unconditional form it learns a single, stationary return distribution.

CONTRIBUTION

- We extend RFM to a **Conditional Joint Riemannian Flow Matching (CJRFM)** model on $S^{d-1} \times \mathbb{R}$ that conditions the next-period return distribution on recent market history.
- CJRFM adapts **both the scale and the joint dependence structure** of returns in a fully data-driven manner.
- A head-to-head comparison against unconditional RFM — same geometry, same architecture — isolates the value of **conditioning itself**.

KEY FINDING

Although RFM and CJRFM share an identical geometric representation and flow-matching architecture, conditioning on recent market observations raises the Sharpe ratio by roughly **1.7×** — distributional flexibility alone is not enough; adapting to current conditions is what matters.

1

Nonlinear dependence
asymmetric, cross-asset tail behavior

2

Time-varying scale
volatility clustering, changing regimes

3

Efficient sampling
many scenarios per rebalance, fast ODE


Three requirements a scenario generator for portfolio decisions should satisfy.

02 Riemannian Flow Matching

A standardized return vector $z \in \mathbb{R}^d$ is mapped one-to-one onto direction $\times$ log-magnitude, $M = S^{d-1} \times \mathbb{R}$:

DECOMPOSITION MAP T

$$T(z) = (x, y), \quad x = z/||z||, \quad y = \log||z|| \quad (z = e^y x)$$



Standardized return decomposed into a spherical direction X and a Euclidean log-radius Y; the flow moves each base point smoothly toward the observed target on M.

Starting from a standard normal base Z(0) and the observed return Z(1), the direction moves along the sphere via **SLERP**; the log-radius moves linearly:

GEODESIC / LINEAR PROBABILITY PATH

$$X(t) = \frac{\sin((1-t)\varphi)}{\sin\varphi} X(0) + \frac{\sin(t\varphi)}{\sin\varphi} X(1) \\ \varphi = \arccos\langle X(0), X(1) \rangle \\ Y(t) = (1-t)Y(0) + tY(1)$$

CONDITIONAL EXTENSION

CJRFM conditions the target Z(1) = $\Xi(t_{n+1})|_{\mathcal{F}_{m:n}}$ on a lookback window. The transformed history is flattened into $C_{\text{flat}} \in \mathbb{R}^{(n-m+1)(d+1)}$ and fed directly into the shared network alongside (t, X(t), Y(t)) — no recurrent encoder is used.

The neural field v_θ additionally takes C_{flat} as input and is trained by regressing on the closed-form conditional tangent velocities at random flow times $\tau \sim U[0,1]$:


CONDITIONAL FLOW-MATCHING LOSS

$$\mathcal{L}(\theta) = \mathbb{E} \left[||v_\theta^\tau(\tau, X(\tau)|_{m:n}, Y(\tau)|_{m:n}, C_{\text{flat}}) - \frac{dX(t)|_{m:n}}{dt}||^2 + ||v_\theta^\tau(\tau, X(\tau)|_{m:n}, Y(\tau)|_{m:n}, C_{\text{flat}}) - \frac{dY(t)|_{m:n}}{dt}||^2 \right]$$

WHY GEOMETRY HELPS

Separating direction from magnitude lets the model place a probability path **directly on the sphere** (via SLERP) instead of forcing a Euclidean flow to approximate curved, high-dimensional dependence — the log-radius coordinate then carries scale/volatility information on its own axis.

03 Architecture & Backtest Design



Neural vector field: sinusoidal time embedding (Kτ=4) + 3-layer shared MLP + tangent-projected spherical head & scalar log-radius head

17-ASSET UNIVERSE

11 U.S. sector ETFs (XLB, XLE, XLF, XLI, XLK, XLP, XLRE, XLU, XLV, XLY, XLC) + SPY (broad equity), TLT (bonds), GLD (gold), USO (oil), VNQ (real estate), BIL (T-bill cash proxy).

- **Rolling window:** est. Jun 19 2018–Dec 31 2019; OOS Jan 2 2020–2024; 5-day rebalance.
- **Benchmarks:** Gaussian, NTS, DCC-GARCH, NTS-GARCH, Equal-weight, unconditional RFM.
- **Portfolio choice:** CRRA utility (γ=10) over M=200,000 scenarios, turnover-penalized, solved via SLSQP; 10 bps transaction cost.

SCENARIO-BASED CRRA OBJECTIVE

$$w^* = \arg \max_{w \in W} \frac{1}{M} \sum_{m=1}^M \frac{(1+w \cdot R^{(m)})^{1-\gamma} - 1}{1-\gamma} - c \cdot T_\ell(w, w^-)$$

TRAINING DETAILS


- Optimizer: **AdamW**, learning rate 10^{-5} , weight decay 10^{-6} .
- Full-batch gradient updates — batch size equals the entire rolling training window.
- **3,000** gradient steps per rolling window; CJRFM additionally conditions on $\ell=3$ lagged days.

04 Results & Takeaways

CJRFM ranks **first** in annualized return, Sharpe ratio, and both reward-to-tail-risk ratios among all seven strategies over the ~4-year out-of-sample window.

Strategy	AnnMean	Sharpe	CVaR1%	CVaR5%	CR1%	CR5%
Equal-weight	9.87%	0.537	5.29%	2.80%	1.867	3.530
Gaussian	8.11%	0.449	4.95%	2.71%	1.640	2.994
DCC-GARCH	7.60%	0.541	3.84%	2.32%	1.980	3.267
NTS	7.56%	0.426	4.91%	2.73%	1.539	2.768
NTS-GARCH	7.99%	0.528	4.02%	2.50%	1.987	3.188
RFM (unconditional)	8.79%	0.372	6.23%	3.47%	1.410	2.533
CJRFM (proposed)	11.80%	0.646	4.22%	2.79%	2.795	4.229

Out-of-sample rolling backtest; DCC-GARCH attains the lowest CVaR, but CJRFM gives the strongest overall return / downside-risk balance.



CJRFM attains the highest Sharpe ratio of all seven strategies — roughly 1.7× the unconditional RFM baseline.

+34%

AnnMean vs.
next-best (EW)

1.7×

Sharpe vs.
unconditional RFM

2.80 / 4.23

CR1% / CR5%
best among all

CONCLUSION

Distributional flexibility alone (unconditional RFM) is **not sufficient**. Adapting the predictive distribution to recent market information — scale *and* dependence jointly — is the key driver of stronger scenario generation and portfolio outcomes. **Future work:** richer conditioning, alternative manifolds, longer horizons, broader universes.

Chen, R.T.Q. & Lipman, Y. (2024). Flow matching on general geometries. *ICLR*.

Lipman, Y. et al. (2023). Flow matching for generative modeling. *ICLR*.

Kim, Y.S. (2022). Portfolio optimization & MCR on multivariate NTS. *Annals of Operations Research*, 312, 853–881.

Engle, R.F. (2002). Dynamic conditional correlation. *JBES*, 20(3), 339–350.

SUBMITTED Park, J. & Kim, Y.S. (2026). Conditional Joint Riemannian Flow Matching... *Working paper, Stony Brook University*.