TailDiff: Tail-Guided Diffusion for Risk-Sensitive Generation

Shuning Zhao[1*] **Patrick Wong**[2*] **Leran Zhang**[3] **Rongshang Li**[4] **Xiaolin Hu**[1,5,6†]
[1]Department of Computer Science and Technology, Institute for AI, BNRist, Tsinghua University
[2]Department of Econometrics and Business Statistics, Monash University
[3]School of Mathematics and Statistics, University of Melbourne [4]Promptlaw
[5]IDG/McGovern Institute for Brain Research, Tsinghua University
[6]Chinese Institute for Brain Research (CIBR)
zsn20@mails.tsinghua.edu.cn patrick.wong@monash.edu
leran.zhang.1@student.unimelb.edu.au
robert.li@promptlaw.com.au xlhu@tsinghua.edu.cn

Abstract

Generative models are increasingly used to simulate risk-sensitive systems, from financial markets and energy demand to healthcare, where downstream decisions depend on rare but consequential events. In this setting, matching average behaviour is not enough: rare tail events drive stress testing, capital allocation, and policy analysis. Yet tail objectives such as Value-at-Risk and Expected Shortfall are difficult to optimise directly because empirical estimators rely on sorting, thresholding, and subset selection. We introduce TAILDIFF[1], a training-free inference-time guidance method for pretrained diffusion models that uses a distributional tail scoring rule to construct differentiable tail-aware objectives. We prove that the population loss uniquely identifies the intended tail law, quantify its finite implementation gap, and bound the terminal sampler perturbation by the accumulated realised guidance updates. On synthetic data, TAILDIFF comes close to the oracle sampling error. On a challenging financial tail-risk simulation benchmark, it outperforms the state of the art on tail metrics while preserving bulk structure.

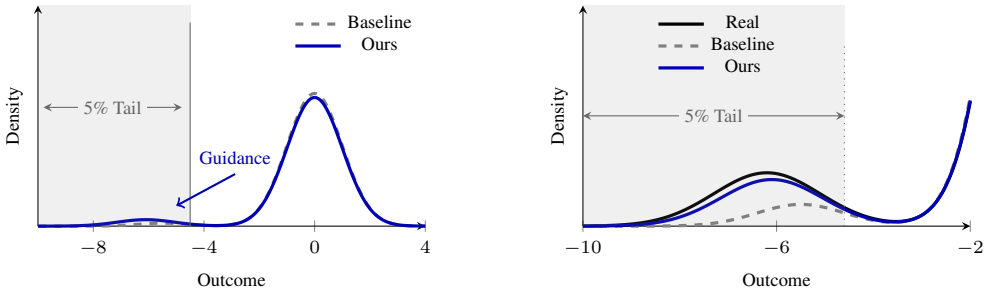
1 Introduction

Generative models address data scarcity by expanding limited observations into synthetic scenarios for learning and decision-making, including financial markets [Li et al., 2025a], healthcare [Ktena et al., 2024], and safety-critical infrastructure [Xu et al., 2025, Lin et al., 2024]. This role is especially important for time series with limited observations [Gonen et al., 2025], yet scarcity is greatest in the tail, where rare events provide the least evidence. Tail-faithful generation therefore addresses the part of the distribution for which synthetic augmentation is needed most.

Modern generators have improved aggregate fidelity through diffusion [Yuan and Qiao, 2024, Naiman et al., 2024] and population-aware modelling [Li et al., 2025b]. However, objectives and evaluations dominated by typical observations can reward bulk fidelity while missing rare events [Liu et al., 2026, Pandey et al., 2025, Huster et al., 2021]. Figure 1a illustrates the desired correction: reallocating probability mass towards underrepresented tail outcomes while preserving bulk fidelity. The need for such correction is evident in Figure 1b, where a simulator can closely match the bulk while underestimating both the frequency and severity of extreme outcomes. In risk-sensitive applications,

*Equal contribution.
†Corresponding author.
[1]Code is available at https://github.com/ShuningZhao/TailDiff/.

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



(a) **Tail-Aware Guidance** (b) **Tail Misspecification**

Figure 1: **TAILDIFF corrects extreme-event misspecification in generative models. (a)** The full distributional view: TAILDIFF selectively reshapes the model distribution by reallocating probability mass towards the tail outcomes while largely preserving the bulk. **(b)** A zoomed-in view of the left tail of the distribution in **(a)**, highlighting that a baseline diffusion model can match the bulk while underestimating both the frequency and severity of extreme events.

this produces optimistic stress tests and risk assessments precisely where data are scarcest and synthetic scenarios are most valuable [Cont et al., 2025, Galib et al., 2024]. The central challenge is therefore to correct tail misspecification without sacrificing the bulk fidelity that makes the simulator useful.

Correcting tail behaviour is operationally difficult because empirical quantiles and tail sets depend on rank selection and hard thresholds, yielding sparse active-set gradients that switch when sample ranks change [He et al., 2023, Hu et al., 2024, Xie et al., 2025]. Our method TAILDIFF addresses these obstacles by targeting the projected tail law, the complete distribution within the worst α fraction of outcomes, rather than a small set of tail summaries. Following Fissler et al. [2026], we construct a differentiable proper score for this target, and our theory shows that its population optimum identifies the intended tail law. TAILDIFF uses this score as a batch-level energy to guide a pretrained diffusion model during sampling, correcting tail misspecification without retraining its backbone or distorting the bulk, as illustrated in Figure 1(a).

We evaluate TAILDIFF under matched finite budgets on controlled synthetic processes and financial market data. Finance serves as our primary benchmark because it provides a well-established framework for evaluating tail-risk. Value-at-Risk (VaR) and Expected Shortfall (ES) are regulator-defined measures used in risk reporting, backtesting, and capital requirements [Basel Committee on Banking Supervision, 2006, 2019]. We evaluate their accuracy alongside preservation of bulk structure. Our theory validates the tail-law objective, connects it to the implemented guidance, and bounds its effect on the pretrained sampler.

Our contribution is threefold:

- **Method:** We introduce TAILDIFF, an inference-time method that guides a pretrained diffusion model through a differentiable score for the complete projected adverse-tail distribution.
- **Theory:** We prove that the population loss uniquely identifies the intended tail law, quantify the finite implementation gap, and bound the resulting sampler intervention.
- **Results:** On established risk-sensitive generation benchmarks, TAILDIFF outperforms the current state of the art on tail-risk measures while preserving bulk structure, and on the synthetic benchmark it recovers VaR and ES almost as accurately as the same number of independent draws from the true law. Additional cross-domain experiments assess portability beyond finance.

2 Related Work

2.1 General-Purpose Generative Models

Generative time-series models have progressed from adversarial sequence generation, such as TimeGAN [Yoon et al., 2019] and Quant GAN [Wiese et al., 2020], to score-based modelling [Song and Ermon, 2019] and denoising diffusion [Ho et al., 2020]. Diffusion models now support general-

purpose multivariate synthesis [Yuan and Qiao, 2024, Naiman et al., 2024, Li et al., 2025b], probabilistic forecasting [Rasul et al., 2021, Liu et al., 2024], imputation [Tashiro et al., 2021], and constrained or self-guided generation [Coletta et al., 2023, Kollovieh et al., 2023]. Financial-market simulators similarly span controlled adversarial models [Xia et al., 2024] and diffusion architectures [Takahashi and Mizuno, 2025, Tanaka et al., 2025].

Generative modelling has also been developed for few-shot anomalies [Jin et al., 2025] and class-imbalanced financial and fraud data [Altman et al., 2023, Yang et al., 2025, Park and Ko, 2026]. However, these methods are typically trained and evaluated using objectives aggregated over the full empirical distribution. Since common observations contribute most of this signal, strong overall fidelity does not by itself establish fidelity in the scarce projected tails that govern downstream risk.

2.2 Tail-Aware and Extreme-Aware Generation

Earlier generators encode Pareto marginals or block maxima in observation space [Huster et al., 2021, Galib et al., 2024]. Recent works such as SHD [Liu et al., 2026] and t-EDM/t-Flow [Pandey et al., 2025] introduce Student-t dynamics. Together, these methods expand a generator's capacity for heavy-tailed data. Our closest comparator TailGAN [Cont et al., 2025] trains a GAN with a statistically consistent score for the jointly elicitable VaR-ES pair. However, its generator objective depends on the generated law only through these two point functionals, so it cannot distinguish projected tail laws that share the same target. TAILDIFF's proper score instead identifies the complete projected tail law at its population optimum. DSI [Zhao et al., 2026] pools checkpoint samples to improve risk estimation, whereas TAILDIFF corrects the generated law itself.

2.3 Diffusion Guidance

Guidance methods modify the reverse diffusion process to favour outputs satisfying a desired condition. Classifier guidance [Dhariwal and Nichol, 2021] obtains the required direction from the gradient of a separately trained classifier. Classifier-free guidance [Ho and Salimans, 2022] removes this external classifier by combining conditional and unconditional predictions learned by the diffusion model. More general training-free methods instead differentiate an external objective through the current estimate of the clean sample. Universal guidance [Bansal et al., 2024] allows this objective to represent a broad range of differentiable conditions. FreeDoM [Yu et al., 2023] expresses the desired condition as a scalar energy, where lower values represent preferred outputs, and uses its gradient to steer the reverse process. Compositional energy-based diffusion [Du et al., 2023] provides a related route through MCMC corrections.

TAILDIFF belongs to this final family, but replaces the usual sample-wise condition with a batch-level distributional energy. Derived from a proper score for the projected tail law, this energy has a precise statistical target: its population optimum identifies the intended adverse-tail distribution.

3 The Tail-Risk Correction Problem

3.1 From Generated Paths to Tail-Risk Evaluation

Tail risk is evaluated on the outcomes induced by generated trajectories, rather than directly on individual coordinates of those trajectories. A generated trajectory X is mapped to a scalar outcome $Y_j = \Pi_j(X)$ through a downstream functional Π_j. In the financial benchmark, Π_j represents a trading strategy and Y_j its profit and loss (PnL), where smaller values correspond to more adverse outcomes (e.g., $Y_j = -2$ represents a loss of two units). VaR, ES, and the tail law are therefore evaluated across the distribution of projected outcomes from simulated trajectories.

To formalise these risk measures, consider a generic scalar outcome Y with cumulative distribution function (CDF) $F(y) = \mathbb{P}(Y \le y)$ and finite first moment. For a chosen tail probability $\alpha \in (0, 1)$, denote the lower u-quantile of F by $Q_u(F) = \inf\{y : F(y) \ge u\}$. The left-tail Value-at-Risk (VaR) and Expected Shortfall (ES) are defined as

$$\begin{aligned} \mathrm{VaR}^L_\alpha(F) &= \inf\{y \in \mathbb{R} : F(y) \ge \alpha\} \qquad & \mathrm{ES}^L_\alpha(F) &= \frac{1}{\alpha}\int_0^\alpha Q_u(F)\,du \\ &= Q_\alpha(F) & &= \mathbb{E}\big[Y \mid Y \le \mathrm{VaR}^L_\alpha(F)\big]\,. \end{aligned} \tag{1}$$

Hereafter, VaR_α and ES_α denote these left-tail quantities. The conditional-mean form of ES holds when $F(\mathrm{VaR}_\alpha(F)) = \alpha$, in particular for continuous F; Appendix A.1 gives the corresponding right-tail convention.

Intuition for VaR and ES. VaR marks the boundary at which outcomes enter a chosen adverse tail, whereas ES measures the average outcome within that tail [McNeil et al., 2015]. VaR therefore describes where extreme outcomes begin, while ES describes their severity once that boundary has been crossed. Since the boundary can remain unchanged even when outcomes deeper in the tail deteriorate, ES complements the information provided by VaR. The coherence framework of Artzner et al. [1999] formalises broader requirements for risk measures. Within this framework, Acerbi and Tasche [2002] establish ES as a coherent alternative to VaR. The practical importance of VaR and ES is reinforced by regulation: VaR remains important for reporting and backtesting, while ES is the headline internal-model measure in the Basel market-risk framework [Basel Committee on Banking Supervision, 2006, 2019].

3.2 Why Empirical Tail Risk Is Difficult to Learn Directly

The empirical versions of these measures are suitable for reporting but provide awkward signals for gradient-based generation. Empirical VaR assigns the boundary to a selected order statistic, while empirical ES averages the observations currently assigned to the tail. Consequently, the VaR signal is concentrated near one boundary observation, the ES signal acts through a small subset of the batch, and both assignments change abruptly when sample ranks cross. The same scarcity that motivates synthetic tail generation therefore also limits the signal available for correcting it.

The projected tail law retains the information discarded by these two summaries. For the outcome law F, the left-α tail law is

$$F^\alpha(x) = \frac{\min\{F(x), \alpha\}}{\alpha}. \tag{2}$$

When F has an atom at $\mathrm{VaR}_\alpha(F)$, F^α keeps only the part of the atom that completes mass α. Its upper boundary is $\mathrm{VaR}_\alpha(F)$ and its mean is $\mathrm{ES}_\alpha(F)$, but it also distinguishes adverse distributions that share both quantities. It therefore provides a richer target than matching the two summaries alone.

A learnable correction requires an objective that identifies this complete tail law at the population level and provides a usable gradient for a finite generated batch. A proper scoring rule [Gneiting, 2011] supplies the first property by assigning the smallest expected score to the correct distribution. The quantile–tail-law construction of Fissler et al. [2026] scores the boundary together with the distribution below it. Section 4 uses this construction to define the differentiable TAILDIFF energy, while Appendix A.2 gives its formal derivation.

4 Tail-Guided Diffusion

TAILDIFF guides a pretrained diffusion model at inference time without changing its parameters. At selected low-noise reverse steps, it computes a soft quantile boundary and an associated left-tail law from a batch of generated paths. It scores this pair against fixed projected reference values using what we call tail-law CRPS (TL–CRPS), a CRPS instantiation of the scoring construction of Fissler et al. [2026, Sec. 5.3].

Figure 2 separates the ordinary reverse-diffusion path, the differentiable generated branch, and the fixed reference branch. The complete finite program and sampling procedure are given in Appendix B and Algorithm 1.

4.1 Denoise Generated Paths and Compute Projections

Denoise + project. For $s \in \{\mathrm{train}, \mathrm{val}, \mathrm{test}\}$, let $\mathscr{P}_s = \{\Pi_j^s\}_{j=1}^{J_s}$ denote the exact scalar projections for split s. TAILDIFF uses differentiable counterparts $\{\Pi_j^{\mathrm{guide}}\}_{j=1}^{J_{\mathrm{guide}}}$ for training projections, while $\mathscr{P}_{\mathrm{val}}$ and $\mathscr{P}_{\mathrm{test}}$ are used only for selection and final evaluation. In the financial benchmark, the projections are trading-strategy PnLs. Let $m = 1, \ldots, B$ index paths in the current guidance

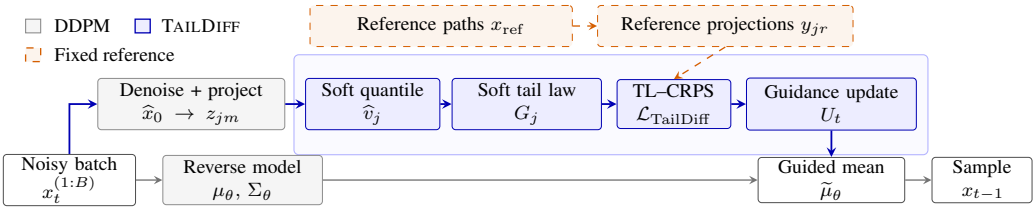


Figure 2: TAILDIFF at an active low-noise reverse step. From left to right, blue operations follow the subsection and paragraph headings in Section 4 and the execution order in Algorithm 1.

batch and $r = 1, \ldots, R$ index the training reference paths $\{x_{\text{ref}}^{(r)}\}_{r=1}^{R}$, and write $\mathbf{x}_t = \text{vec}(x_t^{(1:B)})$ for the stacked batch. At reverse step t, the DDPM backbone forms the clean estimate

$$\widehat{x}_0(x_t) = \frac{x_t - \sqrt{1-\bar{\alpha}_t}\,\epsilon_\theta(x_t, t)}{\sqrt{\bar{\alpha}_t}}, \tag{3}$$

which converts each noisy path into an estimate of its clean counterpart. For each projection, the generated values and projected reference values are

$$z_{jm} = \Pi_j^{\text{guide}}\left(\widehat{x}_0(x_t^{(m)})\right), \qquad y_{jr} = \Pi_j^{\text{guide}}\left(x_{\text{ref}}^{(r)}\right). \tag{4}$$

The generated values z_{jm} remain differentiable with respect to the current batch. The application-specific projection surrogates are given in Appendix D.3.

4.2 Construct the Soft Quantile and Tail Law

The generated quantile–tail-law pair $(\hat{v}_j, G_j)$ consists of a soft quantile boundary $\hat{v}_j$ and an associated soft left-tail law G_j. A detached version of the generated batch is used to compute the median absolute deviation (MAD), which sets the projection-specific smoothing temperatures as specified in Appendix B.2.1.

Soft quantile. Hard empirical quantiles and tail subsets yield sparse, rank-dependent gradients. We replace the hard quantile of the generated batch with the soft quantile boundary $\widehat{v}_j$, a bisection approximation of the root $\bar{v}_j$ of

$$\frac{1}{B}\sum_{m=1}^{B}\sigma(\lambda_{q,j}(\bar{v}_j - z_{jm})) = \alpha, \tag{5}$$

where σ is the logistic CDF and $\lambda_{q,j} > 0$ controls the smoothness of the approximation. As $\lambda_{q,j}$ increases, the distance from $\bar{v}_j$ to the empirical α-quantile set vanishes. For finite $\lambda_{q,j}$, the implicit derivative of $\bar{v}_j$, which backpropagation assigns to $\widehat{v}_j$, distributes the quantile gradient across generated observations near $\bar{v}_j$ rather than through a single order statistic.

Soft tail law. The same soft quantile boundary defines a differentiable left-tail law by setting

$$u_{jm} = \sigma(\lambda_{\text{tl},j}(\widehat{v}_j - z_{jm})), \qquad \pi_{jm} = \frac{u_{jm}}{\sum_{\ell=1}^{B} u_{j\ell}}, \qquad G_j = \sum_{m=1}^{B}\pi_{jm}\delta_{z_{jm}}, \tag{6}$$

where δ_z is a Dirac mass at z and the tail smoothing parameter $\lambda_{\text{tl},j} > 0$, distinct from $\lambda_{q,j}$ in Equation (5), controls how sharply u_{jm} falls from one to zero as z_{jm} crosses $\widehat{v}_j$. Hence, $(\widehat{v}_j, G_j)$ is a differentiable quantile–tail-law pair; Appendix B.2 details the soft-root solver, its implicit Jacobian, and the effective-tail conditions.

4.3 TL–CRPS Tail-Law Objective

Reference paths and projections. The projected reference values y_{jr} in Equation (4) are computed once from the fixed reference paths and detached before entering the objective.

TL–CRPS objective. Following the tail-score construction of Fissler et al. [2026, Sec. 5.3], we score the generated pair $(\widehat{v}_j, G_j)$ against every projected reference value y_{jr}. Let $A_j(u) = \mathrm{CRPS}(G_j, u) - 2u$ denote the reflected left-tail generator score. The indicator masks use the detached forward value $\widehat{v}_j^{\mathrm{detach}}$, while $A_j(\widehat{v}_j)$ remains differentiable. The finite-batch TAILDIFF objective over the active projections $\mathcal{J}_t = \{j : \sum_m u_{jm} \geq \max\{8, \alpha B/2\}\}$ is

$$S_{jr} = \mathbf{1}\{y_{jr} \leq \widehat{v}_j^{\mathrm{detach}}\} A_j(y_{jr}) + \Big(\mathbf{1}\{y_{jr} > \widehat{v}_j^{\mathrm{detach}}\} - (1-\alpha)\Big) A_j(\widehat{v}_j), \tag{7a}$$

$$\mathcal{L}_{\mathrm{TailDiff}} = \sum_{j \in \mathcal{J}_t} \frac{1}{R} \sum_{r=1}^{R} S_{jr}. \tag{7b}$$

The objective acts on more than the two point summaries used for evaluation. Its CRPS component compares the shape of the generated and reference tail laws, while its boundary component adjusts the generated soft quantile boundary. Let $\widehat{F}_{\mathrm{ref},j}$ denote the detached empirical reference CDF defined in Equation (24). Averaging the boundary term over the projected reference values then gives the coefficient $\alpha - \widehat{F}_{\mathrm{ref},j}(\widehat{v}_j^{\mathrm{detach}})$, which changes sign at the empirical reference α-quantile and vanishes there when αR is an integer. Computing $\widehat{v}_j$ from the generated batch therefore preserves this quantile-transport signal.

4.4 Guided Reverse-Mean Update

Guidance update. At an active reverse step, automatic differentiation (AD) of Equation (7b) through the projections and Tweedie's estimate, with $\mathcal{J}_t$ and the detached quantities held fixed, gives the stacked gradient $g_t = \nabla^{\mathrm{AD}}_{\mathbf{x}_t} \mathcal{L}_{\mathrm{TailDiff}}$. Let $g_{t,m}$ and $U_{t,m}$ denote the path-m blocks of g_t and of the update $U_t = \gamma g_t$. For guidance scale $\gamma > 0$, these blocks satisfy

$$g_{t,m} = \nabla^{\mathrm{AD}}_{x_t^{(m)}} \mathcal{L}_{\mathrm{TailDiff}}, \qquad U_{t,m} = \gamma g_{t,m}. \tag{8}$$

Guided mean and sample. The guided sampler applies the resulting shift directly to the reverse mean,

$$\widetilde{\mu}_{\theta,m}(\mathbf{x}_t, t) = \mu_\theta(x_t^{(m)}, t) - U_{t,m}, \qquad (t > 0): \quad x_{t-1}^{(m)} \sim \mathcal{N}(\widetilde{\mu}_{\theta,m}(\mathbf{x}_t, t), \Sigma_\theta(t)). \tag{9}$$

For a start fraction $\tau_{\mathrm{start}} \in (0, 1]$, we apply this shift only at selected steps $t/(T-1) \leq \tau_{\mathrm{start}}$ in the final low-noise part of the reverse process, including the deterministic terminal step $t = 0$, which returns $x_{\mathrm{out}}^{(m)} = \widetilde{\mu}_{\theta,m}(\mathbf{x}_0, 0)$ without adding noise. Appendix B.4 gives the exact scheduling conventions, while batching and the complete sampler are given in Algorithm 1.

5 Theoretical Analysis

Our analysis justifies TAILDIFF at three levels. First, the population energy uniquely identifies the intended projected tail law, validating the statistical target. Second, we quantify the gap between this ideal loss and the finite objective used during sampling. Third, we characterise the resulting backward field and bound its intervention on the pretrained sampler.

Population identification. Let F and F_{ref} denote the generated and reference laws of a projected outcome, let F^α and F^α_{ref} be their left-α tail laws, and write $v_F^\alpha = Q_\alpha(F)$. Specialising the construction of Fissler et al. [2026, Sec. 5.3] to pairs induced by a single candidate, let $L_{F_{\mathrm{ref}}}(v, G)$ denote the expected TL–CRPS for a candidate quantile–tail-law pair (v, G). TAILDIFF restricts this pair to (v_F^α, F^α), as induced by the candidate law, giving the population objective

$$\mathcal{L}(F) = L_{F_{\mathrm{ref}}}(v_F^\alpha, F^\alpha). \tag{10}$$

The resulting objective has the following identification and excess-score decomposition.

Theorem 1 (Population identification of the projected tail law)**.** *Suppose that F_{ref} has a finite first moment and a continuous CDF that is strictly increasing in a neighbourhood of its α-quantile. For any candidate law F with a finite first moment,*

$$\mathcal{L}(F) - \mathcal{L}(F_{\mathrm{ref}}) = \mathcal{B}_\alpha(F; F_{\mathrm{ref}}) + \alpha \int_{\mathbb{R}} \big(F^\alpha(x) - F^\alpha_{\mathrm{ref}}(x)\big)^2 \, dx \geq 0, \tag{11}$$

where the boundary discrepancy $\mathcal{B}_\alpha(F; F_{\mathrm{ref}}) := L_{F_{\mathrm{ref}}}(v_F^\alpha, F^\alpha) - L_{F_{\mathrm{ref}}}(v_{F_{\mathrm{ref}}}^\alpha, F^\alpha)$ *is non-negative and vanishes exactly when the generated and reference* α*-quantiles agree. Equality in Equation* (11) *holds if and only if* $F^\alpha = F_{\mathrm{ref}}^\alpha$.

Interpretation: Theorem 1 justifies the loss used by TAILDIFF. Its two terms correspond directly to the two operations in the TAILDIFF energy: the first corrects the boundary of the generated tail, while the second is α times the squared Cramér distance between the complete generated and reference tail laws. Minimising the population energy therefore recovers the intended projected tail law, with VaR and ES recovered as its boundary and mean, while still distinguishing tail distributions that share both summaries. Proposition 1 in Appendix C formalises this distinction from objectives that depend only on the VaR–ES pair, and Appendix C.2.1 proves the theorem.

Finite implementation. We next connect Theorem 1 to the implemented finite-batch objective used by the sampler. For projection j, let $\widehat{\ell}_j = R^{-1}\sum_r S_{jr}$ be the implemented score, F_j a candidate generated law with α-quantile $v_j^\star = Q_\alpha(F_j)$, and $L_j = L_{F_{\mathrm{ref},j}}$ the population score of Theorem 1 for the projected reference law $F_{\mathrm{ref},j}$. Their gap has the exact decomposition

$$\widehat{\ell}_j - L_j(v_j^\star, F_j^\alpha) = E_{\mathrm{ref},j} + E_{\mathrm{thr},j} + E_{\mathrm{tail},j}, \tag{12}$$

where the three terms arise from the finite projected reference sample, the estimated soft boundary, and the finite soft tail law, respectively. Propositions 2 and 3 decompose and bound the backward field of the finite objective, Proposition 4 controls the soft-quantile error at a fixed scale, and Proposition 5 bounds $E_{\mathrm{thr},j}$ and $E_{\mathrm{tail},j}$.

Controlled sampler intervention. Proposition 6 quantifies how strongly the tail-law energy can alter the pretrained sampler and motivates both the controlled low-noise guidance window and the update-magnitude diagnostic used in the experiments. The population and implementation results validate the objective supplied to the sampler, while the perturbation result controls its effect.

6 Experiments

We evaluate projected tail accuracy, calibration, structural fidelity, and guidance stability. The primary evaluation follows TailGAN [Cont et al., 2025] and DSI [Zhao et al., 2026] on controlled synthetic data and Nasdaq ITCH market data. Unless stated otherwise, we compare full-trajectory simulators using $N = 1000$ equally weighted scenarios.

6.1 Experimental Protocol

The synthetic benchmark is a five-dimensional process with known ground truth, while the market benchmark contains multivariate AAPL, AMZN, GOOG, JPM, and QQQ trajectories. In each setting, the training, validation, and test projection sets $\mathscr{P}_{\mathrm{train}}$, $\mathscr{P}_{\mathrm{val}}$, and $\mathscr{P}_{\mathrm{test}}$ contain different strategy instances generated by the same procedure. Projection-dependent fitting and TAILDIFF guidance use $\mathscr{P}_{\mathrm{train}}$, tuning uses $\mathscr{P}_{\mathrm{val}}$, and final metrics use a common $J_{\mathrm{test}} = 65$ set $\mathscr{P}_{\mathrm{test}}$ across methods. All methods also share the trajectory splits. The primary metric is VaR and ES relative error (RE) at $\alpha = 0.05$, for which lower is better. RE is defined as

$$\mathrm{RE}(M) = \frac{100}{2J_{\mathrm{test}}} \sum_{j=1}^{J_{\mathrm{test}}} \left[\frac{\left|\widehat{\mathrm{VaR}}_{\alpha,j}^{(M)} - \mathrm{VaR}_{\alpha,j}^\star\right|}{\left|\mathrm{VaR}_{\alpha,j}^\star\right|} + \frac{\left|\widehat{\mathrm{ES}}_{\alpha,j}^{(M)} - \mathrm{ES}_{\alpha,j}^\star\right|}{\left|\mathrm{ES}_{\alpha,j}^\star\right|} \right], \tag{13}$$

where M denotes the method, $J_{\mathrm{test}} = 65$ is the number of test projections, hats denote estimates from generated paths, and $\mathrm{VaR}_{\alpha,j}^\star$ and $\mathrm{ES}_{\alpha,j}^\star$ are the reference values. The factor 100 expresses RE as a percentage. We additionally examine statistical calibration, dependence and distributional fidelity, and the behaviour of the guidance field.

We compare with historical simulation (HSM), WGAN [Arjovsky et al., 2017, Gulrajani et al., 2017], TailGAN [Cont et al., 2025], DDPM [Ho et al., 2020], and DSI [Zhao et al., 2026]. TAILDIFF, DDPM, and DSI use the same pretrained diffusion run, with TAILDIFF differing through its inference-time sampling rule and requiring no retraining. Appendix D gives the datasets, projections, baseline protocol, metrics, architecture, hyperparameters, and computational resources.

6.2 Tail-Risk Accuracy and Calibration

Table 1: Average VaR and ES relative error (RE) (%) at $\alpha = 0.05$ from $N = 1000$ generated trajectories. Following Cont et al. [2025], SE(1000) is the matched-budget oracle sampling error. Lower is better.

Method	Synthetic RE	Market RE
HSM	3.4 ± 2.6	11.4 ± 2.3
SE(1000) / Oracle	3.1 ± 2.3	2.6 ± 1.1
WGAN [Arjovsky et al., 2017]	17.3 ± 2.2	29.9 ± 1.9
DDPM [Ho et al., 2020]	7.7 ± 2.5	16.3 ± 4.3
TailGAN [Cont et al., 2025]	5.1 ± 1.4	12.1 ± 1.3
DSI [Zhao et al., 2026]	3.7 ± 0.2	10.6 ± 1.2
TAILDIFF (Ours)	$\mathbf{3.4 \pm 1.1}$	$\mathbf{8.2 \pm 1.4}$

Table 1 reports the primary comparison. On synthetic data, TAILDIFF attains 3.4% RE, close to the 3.1% SE(1000) oracle sampling error. On market data, it attains 8.2% RE, compared with 10.6% for DSI, 12.1% for TailGAN, and 16.3% for the unguided DDPM backbone.

Statistical calibration. Because the required ground-truth tail functionals are unavailable for market data, we conduct the Kupiec coverage and joint Fissler–Ziegel score tests on the synthetic benchmark. Table 2 reports rejection rates across all methods. At the nominal 5% level, the coverage test has a finite-sample rejection rate of 16.6% even for independent draws from the true law, because its threshold is estimated from 1000 generated paths. TAILDIFF obtains 17.17%, close to this reference rate, and a 0.00% rejection rate under the joint score test. Appendix D.5 derives the finite-sample test size.

Table 2: Synthetic backtest rejection rates (%) over 100 trials with $N = 1000$. The coverage test has a finite-sample size of 16.6% for independent draws from the true law.

Test	HSM	WGAN	TailGAN	DDPM	DSI	**TAILDIFF**
Coverage	16.32	70.35	33.55	53.62	20.26	**17.17**
Score	0.00	14.92	0.29	5.06	0.00	**0.00**

6.3 Guidance Mechanism, Stability, and Structural Fidelity

Figure 3 connects the empirical behaviour of TAILDIFF to the gradient and sampler-intervention analyses in Section 5. It compares gradient magnitude, adjacent-step directional coherence, and correlation and autocorrelation distances with those produced by alternative guidance objectives.

Contribution of the tail objective. With the backbone, sampler, tuning budget, and evaluation protocol held consistent, TAILDIFF attains 8.2% market RE, compared with 13.1% for raw CRPS (full-distribution scoring), 23.7% for MSE-ES, and 37.4% for Pinball Loss. Pinball Loss targets the quantile boundary without shaping the distribution below it, MSE-ES reduces the tail to a point summary, and raw CRPS scores the full projected distribution without explicitly targeting the left-tail law. In contrast, TAILDIFF jointly scores the generated quantile boundary and its associated tail law. Table 3 summarises the ablation results, with additional details in Appendix F.1.

Guidance dynamics and structural fidelity. On the market benchmark, Figure 3(a)–(b) show that TAILDIFF maintains controlled gradient magnitudes and coherent directions over neighbouring reverse steps. Figure 3(c) reports correlation and autocorrelation distances for both benchmarks. The full correlation, autocorrelation, and rank-frequency results in Appendix F.4 further show that guidance shifts rare-event frequencies towards the reference distribution without the dependence deterioration observed for the less accurate learned baselines. These diagnostics are consistent with the bounded-guidance and perturbation analysis.

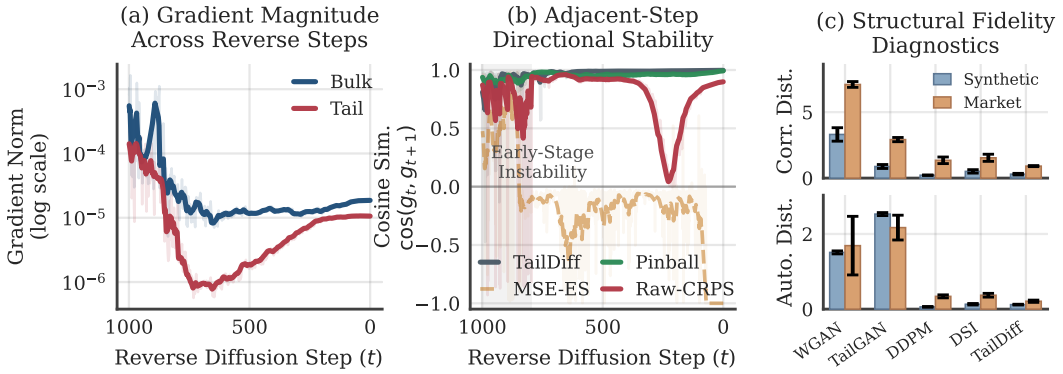


Figure 3: **Guidance dynamics and structural fidelity. (a)** Gradient magnitude across reverse steps. **(b)** Adjacent-step gradient cosine similarity. **(c)** Correlation and autocorrelation distances (lower is better). Compared with MSE-ES, Pinball Loss, and raw CRPS, TAILDIFF has controlled low-noise gradients, coherent adjacent-step directions, and low dependence errors.

Table 3: Market guidance-objective ablation with the same backbone and sampler. Only the guidance loss changes. Values are average VaR/ES relative error (%, mean±std over 5 seeds, $\alpha = 0.05$, $N = 1000$). Adjacent-step cosine measures directional coherence between consecutive guidance gradients.

Objective	RE %	Adj. cosine	Behaviour
MSE-ES $(\widehat{\mathrm{ES}} - r)^2$	23.7 ± 1.1	0.39	diverges
Pinball Loss (quantile only)	37.4 ± 0.5	0.95	stable, no tail-shape control
Raw CRPS (full distribution)	13.1 ± 1.9	0.78	oscillates at low noise
TAILDIFF (Ours)	$\mathbf{8.2 \pm 1.4}$	**0.97**	stable, coherent

Guidance scale and update magnitude. The guidance scale must be chosen jointly with the active low-noise window. Figure 4 reports validation RE and shows that $\tau_{\text{start}} = 0.30$ performs best across nearly the entire sweep. On synthetic data, RE reaches a stable plateau over $\gamma \in [8, 12]$, from which we select $\gamma = 10$. On market data, RE reaches its minimum at $\gamma = 24$ before increasing moderately at larger scales. We fix these configurations before test evaluation. On the market benchmark, the guidance-update magnitude increases with γ, while excessively large updates worsen both RE and dependence diagnostics. Appendix F.3 reports these measurements.

6.4 Robustness and Cross-Domain Portability

Appendices E.1 and E.2 show improving tail accuracy with N and consistent behaviour across $\alpha \in \{0.01, 0.05, 0.10\}$.

Cross-domain portability. To assess portability beyond finance, we apply TAILDIFF to stress functionals from energy, traffic, healthcare, and weather data. Each domain uses a DDPM trained on its own training data. TAILDIFF is then applied to the corresponding pretrained backbone at inference time without further retraining.

Table 4 shows that TAILDIFF improves tail fidelity in all four domains, reducing the average DDPM relative error by approximately 69%. Historical simulation is stronger in several domains because the corresponding DDPM backbones use the default architecture and training configuration. We do not tune those backbones for each domain, since the purpose of this experiment is to test whether TAILDIFF improves a pretrained diffusion model across domains. It is not intended to establish state-of-the-art performance against domain-specific backbone designs. See Appendix G for details.

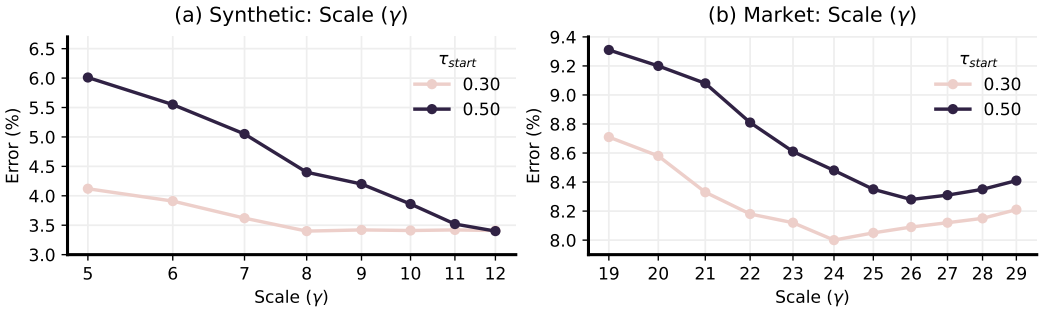


Figure 4: **Guidance-scale sensitivity under low-noise scheduling. (a) Synthetic** and **(b) Market** validation relative error as the guidance scale γ varies for two active windows τ_{start}. Lower is better.

Table 4: Cross-domain portability of TAILDIFF at inference time. Each row uses the corresponding DDPM trained on that domain's training data. SE(1000) is the matched-budget oracle sampling error. Lower relative error is better.

Domain	Dataset	SE(1000)	HSM	DDPM	TAILDIFF
Energy	PJM	2.3 ± 1.3	4.7 ± 2.1	27.6 ± 2.3	9.2 ± 1.4
Traffic	NYC taxi	2.1 ± 0.9	3.2 ± 2.5	36.2 ± 3.7	9.4 ± 1.7
Healthcare	ECG5000	2.4 ± 1.1	7.1 ± 1.2	17.3 ± 3.3	5.6 ± 1.8
Weather	Monash	3.1 ± 1.2	8.6 ± 1.7	37.6 ± 2.4	12.3 ± 3.4

7 Discussion

TAILDIFF addresses a practical failure mode in risk-sensitive generative simulation by guiding pretrained diffusion models towards reference projected tail laws at inference time, without retraining or architectural changes. Its TL–CRPS objective targets the distribution below the generated quantile boundary rather than only point-valued risk measures. Our theory validates this population target, quantifies the finite implementation gap, and bounds the sampler perturbation by the accumulated realised guidance updates. Across the controlled, financial, and cross-domain benchmarks, the results demonstrate improved representation of tail risk in generated scenarios while maintaining key dependence diagnostics, supporting the portability of the guidance mechanism beyond finance.

Limitations and future work. Several limitations suggest natural directions for future work. First, evaluating TL–CRPS gradients at active reverse steps increases inference time, motivating adaptive guidance schedules and computationally cheaper batch estimators. Second, our experiments consider $\alpha \geq 0.01$, whereas more extreme tail levels reduce the effective tail count, which scales with αB, motivating adaptive tail batching and integration with rare-event sampling. Third, extending our population-identification, finite-implementation, and sampler-perturbation results to convergence of the guided sampler towards the target tail law remains an important direction. Finally, benchmarks for tail-risk estimation remain limited, motivating broader evaluation across finance, energy, weather, and health.

Broader impact. Improved tail-faithful simulation can support safer decision-making in domains where rare events carry high societal cost, including financial stress testing, energy-demand planning, climate and weather extremes, infrastructure resilience, insurance, and public-health monitoring. By targeting downstream tail functionals under finite simulation budgets, TAILDIFF provides a practical mechanism for evaluating low-probability, high-impact scenarios that are often underrepresented by standard generative objectives. More reliable synthetic stress scenarios can help practitioners assess robustness, allocate resources, and design safeguards in settings where collecting sufficient real extreme-event data is difficult.

Acknowledgments and Disclosure of Funding

This work was supported in part by the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (No. JYB2025XDXM504) and the National Natural Science Foundation of China (No. U2341228). The authors declare no competing interests.

References

Carlo Acerbi and Dirk Tasche. Expected shortfall: A natural coherent alternative to value at risk. *Economic Notes*, 31(2):379–388, 2002. doi: 10.1111/1468-0300.00091.

Erik Altman, Jovan Blanuša, Luc von Niederhäusern, Beni Egressy, Andreea Anghel, and Kubilay Atasu. Realistic synthetic financial transactions for anti-money laundering models. In *Advances in Neural Information Processing Systems*, volume 36, pages 29851–29874, 2023. doi: 10.52202/075280-1300.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223. PMLR, 2017.

Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999. doi: 10.1111/1467-9965.00068.

Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Roni Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Universal guidance for diffusion models. In *International Conference on Learning Representations*, 2024.

Basel Committee on Banking Supervision. International Convergence of Capital Measurement and Capital Standards: A Revised Framework. Bank for International Settlements, 2006. URL `https://www.bis.org/publ/bcbs128.htm`.

Basel Committee on Banking Supervision. Minimum Capital Requirements for Market Risk. Bank for International Settlements, 2019. URL `https://www.bis.org/bcbs/publ/d457.htm`.

Andrea Coletta, Sriram Gopalakrishnan, Daniel Borrajo, and Svitlana Vyetrenko. On the constrained time-series generation problem. In *Advances in Neural Information Processing Systems*, volume 36, pages 61048–61059, 2023. doi: 10.52202/075280-2668.

Rama Cont, Mihai Cucuringu, Renyuan Xu, and Chao Zhang. Tail-GAN: Learning to simulate tail risk scenarios. *Management Science*, 72(4):2917–2936, 2025. doi: 10.1287/mnsc.2023.00936.

Hoang Anh Dau, Eamonn Keogh, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatana, Yanping Chen, Bing Hu, Nurjahan Begum, Anthony Bagnall, Abdullah Mueen, Gustavo Batista, and Hexagon-ML. The UCR time series classification archive, 2018. URL `https://www.cs.ucr.edu/~eamonn/time_series_data_2018/`.

Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794, 2021.

Yilun Du, Conor Durkan, Robin Strudel, Joshua B. Tenenbaum, Sander Dieleman, Rob Fergus, Jascha Sohl-Dickstein, Arnaud Doucet, and Will Sussman Grathwohl. Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and MCMC. In *Proceedings of the International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 8489–8510. PMLR, 2023.

Tobias Fissler and Johanna F. Ziegel. Higher order elicitability and Osband's principle. *The Annals of Statistics*, 44(4):1680–1707, 2016. doi: 10.1214/16-AOS1439.

Tobias Fissler, Johanna F. Ziegel, and Tilmann Gneiting. Expected shortfall is jointly elicitable with value at risk – implications for backtesting. *Risk*, pages 58–61, 2016.

Tobias Fissler, Fangda Liu, Ruodu Wang, and Linxiao Wei. Elicitability and identifiability of tail risk measures. *Mathematical Finance*, 36(3):451–464, 2026. doi: 10.1111/mafi.70016.

Asadullah Hill Galib, Pang-Ning Tan, and Lifeng Luo. FIDE: Frequency-inflated conditional diffusion model for extreme-aware time series generation. In *Advances in Neural Information Processing Systems*, volume 37, pages 114434–114457, 2024. doi: 10.52202/079017-3634.

Tilmann Gneiting. Making and evaluating point forecasts. *Journal of the American Statistical Association*, 106(494):746–762, 2011. doi: 10.1198/jasa.2011.r10138.

Rakshitha Godahewa, Christoph Bergmeir, Geoffrey I. Webb, Rob J. Hyndman, and Pablo Montero-Manso. Monash time series forecasting archive. In *Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021.

Tal Gonen, Itai Pemper, Ilan Naiman, Nimrod Berman, and Omri Azencot. Time series generation under data scarcity: A unified generative modeling approach. In *Advances in Neural Information Processing Systems*, volume 38, pages 191446–191485, 2025. doi: 10.52202/085713-5750.

Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C. Courville. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems*, volume 30, pages 5769–5779, 2017.

Xuming He, Xiaoou Pan, Kean Ming Tan, and Wen-Xin Zhou. Smoothed quantile regression with large-scale inference. *Journal of Econometrics*, 232(2):367–388, 2023. doi: 10.1016/j.jeconom.2021.07.010.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851, 2020.

Jiaqiao Hu, Meichen Song, and Michael C. Fu. Quantile optimization via multiple-timescale local search for black-box functions. *Operations Research*, 73(3):1535–1557, 2024. doi: 10.1287/opre.2022.0534.

Todd Huster, Jeremy Cohen, Zinan Lin, Kevin Chan, Charles Kamhoua, Nandi O. Leslie, Cho-Yu Jason Chiang, and Vyas Sekar. Pareto GAN: Extending the representational power of GANs to heavy-tailed distributions. In *Proceedings of the International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4523–4532. PMLR, 2021.

Ying Jin, Jinlong Peng, Qingdong He, Teng Hu, Jiafu Wu, Hao Chen, Haoxuan Wang, Wenbing Zhu, Mingmin Chi, Jun Liu, and Yabiao Wang. Dual-interrelated diffusion model for few-shot anomaly image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 30420–30429, 2025. doi: 10.1109/CVPR52734.2025.02832.

Seong-Tae Kim. PJM hourly energy consumption. IEEE Dataport, 2025. URL `https://dx.doi.org/10.21227/wqg6-9917`.

Marcel Kollovieh, Abdul Fatir Ansari, Michael Bohlke-Schneider, Jasper Zschiegner, Hao Wang, and Yuyang Bernie Wang. Predict, refine, synthesize: Self-guiding diffusion models for probabilistic time series forecasting. In *Advances in Neural Information Processing Systems*, volume 36, pages 28341–28364, 2023. doi: 10.52202/075280-1232.

Ira Ktena, Olivia Wiles, Isabela Albuquerque, Sylvestre-Alvise Rebuffi, Ryutaro Tanno, Abhijit Guha Roy, Shekoofeh Azizi, Danielle Belgrave, Pushmeet Kohli, Taylan Cemgil, Alan Karthikesalingam, and Sven Gowal. Generative models improve fairness of medical classifiers under distribution shifts. *Nature Medicine*, 30(4):1166–1173, 2024. doi: 10.1038/s41591-024-02838-6.

Paul H. Kupiec. Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2):73–84, 1995. doi: 10.3905/jod.1995.407942.

Junjie Li, Yang Liu, Weiqing Liu, Shikai Fang, Lewen Wang, Chang Xu, and Jiang Bian. MARS: A financial market simulation engine powered by generative foundation model. In *International Conference on Learning Representations*, 2025a.

Yang Li, Han Meng, Zhenyu Bi, Ingolv T. Urnes, and Haipeng Chen. Population aware diffusion for time series generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 18520–18529, 2025b. doi: 10.1609/aaai.v39i17.34038.

Lequan Lin, Zhengkun Li, Ruikun Li, Xuliang Li, and Junbin Gao. Diffusion models for time-series applications: A survey. *Frontiers of Information Technology & Electronic Engineering*, 25(1): 19–41, 2024. doi: 10.1631/FITEE.2300310.

Haoyu Liu, Tingyu Zhu, Nanshan Jia, Jinghai He, and Zeyu Zheng. Learning to simulate from heavy-tailed distribution via diffusion model. *Operations Research*, 2026. doi: 10.1287/opre.2024.1450.

Jingwei Liu, Ling Yang, Hongyan Li, and Shenda Hong. Retrieval-augmented diffusion models for time series forecasting. In *Advances in Neural Information Processing Systems*, volume 37, pages 2766–2786, 2024. doi: 10.52202/079017-0091.

Alexander J. McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, 2015.

Ilan Naiman, Nimrod Berman, Itai Pemper, Idan Arbiv, Gal Fadlon, and Omri Azencot. Utilizing image transforms and diffusion models for generative modeling of short and long time series. In *Advances in Neural Information Processing Systems*, volume 37, pages 121699–121730, 2024. doi: 10.52202/079017-3868.

New York City Taxi and Limousine Commission. TLC trip record data, 2024. URL `https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page`.

Kushagra Pandey, Jaideep Pathak, Yilun Xu, Stephan Mandt, Michael Pritchard, Arash Vahdat, and Morteza Mardani. Heavy-tailed diffusion models. In *International Conference on Learning Representations*, 2025.

Dae-Young Park and In-Young Ko. Tabular generative model with graph-theoretical methodologies toward effective financial fraud detection. *ACM Transactions on Knowledge Discovery from Data*, 20(6):1–33, 2026. doi: 10.1145/3819586.

Kashif Rasul, Calvin Seward, Ingmar Schuster, and Roland Vollgraf. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In *Proceedings of the International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8857–8868. PMLR, 2021.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, volume 32, pages 11918–11930, 2019.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

Tomonori Takahashi and Takayuki Mizuno. Generation of synthetic financial time series by diffusion models. *Quantitative Finance*, 25(10):1507–1516, 2025. doi: 10.1080/14697688.2025.2528697.

Yuki Tanaka, Ryuji Hashimoto, Takehiro Takayanagi, Zhe Piao, Yuri Murayama, and Kiyoshi Izumi. CoFinDiff: Controllable financial diffusion model for time series generation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 9357–9365, 2025. doi: 10.24963/ijcai.2025/1040.

Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. In *Advances in Neural Information Processing Systems*, volume 34, pages 24804–24816, 2021.

Magnus Wiese, Robert Knobloch, Ralf Korn, and Peter Kretschmer. Quant GANs: Deep generation of financial time series. *Quantitative Finance*, 20(9):1419–1440, 2020. doi: 10.1080/14697688.2020.1730426.

Haochong Xia, Shuo Sun, Xinrun Wang, and Bo An. Market-GAN: Adding control to financial market data generation with semantic context. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 15996–16004, 2024. doi: 10.1609/aaai.v38i14.29531.

Jinhan Xie, Xiaodong Yan, Bei Jiang, and Linglong Kong. Statistical inference for smoothed quantile regression with streaming data. *Journal of Econometrics*, 249:105924, 2025. doi: 10.1016/j.jeconom.2024.105924.

Chejian Xu, Aleksandr Petiushko, Ding Zhao, and Bo Li. DiffScene: Diffusion-based safety-critical scenario generation for autonomous vehicles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 8797–8805, 2025. doi: 10.1609/aaai.v39i8.32951.

Jie Yang, Rui Zhang, Ziyang Cheng, Dawei Cheng, Guang Yang, and Bo Wang. Grad: Guided relation diffusion generation for graph augmentation in graph fraud detection. In *Proceedings of the ACM Web Conference*, pages 5308–5319, 2025. doi: 10.1145/3696410.3714520.

Jinsung Yoon, Daniel Jarrett, and Mihaela van der Schaar. Time-series generative adversarial networks. In *Advances in Neural Information Processing Systems*, volume 32, pages 5508–5518, 2019.

Jiwen Yu, Yinhuai Wang, Chen Zhao, Bernard Ghanem, and Jian Zhang. FreeDoM: Training-free energy-guided conditional diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23174–23184, 2023. doi: 10.1109/ICCV51070.2023.02118.

Xinyu Yuan and Yan Qiao. Diffusion-TS: Interpretable diffusion for general time series generation. In *International Conference on Learning Representations*, 2024.

Shuning Zhao, Patrick Wong, Leran Zhang, and Xiaolin Hu. Diachronic sample integration: Robust tail-risk estimation with generative models. *arXiv preprint arXiv:2607.10810*, 2026.

Appendix Table of Contents

A Formal Background **16**

A.1 Formal Risk Definitions and Tail Conventions 16

A.2 Scoring the Tail Law 17

A.3 DDPM and Guidance Mechanics 17

B TAILDIFF Finite Implementation and Sampling Algorithm **18**

B.1 Denoise and Project Generated and Reference Paths 18

B.1.1 Generated Denoising and Projection 18

B.1.2 Fixed Reference Paths and Projections 18

B.2 Construct the Soft Quantile and Tail Law 18

B.2.1 Soft Quantile 19

B.2.2 Soft Tail Law 19

B.3 TL–CRPS Tail-Law Objective 19

B.3.1 Weighted Empirical CRPS 20

B.3.2 Reference-Conditional Score and Loss 20

B.3.3 Gradient Ownership and Detachment 20

B.4 Guided Reverse-Mean Update 20

B.5 Complete Sampling Algorithm 20

B.6 Batching and Computational Cost 20

C Theoretical Results and Proofs **21**

C.1 Notation and Preliminary Identities 22

C.2 Population Identification 22

C.2.1 Proof of Theorem 1 22

C.2.2 Interpretation and Scope 23

C.2.3 Relation to the Tail-Score Construction 24

C.2.4 Identification Beyond VaR–ES Matching 24

C.2.5 A Concrete Indistinguishable Pair 25

C.3 Finite Implementation 25

C.3.1 Gradient Decomposition of the Finite Objective 25

C.3.2 Conditional Local Gradient Bound 26

C.3.3 Soft-Quantile Error 28

C.3.4 Finite Implementation Gap 29

C.4 Controlled Sampler Intervention 31

C.4.1 Terminal Perturbation Bound 31

D Experimental Protocol and Reproducibility **33**

D.1 Datasets and Splits 33

D.1.1 Synthetic Data Generation Parameters 33

D.1.2 Market Data . . . 33
D.2 Evaluation Projections and Strategy Definitions . . . 34
D.3 Differentiable Guidance Surrogates . . . 34
D.4 Baselines and Comparison Protocol . . . 36
D.5 Metrics and Backtest Calibration . . . 36
D.6 Base Diffusion Architecture . . . 37
D.7 Compute Resources and Hyperparameters . . . 37

E Additional Accuracy and Robustness Results **37**
E.1 Generation-Budget Sensitivity and Guidance-Batch Stability . . . 37
E.2 Robustness across Regulatory Risk Levels . . . 37

F Guidance Mechanism, Stability, and Structural Fidelity **38**
F.1 Contribution of the Tail Objective . . . 38
F.2 Extended Gradient Dynamics . . . 38
F.3 Guidance Scale, Low-Noise Window, and Update Magnitude . . . 39
F.4 Structural and Distributional Fidelity . . . 39

G Cross-Domain Portability **43**
G.1 Experimental Protocol . . . 43
G.2 Datasets and Stress Functionals . . . 43

H Changes from the Reviewed Version **43**

A Formal Background

This appendix collects the formal risk, scoring, and diffusion definitions used in Sections 3 and 4. The finite TAILDIFF program remains in Appendix B.

A.1 Formal Risk Definitions and Tail Conventions

For completeness, let $\mathcal{M}^1$ denote the Borel probability laws on $\mathbb{R}$ with a finite first absolute moment, using the same symbol for a law and its CDF. For $F \in \mathcal{M}^1$, write $v_F^\alpha = Q_\alpha(F) = \inf\{x : F(x) \geq \alpha\}$. When F has an atom at v_F^α, its left-α tail law includes only the fraction of the boundary atom required to obtain total mass α. For every Borel set A,

$$F^\alpha(A) = \frac{1}{\alpha}\left[F\big(A \cap (-\infty, v_F^\alpha)\big) + \big\{\alpha - F(v_F^\alpha-)\big\}\mathbf{1}\{v_F^\alpha \in A\}\right]. \tag{14}$$

For a continuous CDF this reduces to Equation (2). This fractional-boundary convention makes the tail law well-defined for empirical distributions and resolves the distinction between strict and non-strict threshold inequalities. For every $F \in \mathcal{M}^1$, we define $\mathrm{ES}_\alpha(F) = \alpha^{-1}\int_0^\alpha Q_u(F)\,du$, the quantile integral in Equation (1). Since $Q_U(F)$ has law F^α when U is uniform on $(0, \alpha)$, this is the mean $\int y\,dF^\alpha(y)$ of the tail law. The conditional mean $\mathbb{E}[Y \mid Y \leq v_F^\alpha]$ in Equation (1) agrees with it whenever $F(v_F^\alpha) = \alpha$, in particular for every continuous F.

To see the finite-sample object explicitly, let $y_{(1)} \leq \cdots \leq y_{(B)}$ be a sorted batch, set $k = \lfloor \alpha B \rfloor$, and write $\rho = \alpha B - k$. Its empirical tail law is

$$\widehat{F}_B^\alpha = \frac{1}{\alpha B}\left(\sum_{m=1}^{k} \delta_{y_{(m)}} + \rho\,\delta_{y_{(k+1)}}\right), \tag{15}$$

with the second term omitted when $\rho = 0$. The empirical VaR is the corresponding boundary order statistic and the empirical ES is the mean of $\widehat{F}_B^{\alpha}$.

For right-tail applications, where larger outcomes are adverse, we use

$$\mathrm{VaR}_\alpha^R(F) = \inf\{y : F(y) \geq 1-\alpha\}, \qquad \mathrm{ES}_\alpha^R(F) = \frac{1}{\alpha}\int_{1-\alpha}^{1} Q_u(F)\,du. \tag{16}$$

For a continuous law, the corresponding CDF is $F^{R,\alpha}(x) = [F(x) - (1-\alpha)]_+/\alpha$. Applying the left-tail convention to $-Y$, with law F_{-Y}, converts a right-tail loss problem into the left-tail PnL convention used in the main financial experiments: $\mathrm{ES}_\alpha^R(F) = -\mathrm{ES}_\alpha(F_{-Y})$ for every $F \in \mathcal{M}^1$, whereas $-\mathrm{VaR}_\alpha(F_{-Y}) = \inf\{y : F(y) > 1-\alpha\}$ equals $\mathrm{VaR}_\alpha^R(F)$ only when F has a unique $(1-\alpha)$-quantile.

A.2 Scoring the Tail Law

TAILDIFF requires a learning objective that compares a candidate tail law with observed outcomes. A scoring rule $S(G, y)$ assigns a loss to a candidate distribution G after observing y. It is proper if, for every data-generating distribution F,

$$\mathbb{E}_{Y\sim F}[S(F,Y)] \leq \mathbb{E}_{Y\sim F}[S(G,Y)] \qquad \text{for every candidate distribution } G, \tag{17}$$

and it is strictly proper if equality holds only when $G = F$. The continuous ranked probability score (CRPS) is a strictly proper scoring rule relative to $\mathcal{M}^1$ [Fissler et al., 2026, Sec. 5.3]. For a candidate CDF $G \in \mathcal{M}^1$ and observation $y \in \mathbb{R}$, it is

$$\mathrm{CRPS}(G, y) = \int_{\mathbb{R}} \big(G(z) - \mathbf{1}\{y \leq z\}\big)^2\, dz, \tag{18}$$

whose expected value is therefore uniquely minimised by the data-generating law. Following the reflected left-tail construction of Fissler et al. [2026, Sec. 5.3 and Sec. 6], define

$$A_G(y) = \mathrm{CRPS}(G, y) - 2y. \tag{19}$$

If G^- denotes the law of $-X$ for $X \sim G$, the reflection identity

$$\mathrm{CRPS}(G^-, -y) + 2(-y) = \mathrm{CRPS}(G, y) - 2y = A_G(y) \tag{20}$$

converts the right-tail loss construction into the left-tail PnL convention. For a proposed boundary v and candidate tail law G, the resulting score is

$$\mathcal{S}_{\text{TL–CRPS}}(v, G; y) = \mathbf{1}\{y \leq v\}A_G(y) + \big(\mathbf{1}\{y > v\} - (1-\alpha)\big)A_G(v). \tag{21}$$

Under the regularity conditions stated in Section 5, the expected score is minimised at the reference quantile and its associated tail law. Equation (21) is the ideal population score. Section 4 and Appendix B construct its differentiable finite-batch counterpart, including the soft boundary, soft tail membership, and declared gradient ownership.

While proper scoring rules evaluate distributional predictions, the corresponding concept for point-valued statistical functionals is *elicitability*. A functional is elicitable if there exists a scoring function whose expected value is uniquely minimised at the functional's true value. VaR is elicitable, whereas ES is not elicitable on its own [Gneiting, 2011]. However, the pair $(\mathrm{VaR}_\alpha, \mathrm{ES}_\alpha)$ is jointly elicitable [Fissler and Ziegel, 2016], enabling TailGAN [Cont et al., 2025] to construct a training objective based on these two point-valued risk measures. Such an objective cannot distinguish distributions sharing the same VaR and ES. In contrast, the proper distributional score above identifies the quantile boundary together with the complete tail law, distinguishing tail distributions that share both risk measures.

A.3 DDPM and Guidance Mechanics

A DDPM [Ho et al., 2020] defines a forward Markov chain that gradually corrupts clean data $x_0 \sim p_{\mathrm{data}}$ over T steps. With $\bar{\alpha}_t = \prod_{s=1}^{t}(1-\beta_s)$, its marginal forward transition and learned reverse transition are

$$q(x_t \mid x_0) = \mathcal{N}\big(x_t; \sqrt{\bar{\alpha}_t}x_0, (1-\bar{\alpha}_t)I\big), \qquad p_\theta(x_{t-1} \mid x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(t)). \tag{22}$$

The noise predictor parameterises the reverse mean and gives the clean estimate in Equation (3). Score-based diffusion [Song et al., 2021] gives the corresponding continuous-time interpretation. The sampler of Section 4 indexes its reverse steps by $t = T-1, \dots, 0$, starting from $x_{T-1} \sim \mathcal{N}(0, I)$, so its step t uses noise level $t+1$ of Equation (22). In Equations (3) and (9), Algorithm 1 and Proposition 6, $\bar{\alpha}_t$, $\epsilon_\theta(\cdot, t)$, $\mu_\theta(\cdot, t)$ and $\Sigma_\theta(t)$ therefore stand for $\bar{\alpha}_{t+1}$, $\epsilon_\theta(\cdot, t+1)$, $\mu_\theta(\cdot, t+1)$ and $\Sigma_\theta(t+1)$ of Equation (22), and the deterministic terminal step $t = 0$ returns x_{out}, which plays the role of the clean x_0.

Inference-time guidance modifies this reverse transition using a direction evaluated during sampling. TAILDIFF obtains this direction by differentiating its batch-level projected-tail energy through the clean estimate in Equation (3), and shifts the reverse mean according to Equation (9). The precise update, low-noise schedule, and batching scheme are given in Appendix B.4 and Algorithm 1.

B TAILDIFF Finite Implementation and Sampling Algorithm

This appendix follows the four stages of Section 4 and Figure 2: denoise and project the generated and fixed reference paths, construct the differentiable quantile–tail-law pair, evaluate the TL–CRPS objective, and update the reverse mean. Algorithm 1 then collects these stages into the complete sampling procedure.

Notation and gradient convention. Let $t = 0, \dots, T-1$ index diffusion steps, $j = 1, \dots, J_{\mathrm{guide}}$ index scalar projections, $m = 1, \dots, B$ index paths in the current guidance batch, and $r = 1, \dots, R$ index the fixed reference paths. Here, B is the realised size of the current batch and $B_{\max}$ is the configured batch-size cap. We reserve i for indexing individual generated paths and their corresponding terminal marginals. The total number of returned paths is N, which is distinct from B and the number R of reference paths, and $\mathbf{x}_t = \mathrm{vec}(x_t^{(1:B)})$ denotes the stacked state of an active guidance batch. A superscript detach marks a forward value held constant in backpropagation. Appendix C uses the same notation.

B.1 Denoise and Project Generated and Reference Paths

B.1.1 Generated Denoising and Projection

Equations (3) and (4) define the denoising and projection operations. Each Π_j^{guide} is the differentiable counterpart of a selected exact training projection Π_j^{train}, for $j = 1, \dots, J_{\mathrm{guide}}$. Projections in $\mathscr{P}_{\mathrm{val}}$ and $\mathscr{P}_{\mathrm{test}}$ are excluded from guidance. The state Jacobian of the generated projected value is

$$J_{t,m,j} = D_{\mathbf{x}_t} z_{jm}. \tag{23}$$

The generated values remain differentiable with respect to $\mathbf{x}_t$ and enter the score in their native units. Appendix D.3 gives the application-specific forms of the differentiable projections.

B.1.2 Fixed Reference Paths and Projections

The same guidance projections map the fixed reference paths to the values y_{jr} in Equation (4). These projected reference values and all associated reference statistics are computed once and detached. Their empirical CDF is

$$\widehat{F}_{\mathrm{ref},j}(v) = \frac{1}{R} \sum_{r=1}^{R} \mathbf{1}\{y_{jr} \le v\}. \tag{24}$$

This fixed branch joins the differentiable generated branch only when evaluating the TL–CRPS objective.

B.2 Construct the Soft Quantile and Tail Law

The generated pair $(\hat{v}_j, G_j)$ consists of the soft quantile boundary $\widehat{v}_j$ and the soft tail law G_j defined in Equations (5) and (6).

B.2.1 Soft Quantile

Adaptive smoothing scales. Because the projections can have different physical units and scales, let MAD denote the median absolute deviation and let $s_{\min} = 10^{-8}$. At every active reverse step, TAILDIFF computes

$$\begin{aligned} s_j &= \max\{1.4826\ \mathrm{MAD}(z_{j1:B}), s_{\min}\}, \\ \lambda_{q,j} &= \widetilde{\tau}_q / s_j, \qquad \lambda_{\mathrm{tl},j} = \widetilde{\tau}_{\mathrm{tl}} / s_j. \end{aligned} \tag{25}$$

The scale s_j is detached before evaluating the soft quantile and tail weights. The dimensionless parameters $\widetilde{\tau}_q$, for the soft quantile in Equation (5), and $\widetilde{\tau}_{\mathrm{tl}}$, for the tail memberships in Equation (6), therefore adapt the two logistic temperatures to the current generated batch while the score remains in the native units of z_{jm} and y_{jr}.

Forward solver. For each projection j, the soft quantile boundary $\widehat{v}_j$ approximates the root $\bar{v}_j$ of Equation (5). We compute its forward value by $K_q = 30$ bisection iterations, without gradient tracking, over the interval

$$\left[\min_m z_{jm} - \frac{12}{\lambda_{q,j}}, \max_m z_{jm} + \frac{12}{\lambda_{q,j}}\right].$$

The sigmoid arguments are clipped to $[-30, 30]$ during this numerical solve and when evaluating the soft-tail weights to avoid numerical overflow. Each iteration evaluates the clipped soft CDF $H^c_{B,j}$ of Equation (79) at the midpoint of the current bracket and moves the upper endpoint to the midpoint when this value is at least α, and the lower endpoint otherwise; the forward value $\widehat{v}_j$ is the midpoint of the final bracket.

Implicit Jacobian. Write $a_{q,jm} = \lambda_{q,j}(\widehat{v}_j - z_{jm})$ and distinguish the derivative used for the implicit quantile gradient from the ordinary derivative through a clipped sigmoid:

$$d_{\mathrm{ST}}(a) = \sigma'(\mathrm{clip}(a, -30, 30)), \qquad d_{\mathrm{clip}}(a) = \mathbf{1}\{|a| \le 30\}\, \sigma'(\mathrm{clip}(a, -30, 30)). \tag{26}$$

The first derivative in Equation (26) supplies the straight-through implicit Jacobian

$$w_{jm} = \frac{d_{\mathrm{ST}}(a_{q,jm})}{\sum_{\ell=1}^{B} d_{\mathrm{ST}}(a_{q,j\ell})}, \qquad w_{jm} \ge 0, \qquad \sum_m w_{jm} = 1. \tag{27}$$

Its denominator $M_j = \sum_{\ell=1}^{B} d_{\mathrm{ST}}(a_{q,j\ell})$, the derivative mass, satisfies $M_j \ge B\sigma'(30) > 0$. Thus, the forward pass uses the bisection root, while the backward pass uses the displayed implicit derivative. The second derivative, d_{clip}, is the ordinary autograd Jacobian along paths through the soft-tail weights defined below. The unclamped soft-CDF equation has a unique root.

B.2.2 Soft Tail Law

Equation (6) defines the soft memberships u_{jm}, their normalised weights π_{jm}, and the generated tail law G_j. The effective number of observations retained by this law is

$$n_{\mathrm{eff},j} = \sum_{m=1}^{B} u_{jm}. \tag{28}$$

This quantity is the denominator of the normalised weights in Equation (6). As $\lambda_{\mathrm{tl},j}$ increases, the weights concentrate on observations satisfying $z_{jm} \lesssim \widehat{v}_j$, and G_j approaches the corresponding hard empirical tail law.

Effective-tail activation. TAILDIFF activates projection j, placing it in $\mathcal{J}_t$, when

$$n_{\mathrm{eff},j} \ge \max\{8, \tfrac{1}{2}\alpha B\}. \tag{29}$$

This condition prevents guidance from using an under-populated soft tail. An inactive projection contributes the graph-preserving zero $0_{\mathrm{graph}} = z_{j\cdot}^{\mathsf{T}}\mathbf{0}$, where $z_{j\cdot} = (z_{j1}, \ldots, z_{jB})$.

B.3 TL–CRPS Tail-Law Objective

The fixed reference branch joins the differentiable generated branch at this stage. The projected reference values remain detached, whereas the generated pair $(\hat{v}_j, G_j)$ remains differentiable.

B.3.1 Weighted Empirical CRPS

For the generated tail law G_j, the empirical left-tail generator score is

$$A_j(u) = \sum_{m=1}^{B} \pi_{jm}|z_{jm} - u| - \frac{1}{2}\sum_{m=1}^{B}\sum_{\ell=1}^{B} \pi_{jm}\pi_{j\ell}|z_{jm} - z_{j\ell}| - 2u. \tag{30}$$

The first two terms are the weighted empirical CRPS of G_j, while the final linear term gives the reflected left-tail generator used by the tail-score construction. The function A_j is piecewise differentiable in the atoms and differentiable in the weights. The subgradient convention at ties is specified in Appendix C.3.1.

B.3.2 Reference-Conditional Score and Loss

Equation (7a) defines the score contribution S_{jr} for each fixed projected reference value y_{jr}, and Equation (7b) aggregates these contributions into the finite-batch loss. The first term scores projected reference values below the soft quantile boundary, while the second supplies the boundary correction. After averaging over the projected reference values, the latter has coefficient $\alpha - \widehat{F}_{\mathrm{ref},j}(\widehat{v}_j^{\mathrm{detach}})$ and therefore transports the soft quantile boundary towards the corresponding reference quantile.

B.3.3 Gradient Ownership and Detachment

The generated atoms z_{jm} remain differentiable, the soft quantile boundary $\widehat{v}_j$ uses the implicit Jacobian in Equation (27), and the tail weights u_{jm} and π_{jm} use ordinary autograd. The projected reference values, their statistics, the generated-batch MAD, the bisection control flow, and the indicator masks in Equation (7a) are detached. Detaching the masks is valid because the two score branches agree at $y_{jr} = \widehat{v}_j$ [Fissler et al., 2026]. These conventions leave three gradient channels: direct movement of the generated atoms, transport of the soft quantile boundary, and adjustment of the soft-tail weights.

B.4 Guided Reverse-Mean Update

Let $g_{t,m}$ be the restriction of the stacked gradient g_t to path m in the active guidance batch. Equations (8) and (9) define the corresponding update and reverse transition.

Scheduling. Guidance is applied when $t/(T-1) \le \tau_{\mathrm{start}}$ and $t \bmod M = 0$, where the integer $M \ge 1$ is the guidance stride. This schedule includes $t = 0$, where the guided mean is returned without adding Gaussian noise.

B.5 Complete Sampling Algorithm

Algorithm 1 collects the preceding definitions in their execution order. All projected reference values and associated statistics are precomputed, while the generated pair $(\hat{v}_j, G_j)$ and its loss are reconstructed for each active guidance batch.

B.6 Batching and Computational Cost

TAILDIFF returns N paths and evaluates its loss in consecutive batches of at most $B_{\max}$ paths. If N is not divisible by $B_{\max}$, the final guidance batch has realised size $B = N \bmod B_{\max}$. The primary synthetic and market comparisons use $N = B_{\max} = 1000$, so $B = 1000$ and all paths participate in one common empirical quantile and tail-law calculation.

For the reported experiments, the fixed reference paths are training trajectories, whose projected values are precomputed at cost $O(J_{\mathrm{guide}}R)$. With K_q bisection iterations, an active guidance step has cost

$$O(J_{\mathrm{guide}}BK_q) + O(J_{\mathrm{guide}}B\log B) + O\left(\sum_{j=1}^{J_{\mathrm{guide}}} R_{\mathrm{tail},j}B\right), \qquad R_{\mathrm{tail},j} = \#\{r : y_{jr} \le \widehat{v}_j^{\mathrm{detach}}\}. \tag{31}$$

Algorithm 1 TAILDIFF inference-time sampling

Input: pretrained $(\mu_\theta, \Sigma_\theta, \epsilon_\theta)$, projections $\{\Pi_j^{\text{guide}}\}_{j=1}^{J_{\text{guide}}}$, fixed reference paths $\{x_{\text{ref}}^{(r)}\}_{r=1}^{R}$, budget N, batch-size cap $B_{\max}$, and TAILDIFF parameters
compute and detach the projected reference values y_{jr}
sample $x_{T-1}^{(i)} \sim \mathcal{N}(0, I)$ for $i = 1, \dots, N$
for $t = T-1, \dots, 0$ **do**
    compute the unguided reverse mean $\mu_\theta(x_t^{(i)}, t)$ and covariance $\Sigma_\theta(t)$ for all i
    if $t/(T-1) \le \tau_{\text{start}}$ and $t \bmod M = 0$ **then**
        partition the trajectories into batches of at most $B_{\max}$ paths
        for each batch $\mathcal{I} = (i_1, \dots, i_B)$, where $B \le B_{\max}$ **do**
            use local index m for global path i_m, for $m = 1, \dots, B$
            compute $\widehat{x}_0(x_t^{(m)})$ using Equation (3)
            compute z_{jm} using Equation (4)
            compute the detached adaptive scales s_j using Equation (25)
            for $j = 1, \dots, J_{\text{guide}}$ **do**
                solve Equation (5) for $\widehat{v}_j$ and attach the Jacobian in Equation (27)
                compute G_j and $n_{\text{eff},j}$ using Equations (6) and (28)
                evaluate the effective-tail condition in Equation (29)
                compute A_j using Equation (30) and S_{jr} using Equation (7a)
            end for
            compute $\mathcal{L}_{\text{TailDiff}}$ using Equation (7b) and set $g_t \leftarrow \nabla_{\mathbf{x}_t}^{\text{AD}} \mathcal{L}_{\text{TailDiff}}$
            set $U_{t,m} = \gamma g_{t,m}$ and write $\widetilde{\mu}_{\theta,i_m} \leftarrow \mu_\theta(x_t^{(i_m)}, t) - U_{t,m}$ for $m = 1, \dots, B$
        end for
    else
        set $\widetilde{\mu}_{\theta,i} \leftarrow \mu_\theta(x_t^{(i)}, t)$ for all i
    end if
    if $t > 0$ **then**
        sample $x_{t-1}^{(i)} \sim \mathcal{N}(\widetilde{\mu}_{\theta,i}, \Sigma_\theta(t))$ for all i using Equation (9)
    else
        set $x_{\text{out}}^{(i)} \leftarrow \widetilde{\mu}_{\theta,i}$ for all i
    end if
end for
return $x_{\text{out}}^{(1:N)}$

Tail masking of the projected reference values and per-projection evaluation avoid materialising a dense $J_{\text{guide}} \times R \times B$ tensor.

Proposition 4 characterises the sampling error and smoothing bias of an auxiliary deterministic-scale, unclamped soft quantile. For the batch-adaptive, clamped boundary used in sampling, Proposition 5 gives the corresponding exact finite-error decomposition. The guidance batch must be large enough to populate the target tail, while its upper limit is set by memory. In the primary comparisons, we use $N = B = 1000$. For larger generation budgets, we process batches of at most 1000 and estimate the soft quantile boundary separately within each guidance batch.

The reported synthetic experiments use $\widetilde{\tau}_q = \widetilde{\tau}_{\text{tl}} = 20$, a guidance stride of one, and the training trajectories as the fixed reference paths.

C Theoretical Results and Proofs

This appendix follows the three parts of the analysis in Section 5. We first establish the population target and compare it with VaR–ES matching. We then analyse the declared finite-batch backward field and give an exact decomposition linking the finite score to the ideal target. Finally, we derive the terminal perturbation bound from the realised reverse-mean shifts. The finite-batch objective and sampling procedure are defined in Appendix B.

C.1 Notation and Preliminary Identities

Population notation. For the population analysis, let $P \in \mathcal{M}^1$ be a reference law, let $q = Q_\alpha(P)$, and write $H = P^\alpha$ for its left-α tail law. For any tail law $G \in \mathcal{M}^1$, define

$$A_G(u) = \mathrm{CRPS}(G, u) - 2u. \tag{32}$$

The corresponding free-pair population score is

$$L_P(v, G) = \mathbb{E}_{Y\sim P}\mathcal{S}_{\text{TL–CRPS}}(v, G; Y). \tag{33}$$

Finite-batch notation. For the finite analysis, we use the notation of Appendix B, including the generated values z_{jm} and projected reference values y_{jr}, generated threshold $\widehat{v}_j$, soft tail law G_j, state Jacobian $J_{t,m,j}$, guidance gradient g_t, and realised update U_t. Proposition-specific quantities are defined with the result in which they are used. Throughout the finite analysis, D^{AD} and ∇^{AD} denote derivatives returned by the declared backward program of Section 4.4. This program assigns the implicit root Jacobian in Equation (27) while treating the reference branch, generated-batch MAD, bisection decisions, indicator masks, and activation decisions as fixed. Thus $g_t = \nabla^{\mathrm{AD}}_{\mathbf{x}_t}\mathcal{L}_{\mathrm{TailDiff}}$. The active branch of projection j at step t is the case in which j passes the effective-tail condition (29), with all of these decisions held fixed.

Sampler-perturbation notation. Each generated path is a real matrix of series by time steps; write $\mathcal{X}$ for this space, with the Frobenius norm $\|\cdot\|_F$. For a random guidance batch $\mathbf{X}_t = (X_t^{(1)}, \ldots, X_t^{(B)})$ in $\mathcal{X}^B$, write $\mathbf{U}_t(\mathbf{X}_t) = (U_{t,1}, \ldots, U_{t,B})$ and $\|\mathbf{x}\|_B^2 = \sum_{m=1}^B \|x^{(m)}\|_F^2$. The symbols $\mathbf{K}_t^{(g)}$ and $\mathbf{K}_t^{(0)}$ denote the guided and unguided batch kernels, $\boldsymbol{\nu}_t$ denotes the law of $\mathbf{X}_t$ under the guided sampler, and $p_{0,i}^{(g)}$ and $p_{0,i}^{(0)}$ denote the corresponding guided and unguided laws of the returned sample $x_{\mathrm{out}}^{(i)}$ of generated path i. For a law ν and kernels $\mathbf{K}$ and $\mathbf{L}$, $(\nu\mathbf{K})(A) = \int \mathbf{K}(A \mid \mathbf{x})\,\nu(d\mathbf{x})$, the product $\mathbf{KL}$ applies $\mathbf{K}$ first and then $\mathbf{L}$, and $\mathrm{Id}(\cdot \mid \mathbf{x}) = \delta_{\mathbf{x}}$. For a normed space E, $\mathcal{P}_1(E)$ is the set of Borel probability laws on E with a finite first moment, so that $\mathcal{M}^1 = \mathcal{P}_1(\mathbb{R})$, and $W_1(\mathsf{P}, \mathsf{Q}) = \inf_\pi \int \|\mathbf{x} - \mathbf{y}\|\,\pi(d\mathbf{x}, d\mathbf{y})$ is the Wasserstein-1 distance, the infimum running over couplings π of P and Q, with $\|\cdot\|_F$ on $\mathcal{X}$ and $\|\cdot\|_B$ on $\mathcal{X}^B$. We use L_t^{rev} for the propagation constant of the remaining unguided reverse chain.

C.2 Population Identification

C.2.1 Proof of Theorem 1

Consider a candidate law $F \in \mathcal{M}^1$ and a reference law $F_{\mathrm{ref}} \in \mathcal{M}^1$. Their left-$\alpha$ tail laws are denoted by F^α and F_{ref}^α. The candidate tail law induces the threshold

$$v_F^\alpha = Q_\alpha(F) = \inf\{x : F^\alpha(x) = 1\}, \tag{34}$$

where the second equality uses the fractional-boundary convention (14). Let $L_{F_{\mathrm{ref}}}(v, G) = \mathbb{E}_{Y\sim F_{\mathrm{ref}}}\mathcal{S}_{\text{TL–CRPS}}(v, G; Y)$ denote the population score in Equation (21), and define the score at the pair induced by F as $\mathcal{L}(F) = L_{F_{\mathrm{ref}}}(v_F^\alpha, F^\alpha)$.

Proof. We prove Equation (11) by separating the threshold discrepancy from the CRPS discrepancy between the candidate and reference tail laws. Put $P = F_{\mathrm{ref}}$, $q = Q_\alpha(P)$, $H = P^\alpha$, and $A_G(u) = \mathrm{CRPS}(G, u) - 2u$ as in Equation (19). Consider an arbitrary finite threshold v and tail law $G \in \mathcal{M}^1$. By Equation (21), the expected score of this threshold–tail pair is

$$L_P(v, G) = \int_{(-\infty, v]} A_G(y)\,dP(y) + [\alpha - P(v)]A_G(v), \tag{35}$$

which holds for every $P \in \mathcal{M}^1$ because $\mathbb{E}_{Y\sim P}\mathbf{1}\{Y > v\} = 1 - P(v)$. Since $\partial_u \mathrm{CRPS}(G, u) = 2G(u) - 1$ away from atoms,

$$A_G'(u) = 2G(u) - 3 \in [-3, -1], \tag{36}$$

with the corresponding one-sided subgradients at atoms. Thus A_G is strictly decreasing, which is the left-tail sign condition of the construction of Fissler et al. [2026].

Since the theorem assumes that P has a continuous CDF, $P(q) = \alpha$. For $v > q$ and $v < q$ respectively, direct subtraction in Equation (35) gives

$$L_P(v, G) - L_P(q, G) = \int_{(q,v]} \big[A_G(y) - A_G(v)\big]\, dP(y) \geq 0, \tag{37}$$

$$L_P(v, G) - L_P(q, G) = \int_{(v,q]} \big[A_G(v) - A_G(y)\big]\, dP(y) \geq 0, \tag{38}$$

both non-negative because A_G is decreasing. The strict increase of P in a neighbourhood of q is not needed for the equality statement for the induced objective $\mathcal{L}$ proved below. It does make the threshold comparison strict for every $v \neq q$: the open interval between v and q meets that neighbourhood in a non-empty open subinterval of positive P-mass, and the integrands above are strictly positive there because A_G is strictly decreasing.

At $v = q$, the second term of Equation (35) vanishes, so

$$L_P(q, G) = \alpha\, \mathbb{E}_{Y\sim H}\big[\mathrm{CRPS}(G, Y) - 2Y\big], \tag{39}$$

in which the $-2Y$ term does not depend on G. For every x, $\mathbb{E}_H\big[\big(G(x) - \mathbf{1}\{Y \leq x\}\big)^2\big] = \big(G(x) - H(x)\big)^2 + H(x)\big(1 - H(x)\big)$. Using the CRPS decomposition, which follows from Equation (18), this identity and Tonelli's theorem, yields

$$\begin{aligned} L_P(q, G) - L_P(q, H) &= \alpha\big\{\mathbb{E}_H\, \mathrm{CRPS}(G, Y) - \mathbb{E}_H\, \mathrm{CRPS}(H, Y)\big\} \\ &= \alpha \int_{\mathbb{R}} \big(G(x) - H(x)\big)^2\, dx \geq 0. \end{aligned} \tag{40}$$

Equality holds if and only if $G = H$. Thus (q, H) is a minimising threshold–tail pair, and at the threshold q the tail law is uniquely determined.

For a candidate law F, substitute the induced pair (v_F^α, F^α). To make the boundary convention explicit, if $v = v_F^\alpha$ then for every Borel set A,

$$F^\alpha(A) = \frac{1}{\alpha}\left(F\big(A \cap (-\infty, v)\big) + \big[\alpha - F(v-)\big]\mathbf{1}\{v \in A\}\right), \tag{41}$$

so that the required fraction of any atom at v is included. Under this convention, $v_F^\alpha = \inf\{x : F^\alpha(x) = 1\}$. Hence $F^\alpha = H = P^\alpha$ implies $v_F^\alpha = q$. The preceding threshold and CRPS comparisons show that (q, H) is the global minimiser over free threshold–tail pairs, and it is the pair induced by $F = F_{\mathrm{ref}}$. Since $L_P(q, H) = \mathcal{L}(F_{\mathrm{ref}})$, adding and subtracting $L_P(q, F^\alpha)$ gives, for every candidate law F,

$$\begin{aligned} \mathcal{L}(F) - \mathcal{L}(F_{\mathrm{ref}}) &= \big[L_P(v_F^\alpha, F^\alpha) - L_P(q, F^\alpha)\big] + \big[L_P(q, F^\alpha) - L_P(q, H)\big] \\ &= \mathcal{B}_\alpha(F; F_{\mathrm{ref}}) + \alpha \int_{\mathbb{R}} \big(F^\alpha(x) - F_{\mathrm{ref}}^\alpha(x)\big)^2\, dx, \end{aligned} \tag{42}$$

where the first bracketed term is $\mathcal{B}_\alpha(F; F_{\mathrm{ref}})$ by definition and the second is given by Equation (40) evaluated at $G = F^\alpha$. This is the decomposition in Equation (11). Both bracketed terms are non-negative, so $\mathcal{L}(F) \geq \mathcal{L}(F_{\mathrm{ref}})$. The first vanishes exactly when $v_F^\alpha = q$, by the strict threshold comparison, and the second vanishes only when $F^\alpha = H$. Hence $\mathcal{L}(F) = \mathcal{L}(F_{\mathrm{ref}})$ forces $F^\alpha = H$; conversely, when $F^\alpha = H$, the endpoint identity gives $v_F^\alpha = q$, so both bracketed terms vanish. Therefore the minimisers are exactly the candidate laws whose left-α tail law equals F_{ref}^α.

The argument is per scalar guidance projection. It identifies neither the candidate bulk above the left-α tail nor compatibility across a finite projection family. The finite surrogate and reverse transition are analysed separately below. □

C.2.2 Interpretation and Scope

The induced quantile is not a free argument of the objective. By Equation (34), the candidate supplies it. The Cramér-distance reduction therefore applies after the candidate threshold aligns with the reference quantile, at which point the excess score measures tail shape.

The minimiser is a tail equivalence class rather than a complete law. Matching the left-α tail leaves the bulk and cross-projection dependence unrestricted.

C.2.3 Relation to the Tail-Score Construction

The population target specialises the construction of Fissler et al. [2026, Sec. 5.3] to the reflected left-tail convention of Fissler et al. [2026, Sec. 6]. Their score treats the threshold and tail law as independently specified candidate quantities. TAILDIFF restricts those quantities to the induced pairs $\{(Q_\alpha(F), F^\alpha) : F \in \mathcal{M}^1\}$. The source states the corresponding left-tail point-forecast result [Fissler et al., 2026, Thm. 31] and notes that the distributional construction follows by the same reflection. On the induced class, the excess score at an aligned quantile becomes the Cramér distance in Equation (11), which permits a direct comparison with VaR–ES matching.

C.2.4 Identification Beyond VaR–ES Matching

Theorem 1 targets the whole left-α law, where the established tail objective for this benchmark targets two of its functionals: the score of Cont et al. [2025, Eq. (7)] depends on the candidate only through the pair $(\mathrm{VaR}_\alpha, \mathrm{ES}_\alpha)$ shown jointly elicitable by Fissler and Ziegel [2016]. TL–CRPS is the finer objective, and the discrimination it adds is bounded. Two laws agreeing on both functionals need not agree on their tail laws, and by Theorem 1 the score then separates them by α times their squared Cramér distance; the uniform bound in Proposition 1 is $\mathbb{E}[(\mathrm{VaR}_\alpha - Y)^+] = \alpha(\mathrm{VaR}_\alpha - \mathrm{ES}_\alpha)$, the expected amount by which Y falls below Value-at-Risk. This bound sets the scale of what the point pair leaves undetermined, in the units of the guidance objective itself. Proposition 1 below states this formally, and Remark 1 exhibits a pair of ordinary tail shapes, a uniform and a triangle, that the point pair cannot separate.

Proposition 1 (How much the point pair leaves free)**.** *Let $F_{\mathrm{ref}} \in \mathcal{M}^1$ have a continuous CDF strictly increasing in a neighbourhood of $v = Q_\alpha(F_{\mathrm{ref}})$, and write $\Delta_\alpha(F_{\mathrm{ref}}) = \mathrm{VaR}_\alpha(F_{\mathrm{ref}}) - \mathrm{ES}_\alpha(F_{\mathrm{ref}})$. Let $F \in \mathcal{M}^1$ satisfy $\mathrm{VaR}_\alpha(F) = \mathrm{VaR}_\alpha(F_{\mathrm{ref}})$ and $\mathrm{ES}_\alpha(F) = \mathrm{ES}_\alpha(F_{\mathrm{ref}})$, with ES_α the quantile integral of Appendix A.1, and write $\mathfrak{F}_\alpha(F_{\mathrm{ref}})$ for the family of all such F. Then, with $\mathcal{L}$ as in Theorem 1,*

$$\begin{aligned} 0 \;\le\; \mathcal{L}(F) - \mathcal{L}(F_{\mathrm{ref}}) &= \alpha \int_{\mathbb{R}} \big(F^\alpha(x) - F^\alpha_{\mathrm{ref}}(x)\big)^2 dx \\ &\le \alpha\, \Delta_\alpha(F_{\mathrm{ref}}) = \mathbb{E}_{Y\sim F_{\mathrm{ref}}}\big[(\mathrm{VaR}_\alpha(F_{\mathrm{ref}}) - Y)^+\big], \end{aligned} \tag{43}$$

where $a^+ = \max\{a, 0\}$. The displayed upper bound is a diameter bound over $\mathfrak{F}_\alpha(F_{\mathrm{ref}})$, and its constant is sharp uniformly over admissible pairs, namely pairs (F_{ref}, F) with F_{ref} as above and $F \in \mathfrak{F}_\alpha(F_{\mathrm{ref}})$. Any score depending on the candidate only through the pair $(\mathrm{VaR}_\alpha, \mathrm{ES}_\alpha)$ has the same value at every member of $\mathfrak{F}_\alpha(F_{\mathrm{ref}})$ for each fixed observation.

Proof. Write $v = \mathrm{VaR}_\alpha(F_{\mathrm{ref}})$ and $e = \mathrm{ES}_\alpha(F_{\mathrm{ref}})$. Strict monotonicity of F_{ref} immediately below v places positive F^α_{ref}-mass at points strictly below v, so $e < v$ and $\Delta_\alpha(F_{\mathrm{ref}}) = v - e > 0$; the bound below does not use this strict inequality. Since F and F_{ref} share the lower quantile, $v^\alpha_F = Q_\alpha(F) = Q_\alpha(F_{\mathrm{ref}}) = v$, the threshold term of Equation (11) vanishes and Theorem 1 gives the equality in Equation (43).

For the bound, F^α_{ref} and F^α are CDFs supported on $(-\infty, v]$ with mean e, so

$$\int_{-\infty}^{v} F^\alpha_{\mathrm{ref}}(x)\, dx = \int_{-\infty}^{v} F^\alpha(x)\, dx = v - e = \Delta_\alpha(F_{\mathrm{ref}}). \tag{44}$$

Set $f(t) = F^\alpha_{\mathrm{ref}}(v - t)$ and $g(t) = F^\alpha(v - t)$ for $t \ge 0$. These functions are non-increasing, take values in $[0, 1]$, and both integrate to $\Delta_\alpha(F_{\mathrm{ref}})$. For $d = f - g$, let

$$A = \int_0^\infty d_+(t)\, dt = \int_0^\infty d_-(t)\, dt, \qquad a = \|d_+\|_\infty, \qquad b = \|d_-\|_\infty. \tag{45}$$

The equality defining A follows from $\int f = \int g$. To verify $a + b \le 1$, fix $\varepsilon > 0$ and choose $s, t \ge 0$ with $d(s) > a - \varepsilon$ and $-d(t) > b - \varepsilon$. If $s \le t$, then $d(s) - d(t) \le f(s) - f(t) \le 1$; if $t < s$, then $d(s) - d(t) \le g(t) - g(s) \le 1$. Letting $\varepsilon \downarrow 0$ gives the claim. Therefore

$$\int_{\mathbb{R}} (F^\alpha_{\mathrm{ref}} - F^\alpha)^2 dx = \int_0^\infty d(t)^2 dt \le aA + bA \le A \le \int_0^\infty f(t)\, dt = \Delta_\alpha(F_{\mathrm{ref}}), \tag{46}$$

where the first equality also uses $F^{\alpha}_{\mathrm{ref}} = F^{\alpha} = 1$ above v. The argument uses only that both tail laws are CDFs supported on $(-\infty, v]$ with mean e, so it applies to any two members of $\mathfrak{F}_{\alpha}(F_{\mathrm{ref}})$. The constant one is optimal uniformly over admissible pairs. For a fixed $\Delta > 0$, the limiting transforms

$$f_0 = \mathbf{1}_{[0,\Delta]}, \qquad g_{\varepsilon}(0) = 1, \qquad g_{\varepsilon}(t) = \varepsilon \mathbf{1}_{(0,\Delta/\varepsilon]}(t), \quad t > 0, \tag{47}$$

both integrate to Δ and satisfy $\int_0^{\infty} (f_0 - g_{\varepsilon})^2 = \Delta(1-\varepsilon)$. They correspond to a point mass at $v - \Delta$ and to masses $1 - \varepsilon$ at v and ε at $v - \Delta/\varepsilon$. To meet the assumptions on F_{ref}, replace its point mass by an atomless law concentrated in a shrinking interval about $v - \Delta$, add a vanishing continuous component with positive density immediately below v, and recentre the first component to retain mean $v - \Delta$. Appending an atomless mass $1 - \alpha$ above v, with positive density immediately to its right, gives full laws with common VaR_{α} and ES_{α}. Their transforms converge to f_0 in L^2, so first taking this approximation limit and then $\varepsilon \downarrow 0$ proves uniform sharpness.

For the final identity, continuity of F_{ref} gives $F_{\mathrm{ref}}(v) = \alpha$ and $\int_{(-\infty,v]} y \, dF_{\mathrm{ref}}(y) = \alpha e$, hence

$$\mathbb{E}_{Y \sim F_{\mathrm{ref}}}\big[(v - Y)^+\big] = v F_{\mathrm{ref}}(v) - \int_{(-\infty,v]} y \, dF_{\mathrm{ref}}(y) = \alpha(v - e) = \alpha \Delta_{\alpha}(F_{\mathrm{ref}}). \tag{48}$$

Finally, any score depending on the candidate only through $(\mathrm{VaR}_{\alpha}, \mathrm{ES}_{\alpha})$ has the same value at every member of $\mathfrak{F}_{\alpha}(F_{\mathrm{ref}})$ for each fixed observation, since those two functionals are held fixed. □

C.2.5 A Concrete Indistinguishable Pair

Remark 1 (A concrete pair, and how far it sits from the bound). Fix $\alpha \in (0,1)$ and let both laws place mass α on $[-3,-2]$ and mass $1-\alpha$ uniformly on $[-2,0]$, so that $\mathrm{VaR}_{\alpha} = -2$ for both. Let F_{ref} have a uniform tail on $[-3,-2]$ and let F have the symmetric triangular tail on $[-3,-2]$ with mode $-5/2$. Both tail laws are symmetric about $-5/2$, so $\mathrm{ES}_{\alpha} = -5/2$ for both and $\Delta_{\alpha} = 1/2$. In the variable $t = x + 3 \in [0,1]$,

$$F^{\alpha}_{\mathrm{ref}} = t, \qquad F^{\alpha} = 2t^2 \mathbf{1}\{t \le \tfrac{1}{2}\} + \big(1 - 2(1-t)^2\big)\mathbf{1}\{t > \tfrac{1}{2}\}, \tag{49}$$

and by symmetry of the integrand about $t = \frac{1}{2}$,

$$\int_{\mathbb{R}} \big(F^{\alpha} - F^{\alpha}_{\mathrm{ref}}\big)^2 dx = 2 \int_0^{1/2} \big(4t^4 - 4t^3 + t^2\big) dt = \frac{1}{120}, \tag{50}$$

so $\mathcal{L}(F) - \mathcal{L}(F_{\mathrm{ref}}) = \alpha/120$. At $\alpha = 0.05$ this is 4.1667×10^{-4}, and it uses 1.67% of the sharp uniform bound $\alpha\Delta_{\alpha} = 2.5 \times 10^{-2}$ in Equation (43). Two tail laws as ordinary as a uniform and a triangle are thus indistinguishable to the point pair while remaining well inside the maximal tail-shape separation compatible with that pair.

C.3 Finite Implementation

C.3.1 Gradient Decomposition of the Finite Objective

For completeness we state the decomposition summarised in Section 5, and prove it by exhibiting the three channels explicitly.

Proposition 2 (Three-channel TAILDIFF backward field)**.** *For a projection on a fixed active branch, let $\ell_j = R^{-1} \sum_r S_{jr}$, $\mathbf{x}_t = \mathrm{vec}(x_t^{(1:B)})$, and $J_{t,m,j} = D_{\mathbf{x}_t} z_{jm}$. Write $\widetilde{\ell}_j(z_j, \pi_j, v)$ for the same score with its atoms, weights, and threshold treated as independent arguments while the detached masks are fixed. The backward graph is the composition*

$$\ell_j(z_j) = \widetilde{\ell}_j\big(z_j, \pi_j(z_j, \widehat{v}_j(z_j)), \widehat{v}_j(z_j)\big), \tag{51}$$

where $D^{\mathrm{AD}}_{z_{jm}} \widehat{v}_j := w_{jm}$ is the supplied root Jacobian in Equation (27)*. Holding the active branch and masks fixed, the derivative returned by this backward graph is*

$$D^{\mathrm{AD}}_{z_{jm}} \ell_j = D^{\mathrm{dir}}_{jm} + \left.\frac{\partial \widetilde{\ell}_j}{\partial \widehat{v}_j}\right|_{z,\pi} w_{jm} + \sum_{k=1}^{B} \left.\frac{\partial \widetilde{\ell}_j}{\partial \pi_{jk}}\right|_{z,\widehat{v}} \left(\left.\frac{\partial \pi_{jk}}{\partial z_{jm}}\right|_{\widehat{v}} + \frac{\partial \pi_{jk}}{\partial \widehat{v}_j} w_{jm} \right), \tag{52}$$

the atom, quantile-transport, and truncation-weight channels respectively. Each partial derivative of $\widetilde{\ell}_j$ is taken in its three independent arguments; in particular, $D^{\mathrm{dir}}_{jm} = \partial\widetilde{\ell}_j/\partial z_{jm}$ with $(\pi_j, \widehat{v}_j)$ held fixed, and

$$\left.\frac{\partial\widetilde{\ell}_j}{\partial\widehat{v}_j}\right|_{z,\pi} = c_j A'_j(\widehat{v}_j), \qquad c_j = \alpha - \widehat{F}_{\mathrm{ref},j}(\widehat{v}_j^{\mathrm{detach}}), \qquad A'_j(t) = 2G_j(t) - 3 \tag{53}$$

away from ties. The state field is $\nabla^{\mathrm{AD}}_{\mathbf{x}_t}\ell_j = \sum_m J^{\mathsf{T}}_{t,m,j} D^{\mathrm{AD}}_{z_{jm}}\ell_j$, and the sum reduction gives $\nabla^{\mathrm{AD}}_{\mathbf{x}_t}\mathcal{L}_{\mathrm{TailDiff}} = \sum_{j\in\mathcal{J}_t}\nabla^{\mathrm{AD}}_{\mathbf{x}_t}\ell_j$ over the active projections. Let $\bar{a}_{q,jm} = \lambda_{q,j}(\bar{v}_j - z_{jm})$ and let $\bar{w}_j$ be the weights of Equation (27) computed at the exact root $\bar{v}_j$ of Equation (5) in place of $\widehat{v}_j$. If no $\bar{a}_{q,jm}$ is clipped, $\bar{w}_{jm} = \partial\bar{v}_j/\partial z_{jm}$, the implicit-function derivative at fixed $\lambda_{q,j}$; otherwise $\bar{w}_j$ is a straight-through surrogate with $\|\bar{w}_j - \partial\bar{v}_j/\partial z_j\|_1 \le 2B\sigma'(30)/\sum_\ell d_{\mathrm{ST}}(\bar{a}_{q,j\ell})$. In the finite program, the weights w_j computed at $\widehat{v}_j$ define the backward derivative assigned to the detached bisection value.

Proof. For this fixed projection, write $\lambda_q = \lambda_{q,j}$ and $\lambda_{\mathrm{tl}} = \lambda_{\mathrm{tl},j}$. Write $b_{jr} = \mathbf{1}\{y_{jr} \le \widehat{v}_j^{\mathrm{detach}}\}$ and $h_{jr} = \mathbf{1}\{y_{jr} > \widehat{v}_j^{\mathrm{detach}}\} - (1-\alpha)$, so that $\ell_j = R^{-1}\sum_r\{b_{jr}A_j(y_{jr}) + h_{jr}A_j(\widehat{v}_j)\}$, and set $\widehat{p}_j = R^{-1}\sum_r b_{jr}$, $c_j = R^{-1}\sum_r h_{jr} = \alpha - \widehat{p}_j$. Let u'_{jm} be the scalar derivative returned by automatic differentiation for $a \mapsto \sigma(\mathrm{clip}(a, -30, 30))$ at $a = \lambda_{\mathrm{tl}}(\widehat{v}_j - z_{jm})$. Then

$$du_{jm} = \lambda_{\mathrm{tl}}u'_{jm}(d\widehat{v}_j - dz_{jm}), \qquad d\pi_{jm} = \frac{n_{\mathrm{eff},j}\,du_{jm} - u_{jm}\sum_k du_{jk}}{n^2_{\mathrm{eff},j}}, \tag{54}$$

$$\begin{aligned}\left.\frac{\partial A_j(t)}{\partial z_{jm}}\right|_{\pi,t} &= \pi_{jm}\,\mathrm{sgn}(z_{jm} - t) - \pi_{jm}\sum_\ell \pi_{j\ell}\,\mathrm{sgn}(z_{jm} - z_{j\ell}),\\ \frac{\partial A_j(t)}{\partial \pi_{jk}} &= |z_{jk} - t| - \sum_\ell \pi_{j\ell}|z_{jk} - z_{j\ell}|,\end{aligned} \tag{55}$$

$$\left.\frac{\partial\pi_{jk}}{\partial z_{jm}}\right|_{\widehat{v}_j} = -\frac{\lambda_{\mathrm{tl}}u'_{jm}}{n_{\mathrm{eff},j}}(\mathbf{1}\{k=m\} - \pi_{jk}), \qquad \frac{\partial\pi_{jk}}{\partial\widehat{v}_j} = \frac{\lambda_{\mathrm{tl}}u'_{jk} - \pi_{jk}\sum_\ell\lambda_{\mathrm{tl}}u'_{j\ell}}{n_{\mathrm{eff},j}}. \tag{56}$$

At a self-pair (m,m) the pairwise sign derivative is zero. For tied distinct atoms, use antisymmetric selections $\xi_{m\ell} = -\xi_{\ell m} \in [-1,1]$. At an atom tied with the threshold, use the same selection $\xi_{jm} \in [-1,1]$ in the direct atom and threshold derivatives, with $\xi_{jm} = \mathrm{sgn}(z_{jm} - \widehat{v}_j)$ off ties, so that $A'_j(\widehat{v}_j) = -2 - \sum_m \pi_{jm}\xi_{jm}$. Setting tied signs to zero gives the midpoint subgradient. Applying the chain rule to the declared backward graph produces the three channels, and composition with $J_{t,m,j}$ gives the state field. Away from a clamp boundary, $u'_{jm} = d_{\mathrm{clip}}(\lambda_{\mathrm{tl}}(\widehat{v}_j - z_{jm}))$. Proposition 3 uses only $0 \le u'_{jm} \le 1/4$, which also covers the boundary convention selected by the automatic differentiation program.

For the claim about the exact root, implicit differentiation of Equation (5) at fixed λ_q gives $\partial\bar{v}_j/\partial z_{jm} = \sigma'(\bar{a}_{q,jm})/\sum_\ell\sigma'(\bar{a}_{q,j\ell})$, which equals $\bar{w}_{jm}$ when no argument is clipped, since $d_{\mathrm{ST}} = \sigma'$ on $[-30,30]$. In general, since $d_{\mathrm{ST}} \ge \sigma'$, adding and subtracting $\sigma'(\bar{a}_{q,jm})/\sum_\ell d_{\mathrm{ST}}(\bar{a}_{q,j\ell})$ in each component gives

$$\|\bar{w}_j - \partial\bar{v}_j/\partial z_j\|_1 \le \frac{2\sum_m\{d_{\mathrm{ST}}(\bar{a}_{q,jm}) - \sigma'(\bar{a}_{q,jm})\}}{\sum_\ell d_{\mathrm{ST}}(\bar{a}_{q,j\ell})},$$

and each summand in the numerator lies in $[0, \sigma'(30)]$. □

C.3.2 Conditional Local Gradient Bound

For completeness, we state the local bound underlying the finite-implementation summary in Section 5.

Proposition 3 (Conditional local gradient bound)**.** *Let $\mathbf{z}_j = (z_{j1}, \dots, z_{jB})$ and let $\mathscr{J}_{t,j} = D_{\mathbf{x}_t}\mathbf{z}_j$ be the composed Jacobian from the stacked state $\mathbf{x}_t = \mathrm{vec}(x_t^{(1:B)})$ to the projected batch, with*

$\|\mathscr{J}_{t,j}\|_{2\to 2} \le L_{\mathscr{J},j}$. *Define*

$$q_j = \left(D^{\mathrm{AD}}_{z_{j1}}\ell_j, \dots, D^{\mathrm{AD}}_{z_{jB}}\ell_j\right), \qquad r_j = \left.\frac{\partial\widetilde{\ell}_j}{\partial\pi_j}\right|_{z,\widehat{v}},$$
$$u'_{jm} = D^{\mathrm{AD}}_a[\sigma(\mathrm{clip}(a,-30,30))]\big|_{a=\lambda_{\mathrm{tl},j}(\widehat{v}_j - z_{jm})}, \qquad \eta_{w,j} = \|w_j\|_2, \tag{57}$$

and set

$$b_{jr} = \mathbf{1}\{y_{jr} \le \widehat{v}_j^{\mathrm{detach}}\}, \qquad h_{jr} = \mathbf{1}\{y_{jr} > \widehat{v}_j^{\mathrm{detach}}\} - (1-\alpha), \tag{58}$$

$$\widehat{p}_j = \frac{1}{R}\sum_r b_{jr}, \quad \bar{c}_j = \frac{1}{R}\sum_r |h_{jr}|, \quad c_j = \alpha - \widehat{p}_j, \quad C_{a,j} = 2(\widehat{p}_j + \bar{c}_j), \quad C_{q,j} = 3|c_j|. \tag{59}$$

For fixed j, write $\lambda_q = \lambda_{q,j}$ and $\lambda_{\mathrm{tl}} = \lambda_{\mathrm{tl},j}$. Let $D_j = \max_m |z_{jm}|$, and let

$$M_{\mathrm{ref},j} = \frac{1}{R}\sum_{r: y_{jr} \le \widehat{v}_j^{\mathrm{detach}}} |y_{jr}| \tag{60}$$

be the empirical reference-tail first moment. Holding the active branch and masks fixed, the sum of the three channels in Proposition 2 satisfies

$$\|q_j\|_2 \le \frac{C_{a,j}}{\sqrt{n_{\mathrm{eff},j}}} + C_{q,j}\eta_{w,j} + \|r_j\|_2 \frac{\lambda_{\mathrm{tl}}}{n_{\mathrm{eff},j}} \left(\left\|\mathrm{diag}(u'_j) - \pi_j(u'_j)^{\mathsf{T}}\right\|_{2\to 2} + \left\|u'_j - \pi_j \mathbf{1}^{\mathsf{T}} u'_j\right\|_2 \eta_{w,j}\right). \tag{61}$$

The terms left explicit in Equation (61) *obey*

$$\left\|\mathrm{diag}(u'_j) - \pi_j(u'_j)^{\mathsf{T}}\right\|_{2\to 2} \le \frac{1}{4} + \frac{\sqrt{B}}{4\sqrt{n_{\mathrm{eff},j}}},$$
$$\left\|u'_j - \pi_j \mathbf{1}^{\mathsf{T}} u'_j\right\|_2 \le \frac{\sqrt{B}}{4} + \frac{B}{4\sqrt{n_{\mathrm{eff},j}}}, \tag{62}$$

and the numerical bisection bracket gives

$$|\widehat{v}_j| \le D_j + \frac{12}{\lambda_q},$$
$$\|r_j\|_2 \le \sqrt{B}\left[M_{\mathrm{ref},j} + 3D_j\widehat{p}_j + \bar{c}_j\left(4D_j + \frac{12}{\lambda_q}\right)\right]. \tag{63}$$

Recall that $\mathcal{J}_t \subseteq \{1,\dots,J_{\mathrm{guide}}\}$ is the set of projections passing the effective-tail condition at step t. For the reported sum reduction,

$$\left\|\nabla^{\mathrm{AD}}_{\mathbf{x}_t}\mathcal{L}_{\mathrm{TailDiff}}\right\|_2 \le \sum_{j\in\mathcal{J}_t} L_{\mathscr{J},j}\,\|q_j\|_2. \tag{64}$$

The bound is local and data dependent. By Equations (8) *and* (64)*, the applied update satisfies*

$$\|\mathbf{U}_t\|_B = \gamma\left\|\nabla^{\mathrm{AD}}_{\mathbf{x}_t}\mathcal{L}_{\mathrm{TailDiff}}\right\|_2 \le \gamma\sum_{j\in\mathcal{J}_t} L_{\mathscr{J},j}\,\|q_j\|_2. \tag{65}$$

In particular, the update-integrability condition of Proposition 6 holds whenever the right-hand side of Equation (65) *is integrable under $\boldsymbol{\nu}_t$. Since $\sum_m w_{jm} = 1$ by Equation* (27)*, $\eta_{w,j} = \mathrm{ESS}(w_j)^{-1/2}$, where $\mathrm{ESS}(w) = (\sum_m w_m)^2/\sum_m w_m^2$ is Kish's effective sample size.*

Proof. Hold the active branch and masks fixed. Write $q_j = q_j^{\mathrm{atom}} + q_j^{\mathrm{quant}} + q_j^{\mathrm{wt}}$ for the three channels of Proposition 2, with $q_{jm}^{\mathrm{atom}} = D_{jm}^{\mathrm{dir}}$, $q_{jm}^{\mathrm{quant}} = c_j A'_j(\widehat{v}_j)w_{jm}$ and $q_{jm}^{\mathrm{wt}} = \sum_k r_{jk}(\partial\pi_{jk}/\partial z_{jm} + w_{jm}\,\partial\pi_{jk}/\partial\widehat{v}_j)$. Since $|\partial A_j(t)/\partial z_{jm}| \le 2\pi_{jm}$ by Equation (55), the atom and quantile channels satisfy

$$|q_{jm}^{\mathrm{atom}}| \le 2\pi_{jm}R^{-1}\sum_r (b_{jr} + |h_{jr}|) = C_{a,j}\pi_{jm},$$
$$|q_{jm}^{\mathrm{quant}}| = |c_j A'_j(\widehat{v}_j)w_{jm}| \le 3|c_j|w_{jm} = C_{q,j}w_{jm}. \tag{66}$$

Since $C_{q,j} = 3|c_j| \le 3$ and $\|\pi_j\|_2 \le n_{\mathrm{eff},j}^{-1/2}$,

$$\|q_j^{\mathrm{atom}}\|_2 \le \frac{C_{a,j}}{\sqrt{n_{\mathrm{eff},j}}}, \qquad \|q_j^{\mathrm{quant}}\|_2 \le C_{q,j}\eta_{w,j}. \tag{67}$$

Differentiating $\pi_j = u_j/n_{\mathrm{eff},j}$ gives the exact normalised-weight Jacobians

$$D_z\pi_j\big|_{\widehat{v}_j} = -\frac{\lambda_{\mathrm{tl}}}{n_{\mathrm{eff},j}}[\mathrm{diag}(u_j') - \pi_j(u_j')^{\mathsf{T}}], \qquad \partial_{\widehat{v}_j}\pi_j = \frac{\lambda_{\mathrm{tl}}}{n_{\mathrm{eff},j}}[u_j' - \pi_j\mathbf{1}^{\mathsf{T}}u_j']. \tag{68}$$

The chain rule and the triangle inequality then give Equation (61). Since $0 \le u_{jm}' \le 1/4$, $\|u_j'\|_2 \le \sqrt{B}/4$, $\mathbf{1}^{\mathsf{T}}u_j' \le B/4$, and $\|\pi_j\|_2 \le n_{\mathrm{eff},j}^{-1/2}$, the triangle inequality gives Equation (62). The finite-bisection output lies between the bisection endpoints of Appendix B.2.1, which gives $|\widehat{v}_j| \le D_j + 12/\lambda_q$. Each component of r_j is bounded by $M_{\mathrm{ref},j} + 3D_j\widehat{p}_j + \overline{c}_j(|\widehat{v}_j| + 3D_j)$, proving Equation (63). Composing with $\mathscr{J}_{t,j}$ and applying the triangle inequality over active projections gives Equation (64). On an active projection, Equation (29) gives $n_{\mathrm{eff},j} \ge \alpha B/2$, so the factor $\sqrt{B}/n_{\mathrm{eff},j}$ that Equation (63) contributes to Equation (61) is at most $2/(\alpha\sqrt{B})$. □

C.3.3 Soft-Quantile Error

The implemented scale $\lambda_{q,j}$ of Equation (25) is computed from the same batch as the soft quantile, so it couples the summands of Equation (5). This auxiliary result therefore fixes the scale before the batch is drawn and isolates the sampling and smoothing errors of the exact, unclamped soft root. Proposition 5 decomposes the error of the implemented root, whose scale is batch-adaptive and whose soft CDF is clamped.

Proposition 4 (Soft-quantile error at a deterministic scale)**.** *Fix a reverse step t and a projection j, and let $\mathcal{G}_t$ be the σ-field generated by the projected reference sample and by $\mathbf{X}_{t+1}$, the input to the reverse transition that produces the guidance batch $\mathbf{X}_t$. The generated values $Z_m = z_{jm}$, $m = 1, \dots, B$, of $\mathbf{X}_t$ are conditionally independent given $\mathcal{G}_t$; let F_z be the average of their conditional laws, and assume that F_z has a Lebesgue density f_z. The outcome of the effective-tail condition* (29) *at step t is not part of $\mathcal{G}_t$. Let $B \to \infty$, allowing $\mathbf{X}_{t+1}$, and with it F_z, to change with B, and assume that, on a neighbourhood of $Q_\alpha(F_z)$, f_z is bounded above and away from zero and has bounded first three derivatives, with the neighbourhood width and all bounds independent of B.*

In place of the batch statistic $\lambda_{q,j}$, let the scale $\lambda > 0$ be $\mathcal{G}_t$-measurable. Define the population soft quantile $q_{\alpha,\lambda}$ by

$$\mathbb{E}_{Z\sim F_z}[\sigma(\lambda(q_{\alpha,\lambda} - Z))] = \alpha, \tag{69}$$

and let $\bar{v}_{B,\lambda}$ be the exact root of Equation (5) *with λ in place of $\lambda_{q,j}$. There is a λ_0, depending only on the neighbourhood width and the density bounds, such that, conditionally on $\mathcal{G}_t$ and whenever $\lambda \ge \lambda_0$,*

$$\begin{aligned} \bar{v}_{B,\lambda} - q_{\alpha,\lambda} &= O_p(B^{-1/2}), \\ q_{\alpha,\lambda} - Q_\alpha(F_z) &= -\frac{\pi^2}{6\lambda^2}\,\frac{f_z'(Q_\alpha(F_z))}{f_z(Q_\alpha(F_z))} + O(\lambda^{-4}), \end{aligned} \tag{70}$$

where the O_p term is uniform in $\lambda \ge \lambda_0$ and the $O(\lambda^{-4})$ term is uniform in B. The smoothing bias therefore vanishes as $\lambda \to \infty$, but need not vanish as $B \to \infty$ at a fixed λ.

Proof. Given $\mathcal{G}_t$, the preceding guidance update fixes each transition mean, and the paths receive independent Gaussian noise, so the generated values are conditionally independent, though not identically distributed. Write $K_\lambda(a) = \sigma(\lambda a)$, $H_{B,\lambda}(v) = B^{-1}\sum_m K_\lambda(v - Z_m)$, and $H_\lambda = \mathbb{E}H_{B,\lambda}$. Put $q = Q_\alpha(F_z)$ and $q_\lambda = q_{\alpha,\lambda}$. Choose $\delta > 0$ such that $[q - 2\delta, q + 2\delta]$ lies in the common smooth neighbourhood. For constants $0 < c_f < C_f < \infty$ independent of B, the assumptions give

$$f_z \ge c_f \quad \text{and} \quad f_z \le C_f \qquad \text{on } [q - 2\delta, q + 2\delta]. \tag{71}$$

Hence $F_z(q+\delta) - \alpha \ge c_f\delta$ and $\alpha - F_z(q-\delta) \ge c_f\delta$. For a standard logistic random variable L, the representation $H_\lambda(v) = \mathbb{E}_L F_z(v - L/\lambda)$, split according to $|L| \le \lambda\delta$, gives $H_\lambda(q \pm \delta) - F_z(q \pm \delta) = O(\lambda^{-1})$ uniformly in B. The endpoint signs therefore place the unique population root q_λ in $(q - \delta, q + \delta)$ for all $\lambda \ge \lambda_0$, where λ_0 depends only on δ, c_f and C_f.

On this fixed neighbourhood, split the logistic expectation according to whether $|L|/\lambda$ remains inside the common smooth neighbourhood. The complementary probability is exponentially small in λ, and the CDF is bounded by one; the same split applied to the logistic-kernel representation of H'_λ controls its outside-neighbourhood contribution. The local event is symmetric in L, so L and L^3 integrate to zero over it. Taylor expansion on that event, to third order for F_z and to first order for f_z, together with this symmetry, $\mathbb{E}L^2 = \pi^2/3$, $\mathbb{E}L^4 = 7\pi^4/15$ and the bounds on f''_z and f'''_z, gives uniformly in B

$$H_\lambda(q) = \alpha + \frac{\pi^2}{6\lambda^2} f'_z(q) + O(\lambda^{-4}), \qquad H'_\lambda(q) = f_z(q) + O(\lambda^{-2}). \tag{72}$$

Moreover, for $v \in [q-\delta, q+\delta]$,

$$H'_\lambda(v) = \mathbb{E}_L f_z(v - L/\lambda) \geq c_f \ \Pr\{|L| \leq \lambda\delta\}, \tag{73}$$

so H'_λ is uniformly bounded below there for large λ. The second derivative needs one further step, because the assumptions bound f'_z only on the common neighbourhood and say nothing about it elsewhere. Write p_λ for the density of L/λ, so that $p_\lambda(u) = \lambda p_L(\lambda u)$ and

$$H''_\lambda(v) = \int_{\mathbb{R}} f_z(v-u)\, p'_\lambda(u)\, du. \tag{74}$$

Because f_z integrates to one, the region $|u| > \delta$ contributes at most $\sup_{|u|>\delta} |p'_\lambda(u)| = O(\lambda^2 e^{-\lambda\delta})$. Integrating the region $|u| \leq \delta$ by parts leaves $\int_{|u|\leq\delta} f'_z(v-u)p_\lambda(u)\, du$, bounded by the uniform neighbourhood bound on f'_z, together with boundary terms $[f_z(v-u)p_\lambda(u)]^{\delta}_{-\delta} = O(\lambda e^{-\lambda\delta})$. Hence $\sup_v |H''_\lambda(v)| = O(1)$ on a smaller fixed neighbourhood, uniformly in B. The mean-value theorem first gives $q_\lambda - q = O(\lambda^{-2})$; Taylor expansion about q, using the two displays above and the bound on H''_λ, then gives

$$q_\lambda - q = -\frac{\pi^2}{6\lambda^2} \frac{f'_z(q)}{f_z(q)} + O(\lambda^{-4}). \tag{75}$$

In particular, $q_\lambda - q = O(\lambda^{-2})$, and $H'_\lambda \geq c_0 > 0$ on a fixed neighbourhood of q_λ for all $\lambda \geq \lambda_0$, after enlarging λ_0 if necessary.

We next localise the empirical root. For fixed $M > 0$, the mean-value theorem gives

$$H_\lambda(q_\lambda + MB^{-1/2}) - \alpha \geq c_0 MB^{-1/2}, \qquad \alpha - H_\lambda(q_\lambda - MB^{-1/2}) \geq c_0 MB^{-1/2} \tag{76}$$

for all sufficiently large B. At either endpoint, the centred empirical average has variance at most $1/(4B)$. Chebyshev's inequality and monotonicity of $H_{B,\lambda}$ therefore give $\Pr\{|\bar{v}_{B,\lambda} - q_\lambda| > MB^{-1/2} \mid \mathcal{G}_t\} \leq 1/(2c_0^2M^2)$ for all such B. Neither c_0 nor this range of B depends on $\lambda \geq \lambda_0$, so the first line of Equation (70) holds uniformly in $\lambda \geq \lambda_0$, and Equation (75) is its second line. □

For a bisection margin $m_q > |\operatorname{logit}(\alpha)|$, the initial endpoints $\ell^{(0)} = \min_m Z_m - m_q/\lambda$ and $h^{(0)} = \max_m Z_m + m_q/\lambda$ bracket the exact unclamped empirical root. If $\widehat{v}^{(K_q)}_{B,\lambda}$ is the midpoint after K_q bisections, then

$$|\widehat{v}^{(K_q)}_{B,\lambda} - \bar{v}_{B,\lambda}| \leq (h^{(0)} - \ell^{(0)})2^{-(K_q+1)} = \Big(\max_m Z_m - \min_m Z_m + \frac{2m_q}{\lambda}\Big)2^{-(K_q+1)}. \tag{77}$$

This bound concerns the auxiliary unclamped root. Proposition 5 incorporates the clamped root and the active-projection condition used by the finite score.

C.3.4 Finite Implementation Gap

Theorem 1 concerns the ideal score. To compare it with the finite surrogate, let $\ell_j(v, G)$ be the empirical average of Equation (7a) over the projected reference sample, with threshold v, tail law G, and masks evaluated at v. Let $F_{\mathrm{ref},j}$ be the law of $\Pi^{\mathrm{guide}}_j(X)$ for a reference path X, so that $y_{j1}, \ldots, y_{jR}$ is a sample from $F_{\mathrm{ref},j}$. Let

$$L_j(v, G) = \mathbb{E}_{Y\sim F_{\mathrm{ref},j}} \mathcal{S}_{\text{TL–CRPS}}(v, G; Y) \tag{78}$$

be the corresponding population value. The projected reference sample is treated as fixed throughout this comparison. All quantities remain in the native units of guidance projection j.

For the threshold calculation, put $c = 30$ and $m_q = 12$, the implemented clipping level and bisection margin, and write $\lambda_j = \lambda_{q,j} = \widetilde{\tau}_q / s_j$. For the realised generated batch define

$$H_{B,j}(v) = \frac{1}{B} \sum_{m=1}^{B} \sigma(\lambda_j (v - z_{jm})),$$
$$H^c_{B,j}(v) = \frac{1}{B} \sum_{m=1}^{B} \sigma(\mathrm{clip}(\lambda_j (v - z_{jm}), -c, c)). \tag{79}$$

Recall that $\bar{v}_j$ is the unique root of $H_{B,j}(v) = \alpha$, and define the lower crossing of the clamped map by

$$\bar{v}^c_j = \inf\{v : H^c_{B,j}(v) \geq \alpha\}. \tag{80}$$

Finally, for $F \in \mathcal{M}^1$, let $q_{\alpha,\lambda_j}(F)$ be the unique v with $\mathbb{E}_{Z\sim F}\sigma(\lambda_j(v - Z)) = \alpha$.

Proposition 5 (Finite implementation-gap decomposition)**.** *Fix a reverse step and a projection j that passes the effective-tail condition, and assume $F_{\mathrm{ref},j} \in \mathcal{M}^1$. Let $F_j \in \mathcal{M}^1$ be a candidate generated law with induced pair $(v^\star_j, F^\alpha_j)$, where $v^\star_j = Q_\alpha(F_j)$, and let $\widehat{\ell}_j = \ell_j(\widehat{v}_j, G_j)$ be the active finite score. Put $I_j = [\widehat{v}_j \wedge v^\star_j, \widehat{v}_j \vee v^\star_j]$. Then*

$$\widehat{\ell}_j - L_j(v^\star_j, F^\alpha_j) = E_{\mathrm{ref},j} + E_{\mathrm{thr},j} + E_{\mathrm{tail},j}, \tag{81}$$

where

$$E_{\mathrm{ref},j} = \ell_j(v^\star_j, F^\alpha_j) - L_j(v^\star_j, F^\alpha_j),$$
$$E_{\mathrm{thr},j} = \ell_j(\widehat{v}_j, G_j) - \ell_j(v^\star_j, G_j),$$
$$E_{\mathrm{tail},j} = \ell_j(v^\star_j, G_j) - \ell_j(v^\star_j, F^\alpha_j).$$

The threshold and tail-law terms satisfy

$$|E_{\mathrm{thr},j}| \leq \delta_{\mathrm{ref},j}(I_j) \sup_{v\in I_j} |2G_j(v) - 3|\, |\widehat{v}_j - v^\star_j| \leq 3|\widehat{v}_j - v^\star_j|,$$
$$|E_{\mathrm{tail},j}| \leq 4W_1(G_j, F^\alpha_j), \tag{82}$$

where $\delta_{\mathrm{ref},j}(I_j) = \sup_{v\in I_j} |\alpha - \widehat{F}_{\mathrm{ref},j}(v)|$.

Assume $m_q > |\operatorname{logit}(\alpha)|$. The implemented K_q-step bisection threshold further obeys

$$|\widehat{v}_j - v^\star_j| \leq \rho_{\mathrm{num},j} + \rho_{\mathrm{clamp},j} + \rho_{\mathrm{emp},j} + \rho_{\mathrm{smooth},j}, \tag{83}$$

with

$$\rho_{\mathrm{num},j} = |\widehat{v}_j - \bar{v}^c_j|,$$
$$\rho_{\mathrm{clamp},j} = |\bar{v}^c_j - \bar{v}_j|,$$
$$\rho_{\mathrm{emp},j} = |\bar{v}_j - q_{\alpha,\lambda_j}(F_j)|,$$
$$\rho_{\mathrm{smooth},j} = |q_{\alpha,\lambda_j}(F_j) - v^\star_j|. \tag{84}$$

If $\widehat{v}^{\mathrm{ea}}_j$ denotes the midpoint produced by the same K_q bisections in exact arithmetic, then

$$\rho_{\mathrm{num},j} \leq \left(\max_m z_{jm} - \min_m z_{jm} + \frac{2m_q}{\lambda_j}\right) 2^{-(K_q+1)} + \rho_{\mathrm{fp},j}, \qquad \rho_{\mathrm{fp},j} = |\widehat{v}_j - \widehat{v}^{\mathrm{ea}}_j|. \tag{85}$$

These relations are finite-sample statements for the realised adaptive scale $\lambda_j = \widetilde{\tau}_q / s_j$.

Proof. Insert $\ell_j(v^\star_j, G_j)$ and $\ell_j(v^\star_j, F^\alpha_j)$ between $\widehat{\ell}_j$ and $L_j(v^\star_j, F^\alpha_j)$ to obtain Equation (81). As v crosses a projected reference value y_{jr}, the two branches of the score meet. Hence $\ell_j(v, G_j)$ is continuous and is absolutely continuous between its finitely many breakpoints. Its derivative away from those breakpoints is

$$[\alpha - \widehat{F}_{\mathrm{ref},j}(v)] A'_{G_j}(v), \qquad A'_{G_j}(v) = 2G_j(v) - 3. \tag{86}$$

Integration over I_j gives the first inequality in Equation (82). Since $\delta_{\mathrm{ref},j}(I_j) \le 1$ and $|2G_j(v)-3| \le 3$, the second inequality follows.

For the tail-law term, take tail laws $G, \widetilde{G} \in \mathcal{M}^1$. The linear parts of A_G and $A_{\widetilde{G}}$ cancel, and

$$\mathrm{CRPS}(G,u) - \mathrm{CRPS}(\widetilde{G},u) = \int_{\mathbb{R}} (G(x)-\widetilde{G}(x))\big(G(x)+\widetilde{G}(x)-2\mathbf{1}\{u \le x\}\big)\,dx. \tag{87}$$

The second factor lies in $[-2,2]$, so

$$\sup_u |A_G(u) - A_{\widetilde{G}}(u)| \le 2\int_{\mathbb{R}} |G-\widetilde{G}|\,dx = 2W_1(G,\widetilde{G}). \tag{88}$$

Applying this bound with $G = G_j$ and $\widetilde{G} = F_j^{\alpha}$ to the two mask groups in $\ell_j(v_j^\star,\cdot)$, each of total absolute weight at most one, gives $|E_{\mathrm{tail},j}| \le 4W_1(G_j, F_j^{\alpha})$.

Since $m_q \le c$, clipping keeps every argument at the initial endpoints at least m_q in absolute value, so the assumption on m_q ensures that these endpoints bracket the lower crossing of $H^c_{B,j}$. Let $[\ell^{(k)}, h^{(k)}]$ be the bracket after k iterations of the bisection that produces $\widehat{v}_j^{\mathrm{ea}}$. The update rule in Appendix B.2.1 moves $h^{(k)}$ to the midpoint when $H^c_{B,j}$ is at least α there and moves $\ell^{(k)}$ otherwise, so it preserves $H^c_{B,j}(\ell^{(k)}) < \alpha \le H^c_{B,j}(h^{(k)})$, which holds at $k=0$ by this bracketing. Since $H^c_{B,j}$ is continuous and non-decreasing, $\bar{v}_j^c \in (\ell^{(k)}, h^{(k)}]$ for every k. The bracket width after K_q iterations is

$$\left(\max_m z_{jm} - \min_m z_{jm} + \frac{2m_q}{\lambda_j}\right) 2^{-K_q}. \tag{89}$$

Its midpoint $\widehat{v}_j^{\mathrm{ea}}$ is therefore within half this width of $\bar{v}_j^c$. The triangle inequality between the stored threshold, the exact-arithmetic midpoint, and the clamped root gives Equation (85). Adding and subtracting $\bar{v}_j^c$, $\bar{v}_j$, and $q_{\alpha,\lambda_j}(F_j)$ proves Equation (83). The two soft-CDF maps also satisfy

$$\sup_v |H^c_{B,j}(v) - H_{B,j}(v)| \le \sigma(-c). \tag{90}$$

A corresponding root bound requires a positive lower bound on the empirical soft-CDF slope between the two roots, so $\rho_{\mathrm{clamp},j}$ is retained explicitly. □

The four terms in Equation (83) separate numerical solution, clamping, empirical approximation, and smoothing. If $F_j = F_z$ and λ_j is $\mathcal{G}_t$-measurable with $\lambda_j \ge \lambda_0$, where $\mathcal{G}_t$ and λ_0 are as in Proposition 4, that proposition gives $\rho_{\mathrm{emp},j} = O_p(B^{-1/2})$. The signed smoothing displacement has the expansion in Equation (70), and $\rho_{\mathrm{smooth},j} = O(\lambda_j^{-2})$. The same-batch scale $\lambda_j = \widetilde{\tau}_q/s_j$ does not satisfy this deterministic-scale condition, so the two terms remain explicit for the implemented threshold. The fixed projected reference sample makes $E_{\mathrm{ref},j}$ an empirical discrepancy with no rate imposed here, and no rate is assumed for $W_1(G_j, F_j^{\alpha})$. Inactive projections contribute zero to the implemented loss and backward field; Proposition 5 concerns the active branch.

Remark 2 (Effective count under a shared sigmoid scale). If the sigmoids in Equations (5) and (6) use the same inverse scale and clamped sigmoid map, then an exact root $\bar{v}_j^c$ gives $n_{\mathrm{eff},j} = \alpha B$. At the finite-bisection output,

$$|n_{\mathrm{eff},j} - \alpha B| \le B\,\mathrm{Lip}_{I_j^c}(H^c_{B,j})\,|\widehat{v}_j - \bar{v}_j^c|, \tag{91}$$

where I_j^c is the interval between $\widehat{v}_j$ and $\bar{v}_j^c$, and $\mathrm{Lip}_{I_j^c}$ denotes the Lipschitz constant on that interval. This follows from the Lipschitz property of the clamped empirical soft CDF. Up to this residual, the effective-tail condition (29) therefore reduces to the data-independent requirement $\alpha B \ge 8$.

C.4 Controlled Sampler Intervention

C.4.1 Terminal Perturbation Bound

Let $\mathcal{A}$ contain the reverse steps satisfying both the low-noise window and stride conditions. For $t > 0$, L_t^{rev} is the propagation constant of the remaining unguided transition from step $t-1$ to the terminal output. For the deterministic terminal step, set $L_0^{\mathrm{rev}} = 1$.

Proposition 6 (Perturbation control by realised guidance updates)**.** *Let $p_{0,i}^{(g)}$ and $p_{0,i}^{(0)}$ denote the guided and unguided terminal laws of generated path i, let $\mathbf{X}_t$ be the guidance batch containing path i, and let $\boldsymbol{\nu}_t$ be the law of $\mathbf{X}_t$ under the guided sampler. Suppose that, for each $t \in \mathcal{A}$, the one-step output laws $\boldsymbol{\nu}_t\mathbf{K}_t^{(g)}$ and $\boldsymbol{\nu}_t\mathbf{K}_t^{(0)}$ lie in $\mathcal{P}_1(\mathcal{X}^B)$, the expected update norm $\mathbb{E}_{\boldsymbol{\nu}_t}\|\mathbf{U}_t(\mathbf{X}_t)\|_B$ is finite, and the unguided transition $\mathbf{R}_t^{(0)} = \mathbf{K}_{t-1}^{(0)} \cdots \mathbf{K}_0^{(0)}$ remaining after step t, with $\mathbf{R}_0^{(0)} = \mathrm{Id}$, satisfies $W_1(\mathsf{P}\mathbf{R}_t^{(0)}, \mathsf{Q}\mathbf{R}_t^{(0)}) \le L_t^{\mathrm{rev}} W_1(\mathsf{P}, \mathsf{Q})$ for all $\mathsf{P}, \mathsf{Q} \in \mathcal{P}_1(\mathcal{X}^B)$, with W_1 taken under $\|\cdot\|_B$. Then*

$$W_1\Big(p_{0,i}^{(g)}, p_{0,i}^{(0)}\Big) \le \sum_{t\in\mathcal{A}} L_t^{\mathrm{rev}}\, \mathbb{E}_{\mathbf{X}_t\sim\boldsymbol{\nu}_t}\left[\|\mathbf{U}_t(\mathbf{X}_t)\|_B\right]. \tag{92}$$

Proof. We prove Equation (92) first for the terminal batch laws, using one-step coupling and hybrid-chain telescoping, and then pass to each individual path marginal by Wasserstein contraction. When $N > B_{\max}$, the consecutive partition of Appendix B.6 is the same at every reverse step, and the guided and unguided transitions act on each batch separately, so it suffices to treat the batch containing path i. Recall the sampler-perturbation notation of Appendix C.1. For laws on $\mathcal{X}^B$, W_1 uses the product norm $\|\cdot\|_B$. At an active step $t \ge 1$, define

$$\begin{aligned}
\mathbf{K}_t^{(0)}(\cdot \mid \mathbf{x}_t) &= \bigotimes_{m=1}^{B} \mathcal{N}\Big(\mu_\theta(x_t^{(m)}, t), \Sigma_\theta(t)\Big), \\
\mathbf{K}_t^{(g)}(\cdot \mid \mathbf{x}_t) &= \bigotimes_{m=1}^{B} \mathcal{N}\Big(\mu_\theta(x_t^{(m)}, t) - U_{t,m}(\mathbf{x}_t), \Sigma_\theta(t)\Big).
\end{aligned} \tag{93}$$

Here $\bigotimes_{m=1}^{B}$ forms the product measure of its B factors on $\mathcal{X}^B$, so that, given $\mathbf{x}_t$, the B paths receive independent Gaussian noise. The update may depend on the complete batch. At $t = 0$, the transition is deterministic:

$$\begin{aligned}
\mathbf{K}_0^{(0)}(\cdot \mid \mathbf{x}_0) &= \bigotimes_{m=1}^{B} \delta_{\mu_\theta(x_0^{(m)},0)}, \\
\mathbf{K}_0^{(g)}(\cdot \mid \mathbf{x}_0) &= \bigotimes_{m=1}^{B} \delta_{\mu_\theta(x_0^{(m)},0) - U_{0,m}(\mathbf{x}_0)}.
\end{aligned} \tag{94}$$

At inactive steps the two kernels coincide.

The Lipschitz hypothesis of the proposition reads

$$W_1\Big(\mathsf{P}\mathbf{R}_t^{(0)}, \mathsf{Q}\mathbf{R}_t^{(0)}\Big) \le L_t^{\mathrm{rev}} W_1(\mathsf{P}, \mathsf{Q}), \qquad L_0^{\mathrm{rev}} = 1. \tag{95}$$

For any batch input law ν such that $\nu\mathbf{K}_t^{(g)}, \nu\mathbf{K}_t^{(0)} \in \mathcal{P}_1(\mathcal{X}^B)$ and $\mathbb{E}_{\mathbf{X}\sim\nu}\|\mathbf{U}_t(\mathbf{X})\|_B < \infty$, a one-step coupling gives

$$W_1(\nu\mathbf{K}_t^{(g)}, \nu\mathbf{K}_t^{(0)}) \le \mathbb{E}_{\mathbf{X}\sim\nu}\|\mathbf{U}_t(\mathbf{X})\|_B. \tag{96}$$

For $t \ge 1$, couple the transitions using the same input batch and Gaussian noise. Their difference in coordinate m is $-U_{t,m}(\mathbf{X})$, so their product-norm distance is $\|\mathbf{U}_t(\mathbf{X})\|_B$. At $t = 0$, couple the two deterministic images of the same input batch. Their distance is the same quantity. Taking expectations under either coupling proves Equation (96).

Let $\mathbf{P}_0^{(g)}$ and $\mathbf{P}_0^{(0)}$ denote the guided and unguided terminal batch laws. Order the active reverse transitions from $T-1$ down to 0. For each $t \in \mathcal{A}$, consider the hybrid chain that uses guided kernels at all higher reverse-time indices and unguided kernels at all lower indices. Every kernel applied before transition t in this chain is guided, so its law immediately before transition t is $\boldsymbol{\nu}_t$. Replacing the active kernels one at a time gives

$$\begin{aligned}
W_1\Big(\mathbf{P}_0^{(g)}, \mathbf{P}_0^{(0)}\Big) &\le \sum_{t\in\mathcal{A}} W_1\Big((\boldsymbol{\nu}_t\mathbf{K}_t^{(g)})\mathbf{R}_t^{(0)}, (\boldsymbol{\nu}_t\mathbf{K}_t^{(0)})\mathbf{R}_t^{(0)}\Big) \\
&\le \sum_{t\in\mathcal{A}} L_t^{\mathrm{rev}} W_1\Big(\boldsymbol{\nu}_t\mathbf{K}_t^{(g)}, \boldsymbol{\nu}_t\mathbf{K}_t^{(0)}\Big) \\
&\le \sum_{t\in\mathcal{A}} L_t^{\mathrm{rev}} \mathbb{E}_{\mathbf{X}_t\sim\boldsymbol{\nu}_t}\|\mathbf{U}_t(\mathbf{X}_t)\|_B.
\end{aligned} \tag{97}$$

The first inequality follows from hybrid-chain telescoping and the triangle inequality, the second follows from Equation (95), and the third applies Equation (96). Finally, if path i occupies position m of its batch, the coordinate map $\mathbf{x} \mapsto x^{(m)}$ is 1-Lipschitz under the product Euclidean norm. Wasserstein contraction under this map gives Equation (92). □

D Experimental Protocol and Reproducibility

This appendix provides the complete protocol supporting Section 6. Its order follows the evaluation pipeline: trajectory and projection splits, differentiable guidance surrogates, baselines, metrics and backtests, and computational configuration.

D.1 Datasets and Splits

D.1.1 Synthetic Data Generation Parameters

The final training and test sets follow the TailGAN protocol [Cont et al., 2025], including its code for processing, generation, splitting, and evaluation. For hyperparameter selection, we generate a further independent validation sample of 10,000 trajectories from the same synthetic data-generating process. This additional sample is used only for hyperparameter selection, while all reported test results use the original test set.

To ensure controlled evaluation of extreme tail generation, we construct a 5-dimensional synthetic process combining AR(1) dynamics with heavy-tailed GARCH(1,1) volatility. At each step $t \in \{1, \dots, 100\}$, we sample latent innovations $\mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, \Sigma)$, where Σ is constructed using randomised uniform correlations and annualised standard deviations $s_i \sim \mathrm{Unif}(0.3, 0.5)$ for assets $i = 1, \dots, 5$. We write $\sigma_{i,t}$ for the conditional volatility and $\eta_{i,t}$ for a heavy-tailed innovation.

The five asset price increments $\Delta p_{i,t}$ are generated as follows:

$$\begin{aligned}
\Delta p_{1,t} &= u_{1,t}, \\
\Delta p_{2,t} &= 0.5\Delta p_{2,t-1} + u_{2,t}, \\
\Delta p_{3,t} &= -0.12\Delta p_{3,t-1} + u_{3,t}, \\
\Delta p_{4,t} &= \sigma_{4,t}\eta_{1,t}, \\
\Delta p_{5,t} &= \sigma_{5,t}\eta_{2,t}.
\end{aligned}$$

Assets 4 and 5 use GARCH(1,1) conditional variances $\sigma^2_{i,t} = \omega_i + \kappa_i \varepsilon^2_{i,t-1} + \beta_i \sigma^2_{i,t-1}$, where $(\omega_i, \kappa_i, \beta_i)$ are the GARCH parameters and $\varepsilon_{i,t}$ is the corresponding innovation. To induce Student-t tails, let $v_{1,t} \sim \chi^2_5$ and $v_{2,t} \sim \chi^2_{10}$ be independent of $\mathbf{u}_t$ and set $\eta_{1,t} = u_{4,t}/\sqrt{v_{1,t}/5}$ and $\eta_{2,t} = u_{5,t}/\sqrt{v_{2,t}/10}$.

D.1.2 Market Data

Table 5: Empirical Value-at-Risk (VaR) and Expected Shortfall (ES) values on the market training data (Nasdaq ITCH) for the $\alpha = 5\%$ left tail.

	Static Buy-and-Hold		Mean Reversion		Trend Following	
Asset	VaR	ES	VaR	ES	VaR	ES
AAPL	-0.366	-0.591	-0.307	-0.491	-0.382	-0.527
AMZN	-0.420	-0.606	-0.347	-0.508	-0.330	-0.453
GOOG	-0.377	-0.572	-0.340	-0.510	-0.347	-0.462
JPM	-0.360	-0.519	-0.287	-0.405	-0.350	-0.473
QQQ	-0.223	-0.447	-0.192	-0.349	-0.191	-0.289

The finance benchmark uses Nasdaq ITCH limit-order-book data for five assets (AAPL, AMZN, GOOG, JPM, QQQ), sampled at 9-second intervals into trajectories of length $H = 100$. Following TailGAN [Cont et al., 2025] and DSI [Zhao et al., 2026], the final model uses November 2019 for training and the first week of December 2019 for testing. For hyperparameter selection, we partition

November chronologically by complete trading days, using the earlier observations for fitting and the final week for validation. After selecting $(\gamma, \tau_{\text{start}})$, we retrain the backbone on the complete November training set and evaluate the frozen configuration on the original test set. See Table 5 for dataset statistics.

D.2 Evaluation Projections and Strategy Definitions

For $s \in \{\text{train}, \text{val}, \text{test}\}$, let $\mathscr{P}_s = \{\Pi_j^s\}_{j=1}^{J_s}$ denote the exact strategy projections assigned to split s. The sets contain distinct instances produced by the same construction procedure. Projection-dependent fitting and guidance use $\mathscr{P}_{\text{train}}$, tuning uses $\mathscr{P}_{\text{val}}$, and common final evaluation uses $\mathscr{P}_{\text{test}}$ with $J_{\text{test}} = 65$. Each Π_j^s maps a multivariate price trajectory $\mathbf{S}_{1:H}$ to a scalar PnL value X. Here, $d = 1, \ldots, D$ indexes assets, $h = 1, \ldots, H$ indexes data time, and H differs from the number T of diffusion steps. The families are shared across splits, so we suppress s below.

1. Buy-and-Hold (Multivariate Dependence): Evaluates static allocations. For a weight vector $\mathbf{w}$, the portfolio value is $V_h = \sum_{d=1}^{D} w_d S_{d,h}$. With capital multiplier C, the terminal PnL is $X = C(V_H - V_0)$.

2. Mean Reversion (Short-Term Temporal): Trades are based on deviations from the mean of the first $L_{\text{init}} + 1$ price levels. For the normalised signal $\xi_h = (S_h - \mu_{\text{init}})/0.01$, an upward crossing of the upper threshold opens a short position and a downward crossing of the lower threshold opens a long position. The position is carried forward until the next zero crossing closes it; the terminal position is set to zero before PnL aggregation, and PnL sums the position at h times $S_{h+1} - S_h$.

3. Trend Following (Long-Term Temporal): Uses simple moving average (SMA) crossovers. We define $\xi_h = (\text{SMA}_{\text{short},h} - \text{SMA}_{\text{long},h})/0.01$. The logic is inverted compared with mean reversion to follow momentum: long if $\xi_h > \tau_{\text{upper}}$ and short if $\xi_h < \tau_{\text{lower}}$.

While the discrete strategies defined in this subsection are required for rigorous financial evaluation, their reliance on Heaviside step functions and hard sorting operations blocks the gradient flow required for the guided mean update in Equation (9). To enable inference-time guidance, we construct continuous, differentiable surrogates for Mean Reversion and Trend Following using temperature-scaled sigmoids.

D.3 Differentiable Guidance Surrogates

Each selected Π_j^{train} has a differentiable counterpart Π_j^{guide}, which enters the guidance gradient g_t of Equation (8) through the generated and projected reference values in Equation (4). Validation and test projections are excluded.

Let $x \in \mathbb{R}^{D \times H}$ denote a generated return path for D assets over H data-time steps. We recover the unnormalised price path $P \in \mathbb{R}^{D \times (H+1)}$ by assuming an initial price of 1 and accumulating the returns:

$$P_{d,h} = 1 + \sum_{s=1}^{h} x_{d,s} \tag{98}$$

for $d = 1, \ldots, D$ and $h = 1, \ldots, H$. The price increment from h to $h+1$ is $\Delta P_{d,h} = P_{d,h+1} - P_{d,h} = x_{d,h+1}$.

Linear Projections: Buy-and-Hold and Static Portfolios. The simplest projections are linear combinations of the start and end prices. Let C denote the base capital allocation scaling factor.

The individual asset Buy-and-Hold (Stk) guidance PnL is given by

$$\Pi_{Stk,d}^{\text{guide}}(x) = C(P_{d,H} - P_{d,0}). \tag{99}$$

For the selected training-side static portfolios, let $\mathbf{W}^{\text{train}} \in \mathbb{R}^{D \times K}$ denote their weight matrix, where K is the number of portfolios (e.g., $K = 50$). The guidance PnL for the k-th portfolio is the weighted sum of individual asset returns:

$$\Pi_{Trans,k}^{\text{guide}}(x) = C \sum_{d=1}^{D} \mathbf{W}_{d,k}^{\text{train}}(P_{d,H} - P_{d,0}). \tag{100}$$

Both of these operations are trivially differentiable with respect to x.

Differentiable Mean Reversion (MR) Surrogate. The standard Mean Reversion strategy takes a short position when the asset price far exceeds its historical mean and a long position when it falls significantly below it. Because standard implementations rely on hard Boolean thresholds (e.g., $\mathbb{I}_{\{z>ST\}}$), they block gradient flow. We formulate a differentiable soft surrogate using sigmoid functions.

First, we compute a static moving average MA_d over a fixed initial window of length L_{MA}:

$$MA_d = \frac{1}{L_{\mathrm{MA}}+1} \sum_{s=0}^{L_{\mathrm{MA}}} P_{d,s}. \tag{101}$$

We define a normalised deviation (pseudo z-score) for the price at data time h:

$$\xi_{d,h} = \frac{P_{d,h} - MA_d}{0.01}. \tag{102}$$

Let ST^{MR} and LT^{MR} denote the pooled short and long entry thresholds (the 69th and 31st percentiles, respectively), broadcast to every asset. We approximate the discrete trading positions using the logistic CDF σ with projection steepness κ_Π (we use $\kappa_\Pi = 10$), together with long- and short-position multipliers LR and SR:

$$Pos^{MR}_{d,h} = \underbrace{C\,LR\,\sigma\big(\kappa_\Pi(LT^{MR} - \xi_{d,h})\big)}_{\text{soft long position}} - \underbrace{C\,SR\,\sigma\big(\kappa_\Pi(\xi_{d,h} - ST^{MR})\big)}_{\text{soft short position}}. \tag{103}$$

The differentiable MR PnL is the sum of the products of these soft positions and the subsequent price increments:

$$\Pi^{\mathrm{guide}}_{MR,d}(x) = \sum_{h=0}^{H-1} Pos^{MR}_{d,h} \Delta P_{d,h}. \tag{104}$$

Differentiable Trend Following (TF) Surrogate. The Trend Following strategy relies on the crossover of fast and slow moving averages. We use 1D average pooling to compute the rolling fast and slow moving averages at each data-time step $h \geq L_{\mathrm{MA}}$:

$$MA_{fast,d,h} = \frac{1}{L_{\mathrm{MA}}} \sum_{s=h-L_{\mathrm{MA}}+1}^{h} P_{d,s}, \quad MA_{slow,d,h} = \frac{1}{2L_{\mathrm{MA}}} \sum_{s=h-2L_{\mathrm{MA}}+1}^{h} P_{d,s}. \tag{105}$$

The corresponding trend signal is $\xi_{d,h} = (MA_{fast,d,h} - MA_{slow,d,h})/0.01$.

Unlike Mean Reversion, Trend Following takes a long position during strong uptrends ($\xi_{d,h} \gg 0$) and a short position during strong downtrends ($\xi_{d,h} \ll 0$). Using the pooled empirical thresholds LT^{TF} (uptrend) and ST^{TF} (downtrend), broadcast to every asset, we define the continuous TF position as

$$Pos^{TF}_{d,h} = \underbrace{C\,LR\,\sigma\big(\kappa_\Pi(\xi_{d,h} - LT^{TF})\big)}_{\text{soft long position}} - \underbrace{C\,SR\,\sigma\big(\kappa_\Pi(ST^{TF} - \xi_{d,h})\big)}_{\text{soft short position}}. \tag{106}$$

The differentiable TF PnL is similarly aggregated over time:

$$\Pi^{\mathrm{guide}}_{TF,d}(x) = \sum_{h=0}^{H-1} Pos^{TF}_{d,h} \Delta P_{d,h}. \tag{107}$$

By replacing step functions with temperature-scaled sigmoids, $\Pi^{\mathrm{guide}}_{MR}(x)$ and $\Pi^{\mathrm{guide}}_{TF}(x)$ provide smooth, continuous gradients $\nabla_x \Pi^{\mathrm{guide}}_j(x)$ for the reverse diffusion trajectory.

D.4 Baselines and Comparison Protocol

We enforce a like-for-like comparison protocol across methods. TAILDIFF, DDPM, and DSI all draw on one DDPM training run and differ only in how they sample from it: DDPM uses unguided ancestral sampling from a single checkpoint, TAILDIFF adds TL–CRPS tail-law guidance to that same sampler at that same checkpoint, and DSI pools samples across several checkpoints saved along the same run. No method receives additional training, additional training data, or privileged labels.

Following the TailGAN [Cont et al., 2025] and DSI [Zhao et al., 2026] setups, all methods share the trajectory splits, each split-specific projection set, and the tail-risk metrics (VaR/ES RE at matched α). Final evaluation uses $\mathscr{P}_{\text{test}}$. The primary comparison uses the same finite generation budget ($N = 1000$). Reported means and standard deviations are computed over matched seed sets. Following TailGAN [Cont et al., 2025], SE(1000) is the matched-budget oracle sampling error obtained from the known data-generating law on Synthetic and the held-out empirical distribution on Market. HSM is included as a transparent nonparametric reference that resamples observed trajectories. TailGAN and WGAN are evaluated using their released protocols under the same downstream evaluation pipeline.

For the synthetic experiments, we use the official TailGAN code [Cont et al., 2025] for data generation, evaluation, and the released GAN baselines. The DDPM and DSI baselines use the official DSI implementation [Zhao et al., 2026]. We harmonise data splits, evaluation metrics, and generation budgets across these codebases so that differences are attributable to method behaviour rather than protocol drift.

This setup is intended to isolate the effect of inference-time tail guidance from changes in architecture, training objective, or data access.

D.5 Metrics and Backtest Calibration

Both backtests follow Cont et al. [2025, Sec. 4.2.3] and use $\mathscr{P}_{\text{test}}$ at the 5% level over 100 trials, each drawing $N_{\text{gen}} = 1000$ generated paths to estimate the tail functionals and $N_{\text{real}} = 1000$ held-out observations from the true law. The Kupiec coverage test [Kupiec, 1995] counts exceedances $C_j = \sum_{i=1}^{N_{\text{real}}} \mathbf{1}\{\Pi_j^{\text{test}}(p_i) < \widehat{\text{VaR}}_{\alpha,j}\}$, where p_i is held-out path i and $\widehat{\text{VaR}}_{\alpha,j}$ is the empirical VaR of test projection $j = 1, \ldots, J_{\text{test}}$ over the N_{gen} generated paths. The test refers the likelihood-ratio statistic to χ_1^2, while the joint VaR–ES score test [Fissler et al., 2016] compares the realised Fissler–Ziegel scores of the generated VaR–ES pair with those of the ground-truth estimates of VaR and ES.

The coverage test is oversized under this protocol. Its χ_1^2 calibration presumes a known threshold, whereas a single estimate $\widehat{\text{VaR}}_{\alpha,j}$ is shared across all N_{real} observations, so the exceedance indicators are positively correlated. Let F be the true CDF of the projected outcome and write $U = F(\widehat{\text{VaR}}_{\alpha,j})$, so that $C_j \mid U \sim \text{Bin}(N_{\text{real}}, U)$. When the generated paths are independent draws from the true law, $\text{Var}(U) \approx \alpha(1-\alpha)/N_{\text{gen}}$, and the law of total variance gives

$$\text{Var}(C_j) \;\approx\; N_{\text{real}}\alpha(1-\alpha)\left(1 + \frac{N_{\text{real}}}{N_{\text{gen}}}\right), \tag{108}$$

which doubles at $N_{\text{gen}} = N_{\text{real}}$. For such draws, U is a uniform order statistic for every continuous F, so the resulting size depends only on $(\alpha, N_{\text{gen}}, N_{\text{real}})$ and not on the data-generating process. Table 6 reports this size.

Table 6: Finite-sample size (%) of the Kupiec coverage test at the nominal 5% level for independent draws from the true law, with $\alpha = 0.05$ and $N_{\text{real}} = 1000$.

N_{gen}	500	1000	2000	5000	10^4	10^5	known
Size (%)	26.0	16.6	11.1	7.5	6.3	5.3	5.1

At the operating point $N_{\text{gen}} = N_{\text{real}} = 1000$ the test rejects independent draws from the true law 16.6% of the time, and direct simulation under Gaussian, $t(4)$ and $t(3)$ marginals returns 16.5 to 17.4%, confirming the law-independence implied by Equation (108). This is a reference rate, not

a lower bound: a generator whose VaR estimate concentrates more tightly around the true VaR is rejected less often, down to the 5.1% of a known threshold.

D.6 Base Diffusion Architecture

All diffusion methods, including our TAILDIFF framework and the DDPM baselines, share an identical base architecture. We implement a continuous-state DDPM over $T = 1000$ diffusion steps with a linear variance schedule from $\beta_1 = 10^{-4}$ to $\beta_T = 0.02$. The reverse process $\epsilon_\theta(\mathbf{x}_t, t)$ is parameterised by a 1D U-Net with a base channel dimension of 64. The network employs sinusoidal time embeddings projected via an MLP, 1D residual blocks with Group Normalisation (groups=8) and SiLU activations, and symmetric downsampling/upsampling paths with skip connections to preserve high-frequency temporal details. Models are optimised using AdamW with a learning rate of 10^{-5}.

D.7 Compute Resources and Hyperparameters

Table 7: **Inference-time cost of tail guidance.** Wall time and peak GPU memory for generating $N = 1000$ trajectories on a single NVIDIA RTX 3090.

	Synthetic		Market	
	Guided	Unguided	Guided	Unguided
Wall time	1m 15s	24s	50s	24s
Peak allocated	2.01 GB	1.88 GB	1.88 GB	1.88 GB
Peak reserved	2.50 GB	2.10 GB	2.16 GB	2.10 GB

All models were trained and evaluated on a single NVIDIA RTX 3090 GPU. The $J_{\text{guide}}R$ projected reference values are computed once per run at cost $O(J_{\text{guide}}R)$, after which each active guidance step has the cost given in Equation (31) and depends on B, R, and the fixed bisection count rather than on the full training data size. Table 7 reports wall time and peak GPU memory for generating $N = 1000$ trajectories, against unguided ancestral sampling from the same backbone, where the gap between the two regimes is driven by τ_{start} rather than by R. TAILDIFF uses $\gamma = 10$ and $\tau_{\text{start}} = 0.3$ on Synthetic data, and $\gamma = 24$ and $\tau_{\text{start}} = 0.3$ on Market data.

E Additional Accuracy and Robustness Results

E.1 Generation-Budget Sensitivity and Guidance-Batch Stability

Table 8 reports finite-budget behavior as the generated path count N increases. TAILDIFF uses active guidance batches of size $B = \min(N, 1000)$. For $N \leq 1000$, increasing N also increases the batch used to estimate the empirical tail thresholds during guidance. For $N > 1000$, generation is split into batches of size at most 1000, so additional samples primarily reduce Monte Carlo error. The improvement from $N = 100$ to $N = 1000$ shows under-populated guidance batches give noisier empirical thresholds, whereas larger batches stabilize the guidance field. The continued variance reduction beyond $N = 1000$ reflects standard finite-sample averaging.

Table 8: Generation-budget sensitivity and guidance-batch stability: mean relative error (%) $\pm$ standard deviation vs. generated paths (N). For $N \leq 1000$, increasing N also increases the active guidance batch $B = \min(N, 1000)$.

Dataset	$N = 100$	$N = 500$	$N = 1000$	$N = 2000$	$N = 5000$
Synthetic	9.5 ± 1.7	4.8 ± 0.3	3.4 ± 1.1	3.3 ± 0.8	3.3 ± 0.3
Market	20.7 ± 8.7	12.0 ± 1.8	8.2 ± 1.4	8.0 ± 0.9	7.9 ± 0.4

E.2 Robustness across Regulatory Risk Levels

This subsection reports the extended risk-level sweep for TAILDIFF across the synthetic and market regimes. The pattern is consistent with Section 6.4: performance is stable across regulatory levels, with expected degradation only at the most extreme tail due to reduced effective sample mass.

Table 9: TAILDIFF robustness across regulatory risk levels (α), reported as mean RE (%) $\pm$ standard deviation.

	Synthetic (RE)			Market (RE)		
Method	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$
TAILDIFF (Ours)	4.1 ± 1.3	3.4 ± 1.1	3.1 ± 0.7	9.2 ± 2.1	8.2 ± 1.4	7.9 ± 0.9

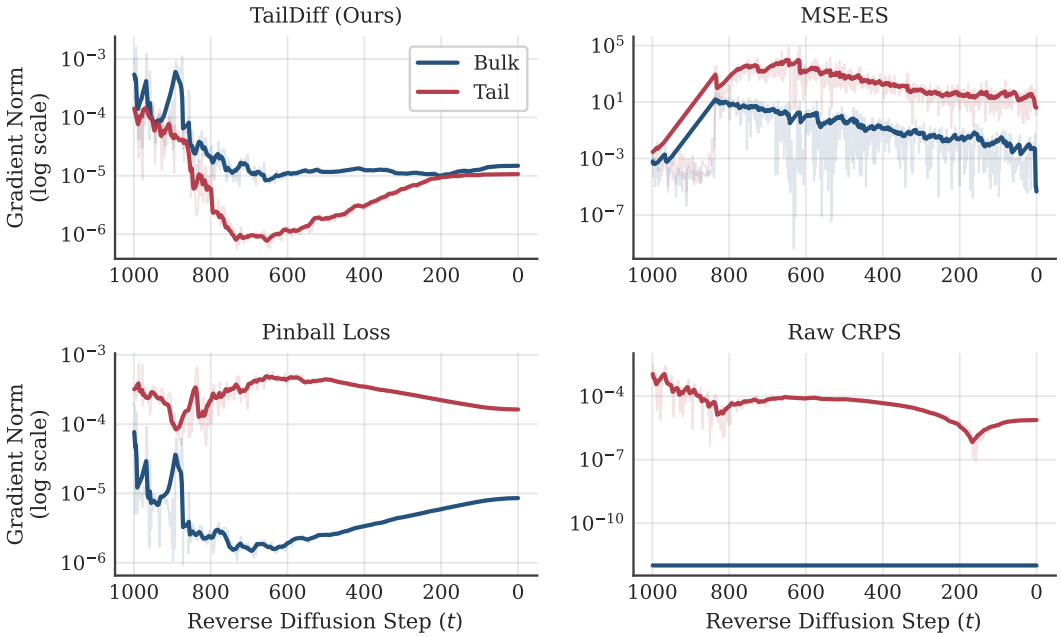


Figure 5: **Guidance gradient magnitudes.** Expected gradient norms (log scale) for bulk (blue) and tail (red) samples across the reverse process. TAILDIFF (top left) shows stable, noise-adaptive behaviour: tail gradients remain active and the reference-threshold term evaluates the boundary in the low-noise regime. Baselines show larger instability (MSE-ES), late-stage detachment (raw CRPS), or incomplete tail-law shaping (Pinball Loss). These are diagnostics, not theorem verification.

F Guidance Mechanism, Stability, and Structural Fidelity

F.1 Contribution of the Tail Objective

We compare TAILDIFF against three alternative guidance objectives using the same pretrained diffusion backbone and reverse-update rule. Each objective is tuned independently under the same validation budget, with a grid search over $\gamma \in [10, 40]$ and $\tau_{\text{start}} \in \{0.1, 0.3, 0.5\}$. The selected configurations are evaluated over five seeds using the common test projections and evaluation pipeline.

The alternatives include MSE-ES, which targets ES as a point-valued summary, and Pinball Loss, which targets the quantile boundary without explicitly shaping the tail distribution. We also consider raw CRPS (full-distribution scoring), which evaluates the complete projected distribution against reference outcomes rather than explicitly targeting the generated quantile boundary and its associated left-tail law. Table 3 reports the resulting VaR/ES relative errors and gradient-direction diagnostics.

F.2 Extended Gradient Dynamics

Figure 5 separates the expected gradient norms of bulk and tail samples across the four guidance objectives, complementing the aggregate diagnostics in Section 6.3.

The baselines exhibit several pathologies that render them less suitable for stable diffusion guidance:

- **MSE-ES:** Exhibits severe numerical instability. Gradient magnitudes spike orders of magnitude above standard diffusion scores, consistent with the trajectory collapse observed in experiments.
- **Raw CRPS:** Scores the full projected distribution without explicitly targeting the left-tail law. Its gradients exhibit reduced magnitude and directional coherence in the low-noise regime, indicating less stable guidance.
- **Pinball Loss:** Exhibits a smoother magnitude curve, but its quantile-only focus does not shape the distribution below the boundary, making it insufficient as a tail-law guidance signal.

F.3 Guidance Scale, Low-Noise Window, and Update Magnitude

The guidance scale γ and the active low-noise window τ_{start} jointly determine the strength and duration of the reverse-mean perturbation. Figure 4 in the main paper reports the complete sweep used to select the operating points in the reported experiments.

Table 10 reports the realised guidance-update magnitude $\|U_t\|$ (the batch average $B^{-1}\sum_m \|U_{t,m}\|_F$ of the per-path update norms), tail-risk accuracy, and structural diagnostics as γ varies on market validation trajectories and $\mathscr{P}_{\mathrm{val}}$. The update increases across the range considered. Structural metrics remain stable where tail accuracy is best. These measurements provide empirical context for the realised guidance updates in Proposition 6, rather than theorem verification. By the Cauchy–Schwarz inequality, this per-path average is at most $\|\mathbf{U}_t\|_B/\sqrt{B}$, where $\|\mathbf{U}_t\|_B$ is the batch norm in Equation (92).

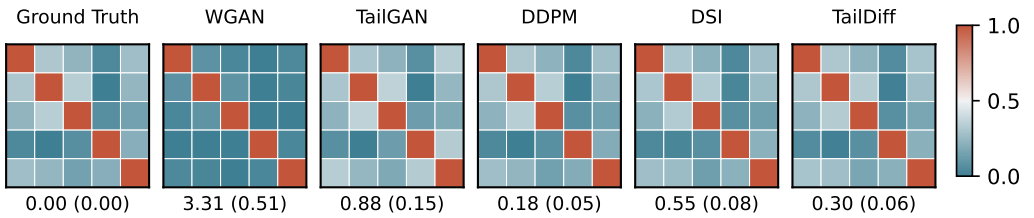
Table 10: **Guidance-scale sensitivity on market validation data.** Realised guidance-update magnitude $\|U_t\|$ (the batch average $B^{-1}\sum_m \|U_{t,m}\|_F$ of the per-path update norms), averaged over active steps, VaR/ES RE (%) on $\mathscr{P}_{\mathrm{val}}$, cross-asset correlation distance (Corr Δ), and lag-1 autocorrelation distance (AC Δ) as functions of γ. We fix $\tau_{\text{start}} = 0.3$. Lower is better. The best RE is highlighted.

γ	$\|U_t\|$ ($\times 10^{-2}$)	RE (%) $\downarrow$	Corr $\Delta\downarrow$	AC $\Delta\downarrow$
10	0.33 ± 0.05	14.6 ± 2.9	0.040 ± 0.010	0.039 ± 0.012
13	0.78 ± 0.06	12.1 ± 2.2	0.041 ± 0.011	0.036 ± 0.011
15	1.24 ± 0.11	10.3 ± 1.6	0.045 ± 0.012	0.037 ± 0.011
19	1.72 ± 0.13	8.4 ± 1.5	0.044 ± 0.011	0.041 ± 0.014
24	1.95 ± 0.11	$\mathbf{8.0 \pm 1.6}$	0.046 ± 0.013	0.042 ± 0.013
28	2.28 ± 0.21	8.3 ± 1.8	0.043 ± 0.013	0.044 ± 0.015
32	2.81 ± 0.26	10.8 ± 2.4	0.054 ± 0.012	0.052 ± 0.013
36	3.64 ± 0.32	12.9 ± 3.1	0.061 ± 0.014	0.067 ± 0.019

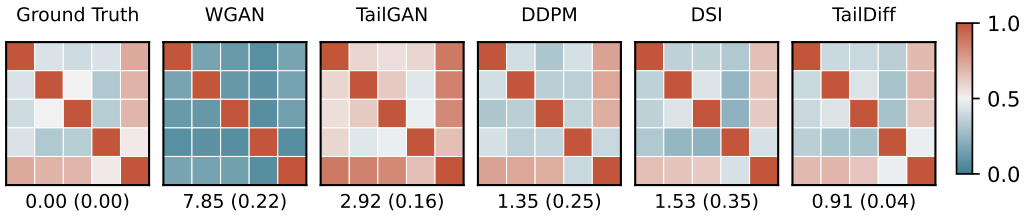
F.4 Structural and Distributional Fidelity

To supplement the sampler-perturbation analysis in Section 5, we report dependency diagnostics. Figures 6a and 6b display the multivariate cross-asset correlation heatmaps for the Synthetic and Market regimes, respectively. Figures 6c and 6d detail temporal autocorrelation decay. Across all metrics, TAILDIFF preserves the macroscopic dependence structure of the underlying data without the mode-collapse behaviour observed in adversarial baselines.

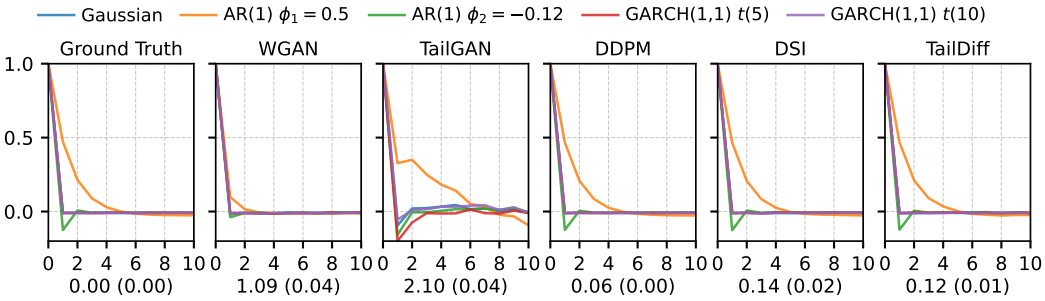
Rank-frequency distributions. Figures 7 and 8 visualise empirical rank-frequency distributions of generated PnL across all test strategy projections. These log-scale plots show where unguided DDPM sampling under-captures extreme mass and how TAILDIFF reshapes the tail towards the ground-truth frequency of rare events.



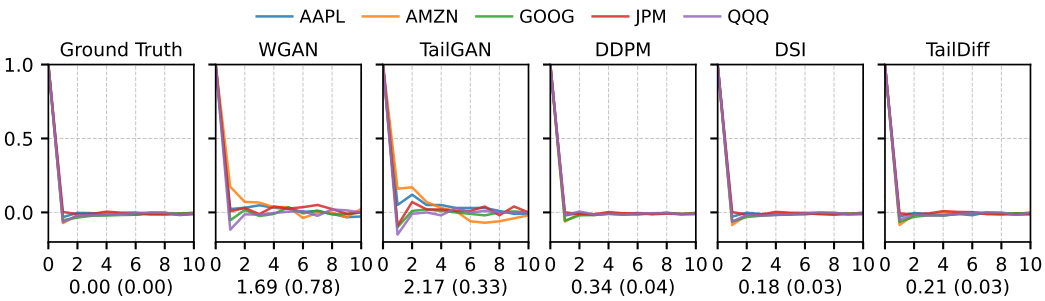
(a) Cross-asset correlations, synthetic regime.



(b) Cross-asset correlations, market regime.



(c) Temporal autocorrelations, synthetic regime.



(d) Temporal autocorrelations, market regime.

Figure 6: **Dependence structure across regimes.** Panels (a) and (b) report cross-asset correlations, while panels (c) and (d) report temporal autocorrelation decay. Numbers give the mean and standard deviation (in parentheses) across model runs of the summed absolute differences between the reference and generated coefficient arrays. Each generated estimate uses 1000 paths.

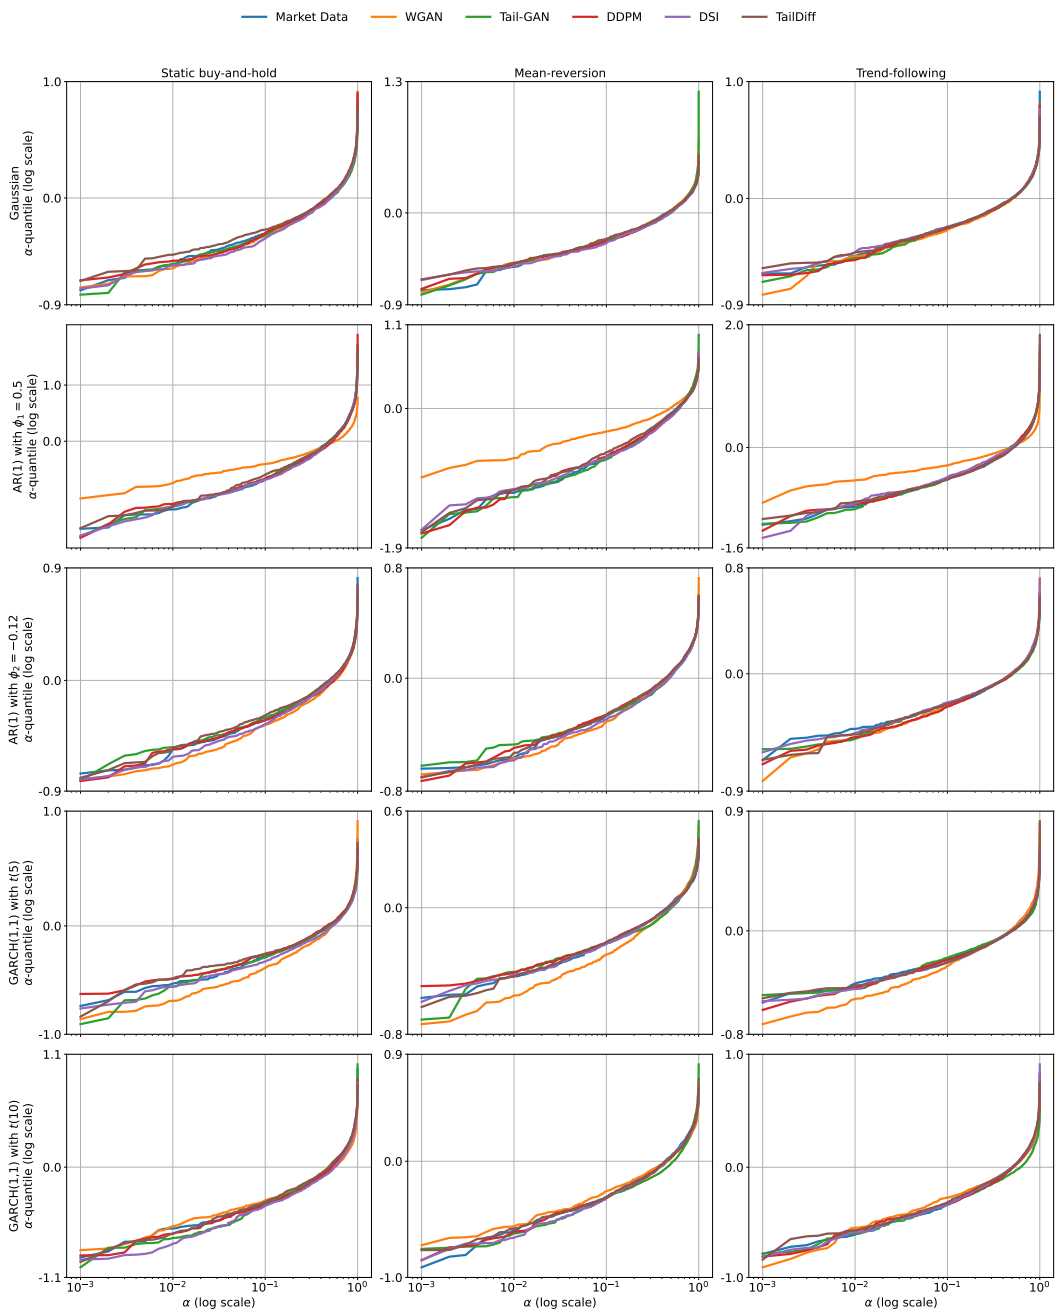


Figure 7: Tail behaviour via the synthetic rank-frequency distribution of the test-strategy PnL. The columns represent the strategy types.

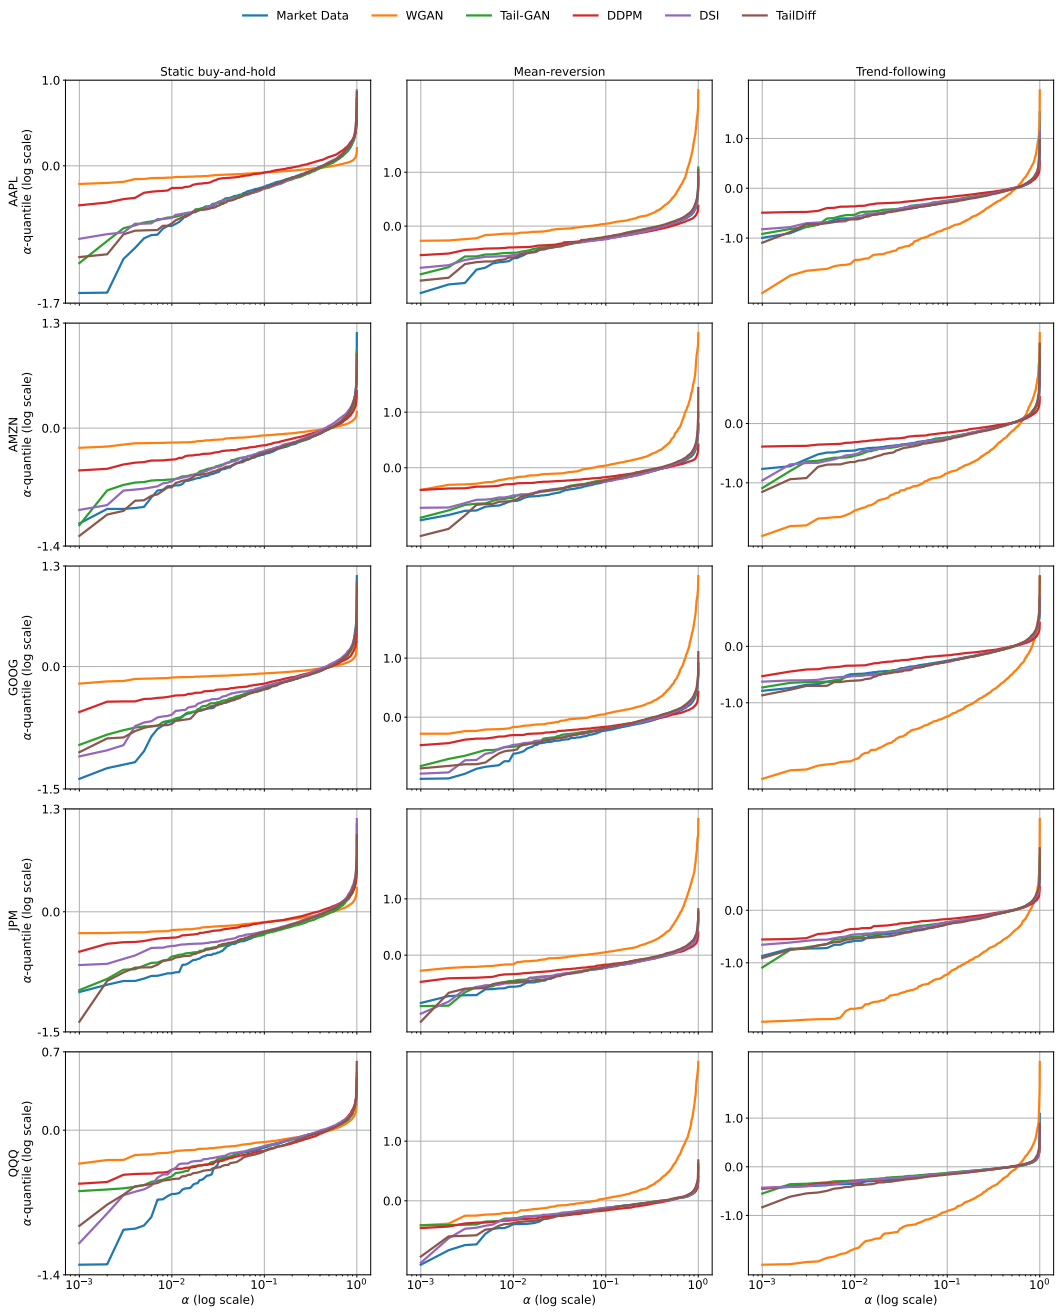


Figure 8: Tail behaviour via the market rank-frequency distribution of the test-strategy PnL. The columns represent the strategy types.

G Cross-Domain Portability

This appendix reports cross-domain portability results supporting Section 6.4. We apply TAILDIFF guidance at inference time to diffusion backbones trained on four public time-series datasets spanning energy, traffic, healthcare, and weather settings that differ structurally from finance but depend equally on accurate extremes. This experiment assesses whether the guidance mechanism transfers without domain-specific retraining across these domains.

G.1 Experimental Protocol

For each benchmark, we follow the default training configuration of the DDPM backbone in Appendix D.6. We use the same experimental setup as the main experiments and generate $N = 1000$ synthetic trajectories. For a unified criterion, we define a scalar stress functional $g(x)$ on each trajectory and compute tail summaries from $\{g(x_i)\}_{i=1}^{N}$. We use Equation (1) for adverse low outcomes (e.g. PnL losses and abnormal electrocardiogram excursions) and the right-tail convention in Equation (16) for exceedance risks (e.g. peak electricity demand and traffic congestion spikes). Relative error is reported against the same statistic from held-out real data.

G.2 Datasets and Stress Functionals

Energy: PJM Electricity Demand [Kim, 2025]. The PJM dataset contains hourly electricity demand across the PJM Interconnection grid. We use stress functionals targeting peak-demand events: maximum horizon demand, cumulative overload above a threshold, and worst-period ramp increase.

Traffic: NYC Taxi Demand [New York City Taxi and Limousine Commission, 2024]. The NYC Taxi and Limousine Commission dataset contains trip records aggregated into demand time series. Stress functionals capture congestion extremes: maximum congestion level, cumulative severe congestion, and largest short-horizon spike.

Healthcare: ECG5000 [Dau et al., 2018]. ECG5000 contains electrocardiogram recordings from the UCR Time Series Archive. Stress functionals target abnormal excursions: extreme deviation scores and maximum abnormal excursion over the horizon.

Weather: Monash Weather Dataset [Godahewa et al., 2021]. The Monash Weather Dataset contains multivariate meteorological time series. Stress functionals capture temperature extremes: peak temperature exceedance and maximum short-horizon anomaly.

H Changes from the Reviewed Version

Following the review discussion, we further investigated the population interpretation and numerical stability of the guidance objective, which led to a revised finite formulation. The reviewed version used a fixed reference-derived threshold together with the generated empirical distribution; the camera-ready version instead constructs the quantile boundary and tail law jointly from the generated samples, with the sampling update revised accordingly. This formulation retains the same inference-time guidance framework and the same tail-scoring motivation of Fissler et al. [2026], but enables the stronger population-identification result in Theorem 1 and the accompanying finite-implementation analysis. The research problem, pretrained diffusion backbone, experimental setting, evaluation protocol, and primary empirical conclusions remain unchanged. These revisions were developed after the review discussion in response to issues raised during that process.

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 made matches theoretical and experimental results, and reflect how much the results can be expected to generalise to other settings.

   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: See 7.

   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: Please see Appendix C.

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: Please see Section 6 and Appendix D.

   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 benchmark construction and evaluation protocol follow the publicly available TailGAN codebase: https://github.com/chaozhang-ox/Tail-GAN, which provides the synthetic data generator, nonparametic, WGAN, TailGAN baselines, and evaluation pipeline. All datasets used in our experiments are publicly available. Our implementation of the TAILDIFF, diffusion-based baselines, experiment scripts, and configuration files are available at https://github.com/ShuningZhao/TailDiff/.

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: Please see Section 6 and Appendix D.

   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: All results and plots contain confidence interval and error bars. Additional statistical tests are available in Table 2.

   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: See D.7.

   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: The research uses publicly available data, does not involve human-subject experiments or personal data collection, and the work follows the NeurIPS Code of Ethics.

   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: [Yes]

    Justification: See section 7.

    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: The paper poses no such risks.

    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: We credit and cite the original sources for all existing assets used in the paper, including benchmark datasets, baseline implementations, and evaluation code where applicable. The supplemental material documents the relevant asset sources, versions or access dates when available, licences/terms of use, and any modifications made for our experiments. We use these assets only for research and benchmarking purposes and respect their stated licences and terms of use.

    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: [Yes]

    Justification: The paper introduces a software artefact implementing TailDiff and its evaluation protocol. The implementation and supporting documentation are publicly available at https://github.com/ShuningZhao/TailDiff/, including installation instructions, dependencies, data preprocessing, experiment scripts, configuration files, evaluation metrics, and instructions for reproducing the main results. The artefact does not introduce new human-subject data or identifiable personal information; therefore, consent documentation is not applicable.

    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: The paper does not involve crowdsourcing nor research with human subjects.

    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: The paper does not involve crowdsourcing nor research with human subjects.

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: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

    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.