ICASSP2026 Barcelona

IEEE Signal Processing Society

IMPERIAL

RefineBridge

Generative Bridge Models Improve Financial Forecasting by Foundation Models

Presenter: *Anthony Bolton*

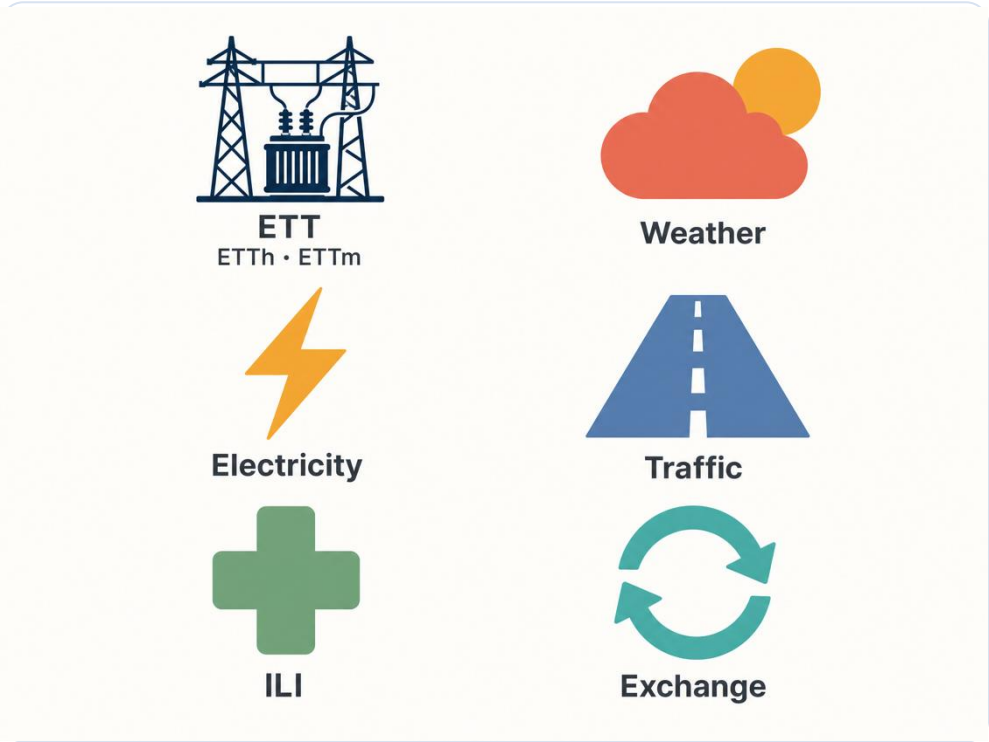
Anthony Bolton * *Wuyang Zhou* *

Zehua Chen *Giorgos Iacovides* *Danilo Mandic*

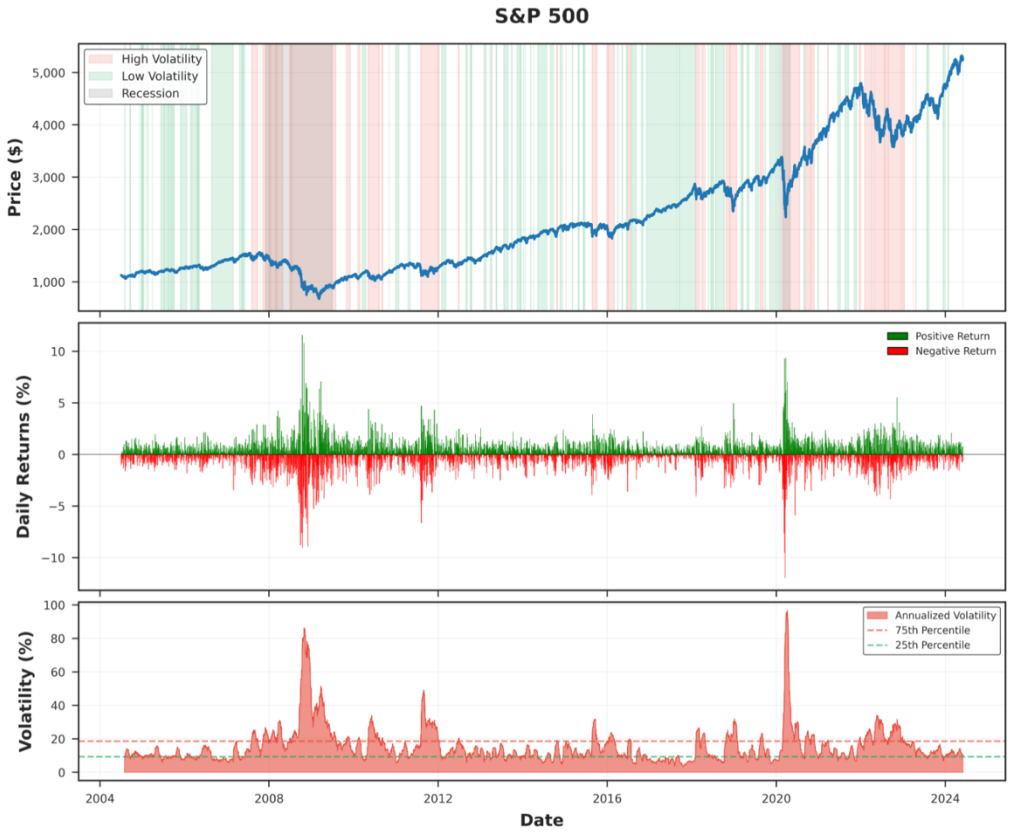
TSFMs on finance: a different problem

IMPERIAL

Standard benchmarks



Financial data: a different beast



Three structural properties standard benchmarks don't test:

- Non-stationary regimes
- Heavy-tailed distributions
- High frequency noise

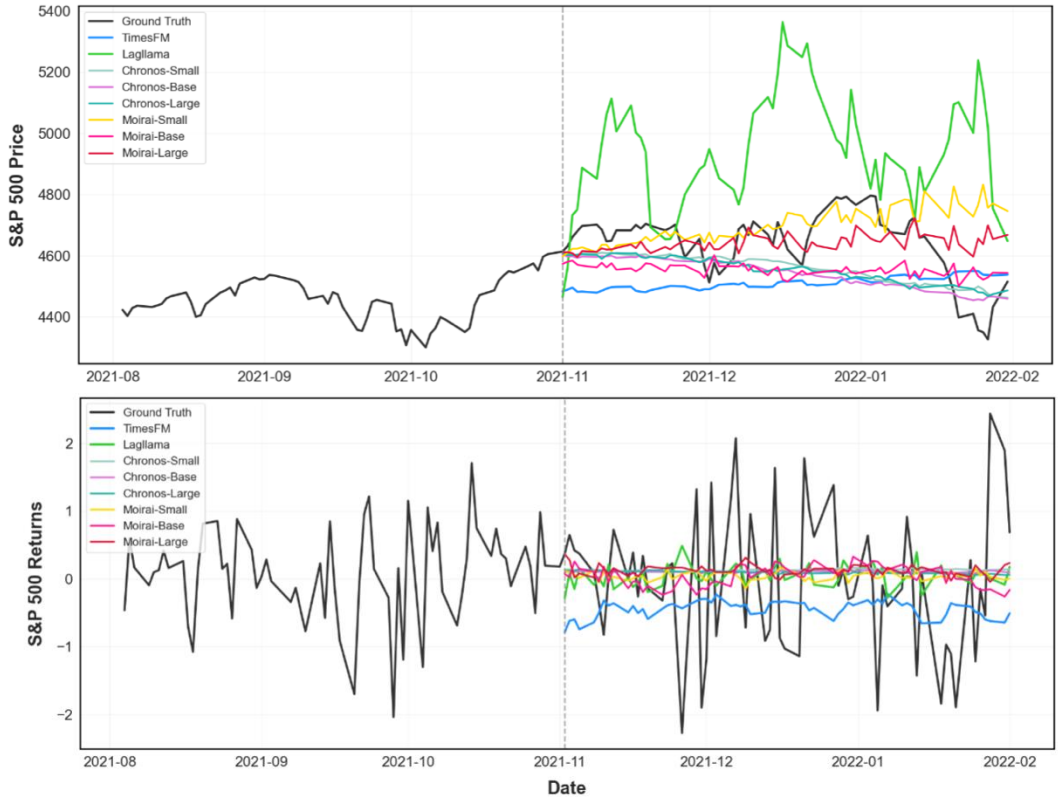
Universal failure: calm markets and crashes alike

IMPERIAL

8 TSFMs Chronos S/B/L · Moirai S/B/L · TimesFM · Lag-Llama Context window 252 days · Forecast horizon 63 days · S&P 500

NORMAL TRADING DAY

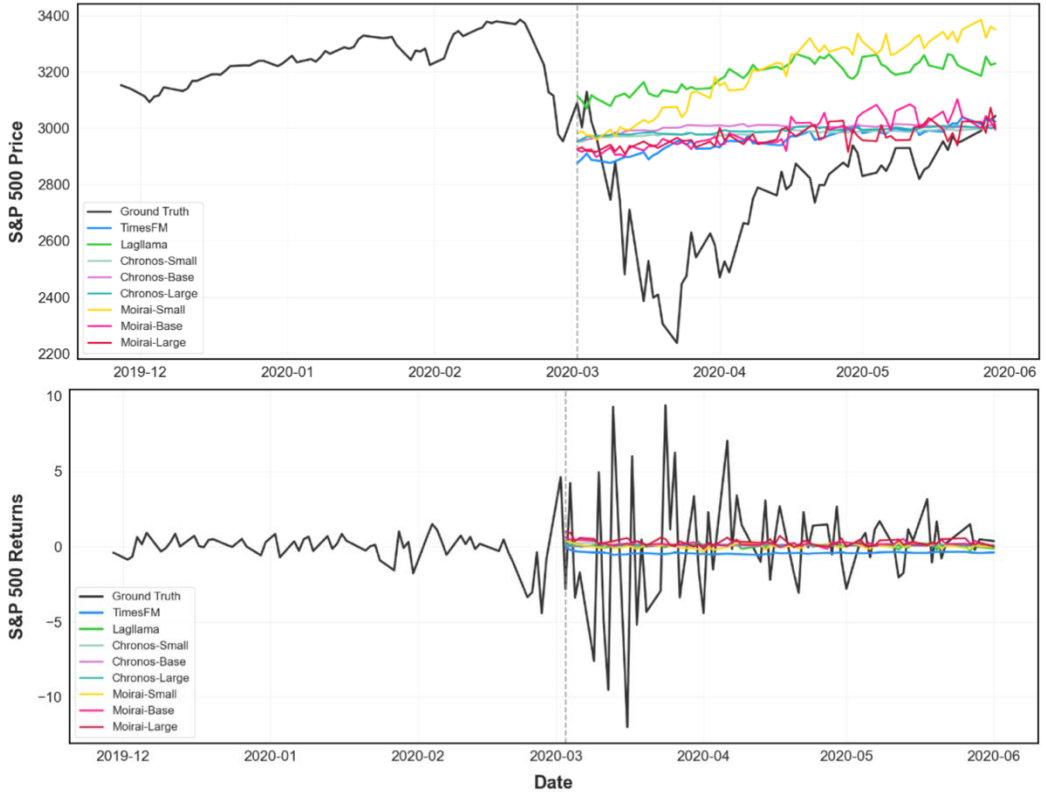
Even in calm markets, predictions diverge architecturally



Top: prices · Bottom: returns · Black = ground truth

COVID-19 CRASH · MARCH 2020

S&P 500 falls 35% in 23 days — every model predicts flat or upward



Top: prices · Bottom: returns · Black = ground truth

Different architectures. Different scales. Same failure.

Diagnosis: three losses, one shared smoothing pressure

IMPERIAL

TSFM	What it optimizes	What it converges to on noisy data
Chronos	token cross-entropy on quantized bins	most-probable bin → context mean
Moirai	mixture-distribution log-likelihood	mean of the predictive mixture
Time-MoE	multi-resolution Huber regression	robust mean (Huber down-weights tails)

Three different objectives. All converge toward the mean.

The natural follow-up: just fine-tune?

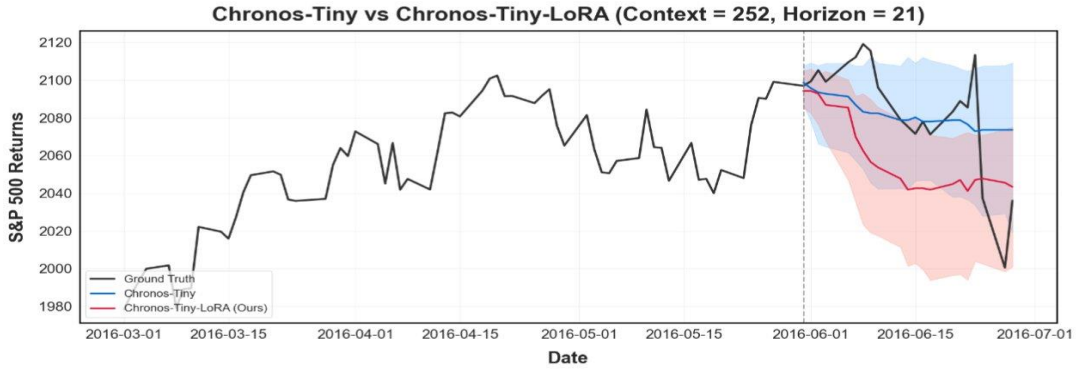
IMPERIAL

We fine-tune each TSFM with LoRA on financial data — and benchmark against the original frozen model across all 45 configurations.

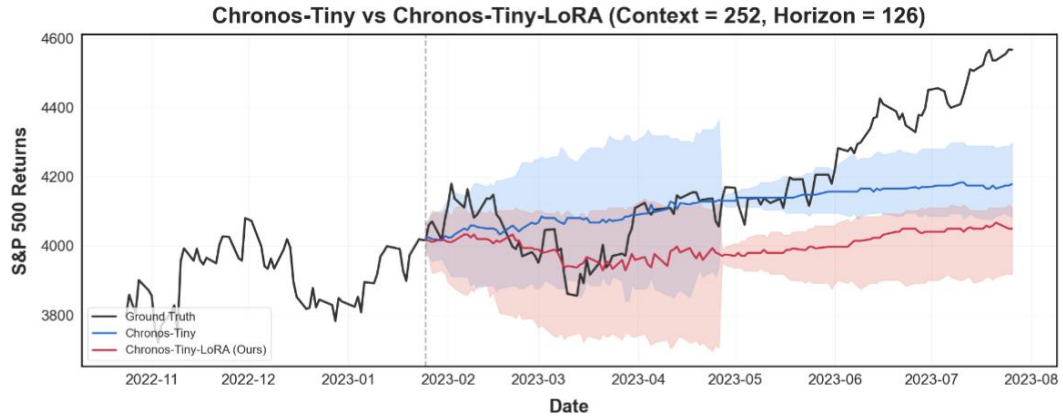
ILLUSTRATIVE EXAMPLE · CHRONOS-TINY ON S&P 500

LoRA shifts the prediction further from ground truth

Horizon = 21 days · Context = 252 days · Shaded band = 10–90% confidence interval



Horizon = 63 days · Same setup, longer horizon



FULL BENCHMARK · ORIGINAL vs LoRA

Three TSFMs · three assets · five horizons

Bold = better in row · Underlined = worse

Metrics		Chronos$_{large}$				Moirai$_{large}$				Time-MoE$_{base}$			
		Original		LoRA		Original		LoRA		Original		LoRA	
Asset	H	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
S&P500	5	**0.009**	**0.075**	0.011	0.078	**0.016**	**0.095**	0.023	0.110	0.012	0.088	**0.011**	**0.079**
	10	**0.016**	0.101	0.017	**0.099**	**0.021**	**0.111**	0.026	0.122	0.022	0.121	**0.017**	**0.103**
	21	**0.157**	**0.324**	1.894	1.339	**0.115**	**0.267**	0.898	0.329	**0.357**	**0.491**	1.330	1.098
	63	**1.164**	**0.894**	3.308	1.025	**0.388**	**0.487**	3.972	1.906	2.542	1.391	**2.293**	**1.389**
	126	**5.536**	**2.026**	6.493	2.457	**1.566**	**0.996**	5.550	2.188	9.257	2.761	**2.511**	**1.326**
WTI	5	**0.175**	**0.319**	0.252	0.373	0.351	0.447	**0.291**	**0.412**	**0.198**	**0.343**	0.202	0.341
	10	**0.278**	**0.410**	0.445	0.497	0.554	0.574	**0.458**	**0.522**	**0.313**	**0.438**	0.319	0.441
	21	**0.407**	**0.503**	0.586	0.627	**0.409**	**0.502**	1.338	0.965	**0.370**	**0.478**	0.395	0.501
	63	0.907	0.771	**0.819**	**0.715**	**1.338**	**0.943**	3.319	1.505	2.092	1.223	**1.897**	**1.134**
	126	**1.100**	**0.844**	1.176	0.876	**1.813**	**1.047**	2.174	1.190	4.010	**1.523**	**3.630**	1.569
EURUSD	5	**0.295**	**0.418**	0.315	0.432	0.627	0.629	**0.596**	**0.619**	0.419	0.500	**0.372**	**0.476**
	10	**0.555**	0.584	0.631	**0.593**	1.093	0.845	**0.917**	**0.779**	0.786	0.696	**0.756**	**0.685**
	21	**1.017**	**0.802**	1.347	0.896	**1.052**	**0.801**	1.733	1.030	1.224	0.891	**1.018**	**0.773**
	63	**3.381**	**1.404**	8.423	2.119	**3.263**	**1.298**	4.872	1.714	**4.626**	**1.614**	4.683	1.801
	126	**3.328**	**1.384**	4.833	1.736	**2.621**	**1.267**	6.181	2.066	**2.956**	**1.284**	6.405	2.067

LoRA degrades performance in 29 of 45 configurations.

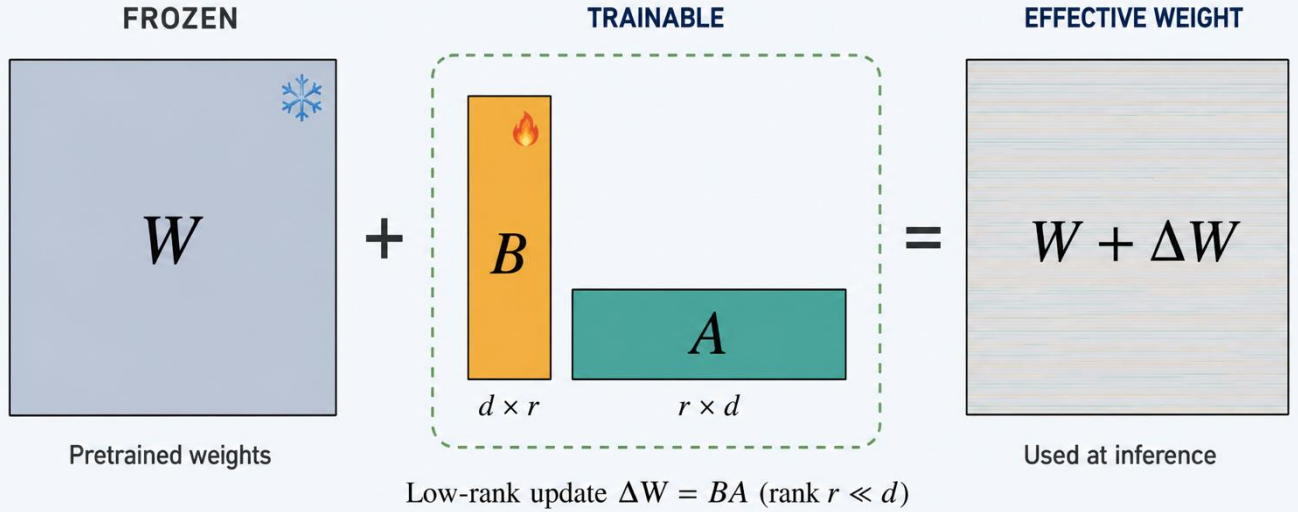
LoRA inherits the smoothing pressure

IMPERIAL

Fine-tuning adjusts the weights — but keeps every other structural choice that biases the model toward smoothing.

THE LoRA UPDATE

A low-rank perturbation added to the same frozen weights



LoRA modifies the weights of the same model. Architecture and training objective are unchanged.

WHAT FINE-TUNING CAN'T CHANGE

Three structural inheritances

- **Same architecture**

  LoRA preserves the transformer stack — the same inductive bias toward smooth, mean-reverting outputs.

- **Same training objective**

  Token-CE / Huber / mixture-likelihood — all still reward predicting the conditional mean.

- **Same family of solutions**

  ΔW perturbs the weights, but the model can only reach outputs the original architecture and loss can produce.

LoRA is a small perturbation on a model that fundamentally doesn't want to predict financial dynamics.

TSFM forecasts: suboptimal but informative

IMPERIAL

The forecast isn't noise. It's a coarse signal — wrong in predictable ways, useful in others.



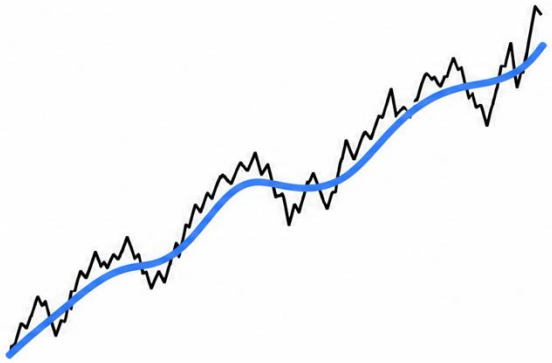
BLURRY PHOTO

the scene is recognizable;
fine detail is lost.



NOISY AUDIO

the speech is intelligible;
fine texture is lost.



COARSE FORECAST

the trend is approximately right;
fine variation is lost.

01

Predictable failure modes

Wrong in specific, structural ways — smoothing, mean reversion. Not random noise.

02

Real signal underneath

Trend direction and scale are usually approximately right. Exploitable information.

03

A starting point worth refining

Better than the unconditional mean. Worth using as input to a separate refinement.

Don't change the model — repair its output.

We propose RefineBridge

IMPERIAL

A Schrödinger Bridge module that transports a frozen TSFM forecast toward the data distribution. The foundation model never moves; only the bridge is trained.

OUR PROPOSAL · ONE TRAINED BRIDGE OVER A FROZEN FOUNDATION MODEL

RefineBridge

A small **Schrödinger Bridge** learns the optimal transport from a **coarse-but-informative TSFM forecast** toward the data distribution — **only the bridge is trained.**

WHY A SCHRÖDINGER BRIDGE · SAME MACHINERY AS DIFFUSION, BUT THE PRIOR IS NOT NOISE

DIFFUSION

Pure noise · N(0, I)



reverse a known forward noising

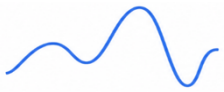


Data

The model must invent every structural feature of the data from scratch.

SCHRÖDINGER BRIDGE

TSFM forecast · any prior



learn KL-minimal optimal transport



Data

The prior already has trend and scale roughly right — the bridge only learns the residual.

Same generative machinery — but starting from an informative prior, the bridge has far less to learn.

The Schrödinger Bridge, formally

IMPERIAL

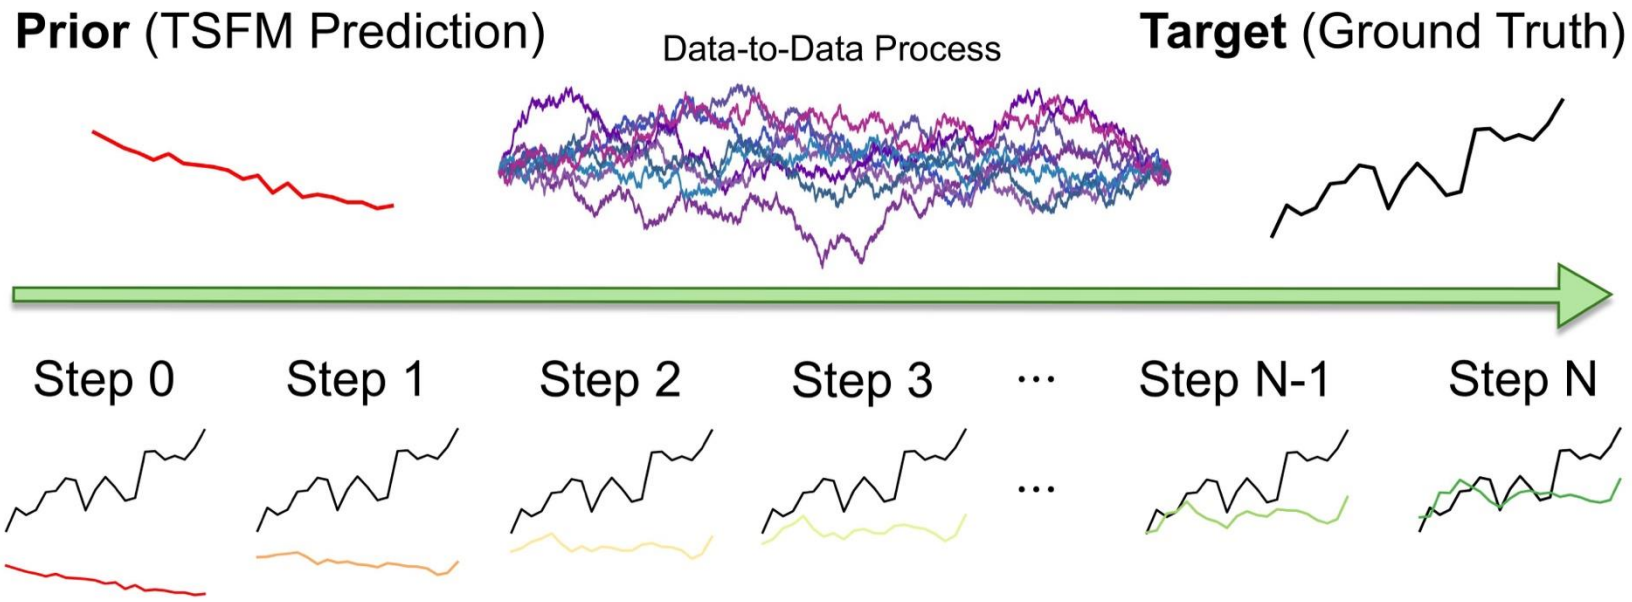
Among all stochastic processes that bridge prior → data, the SB is the one closest to a reference Brownian motion.

WHAT THE BRIDGE LEARNS · THE OPTIMIZATION PROBLEM

$$\pi^* = \arg\min_{\pi} D_{\mathrm{KL}}(\pi \,\|\, \pi_{\mathrm{ref}}) \quad \text{s. t.} \quad \pi_0 = p_{\mathrm{prior}}, \quad \pi_1 = p_{\mathrm{data}}$$

Read it as: among all path measures π that bridge prior → data, pick the one closest in KL divergence to a reference Brownian motion.

A LEARNED TRAJECTORY · STEP 0 = TSFM PRIOR (RED), STEP N = DATA (BLACK), BROWNIAN-LIKE TRANSPORT IN BETWEEN



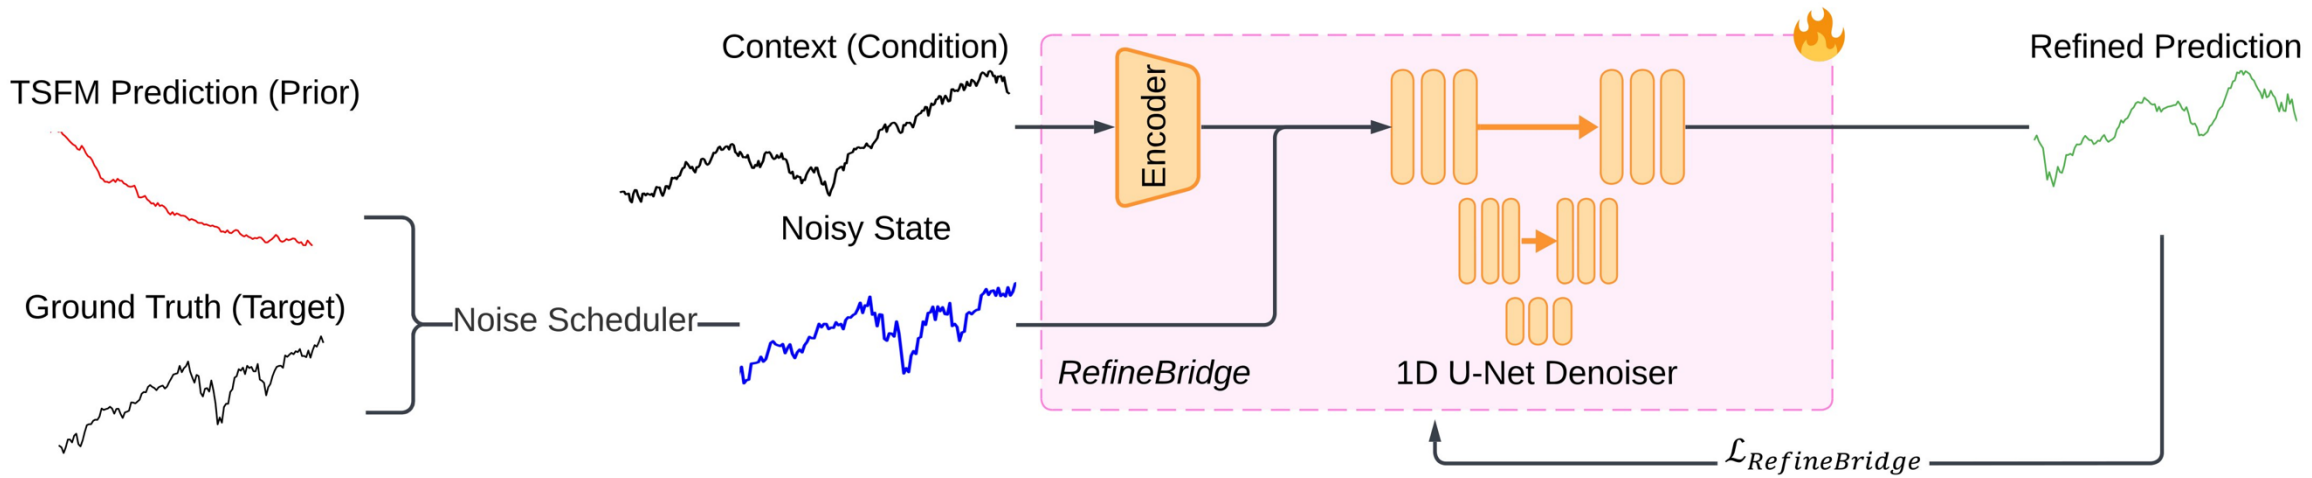
SB keeps what the prior already got right — the model only learns the residual to data.

Training RefineBridge

IMPERIAL

The Schrödinger bridge made concrete: a conditional 1D U-Net learns the transport from a frozen TSFM forecast to ground truth.

THE TRAINING PIPELINE · PRIOR + GROUND TRUTH → INTERPOLANT → DENOISER → REFINED PREDICTION



WHAT TO NOTICE · THREE DESIGN CHOICES THAT MAKE IT WORK

01

Stochastic interpolant

The noise scheduler interpolates between TSFM prior and ground truth, producing a noisy state along the bridge path.

02

Context conditioning

A learned encoder embeds the historical window — the denoiser sees what came before, not just the noisy interpolant.

03

Frozen TSFM, trained bridge

The foundation model never moves. Only the encoder and 1D U-Net are updated against the ground truth.

The bridge learns the residual. The foundation model never moves — train once, deploy across any TSFM forecast.

Training objective: one loss for both networks

IMPERIAL

The pipeline above is trained with a single loss — conditional denoising on the bridge marginal.

THE LOSS · CONDITIONAL DENOISING ON THE BRIDGE MARGINAL

$$\mathcal{L} = \mathbb{E}_{(c,\, x_0) \sim p_{\text{data}},\; z_c = \mathcal{E}(c)}\, \mathbb{E}_t\big[\, \| x_\theta(x_t,\, t,\, x_T,\, z_c) - x_0 \|_2^2 \,\big]$$

x_t

Bridge marginal — a noisy interpolation between prior and target at random time t.

x_T

TSFM prior — the frozen forecast that anchors one endpoint of the bridge.

z_c

Encoded context — DLinear features of the historical window c.

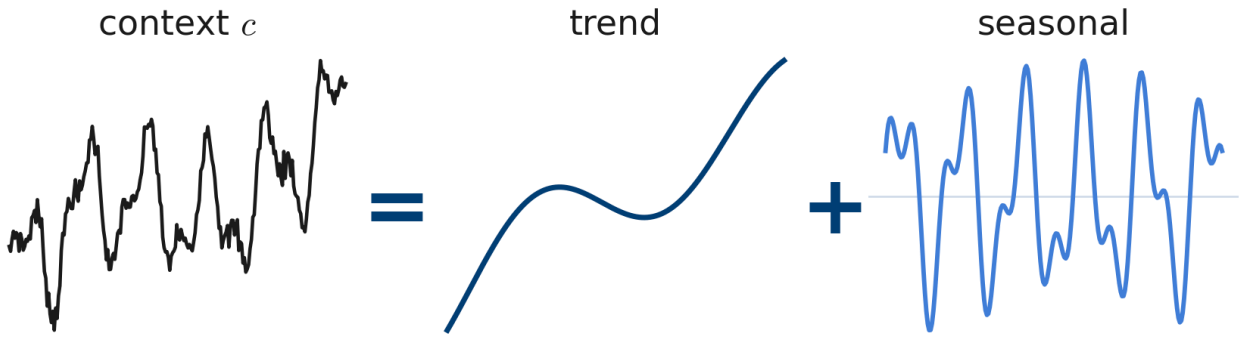
x_0

Ground truth target — the actual future the bridge transports toward.

THE CONTEXT ENCODER · DLINEAR DECOMPOSITION

Dlinear [1] splits the historical context into a **trend** (moving-average smoothing) and a **seasonal** (residual) component. Each is mapped through its own linear projection.

The recombined latent provides **one feature vector per refinement step**, conditioning the U-Net at every step of the bridge.



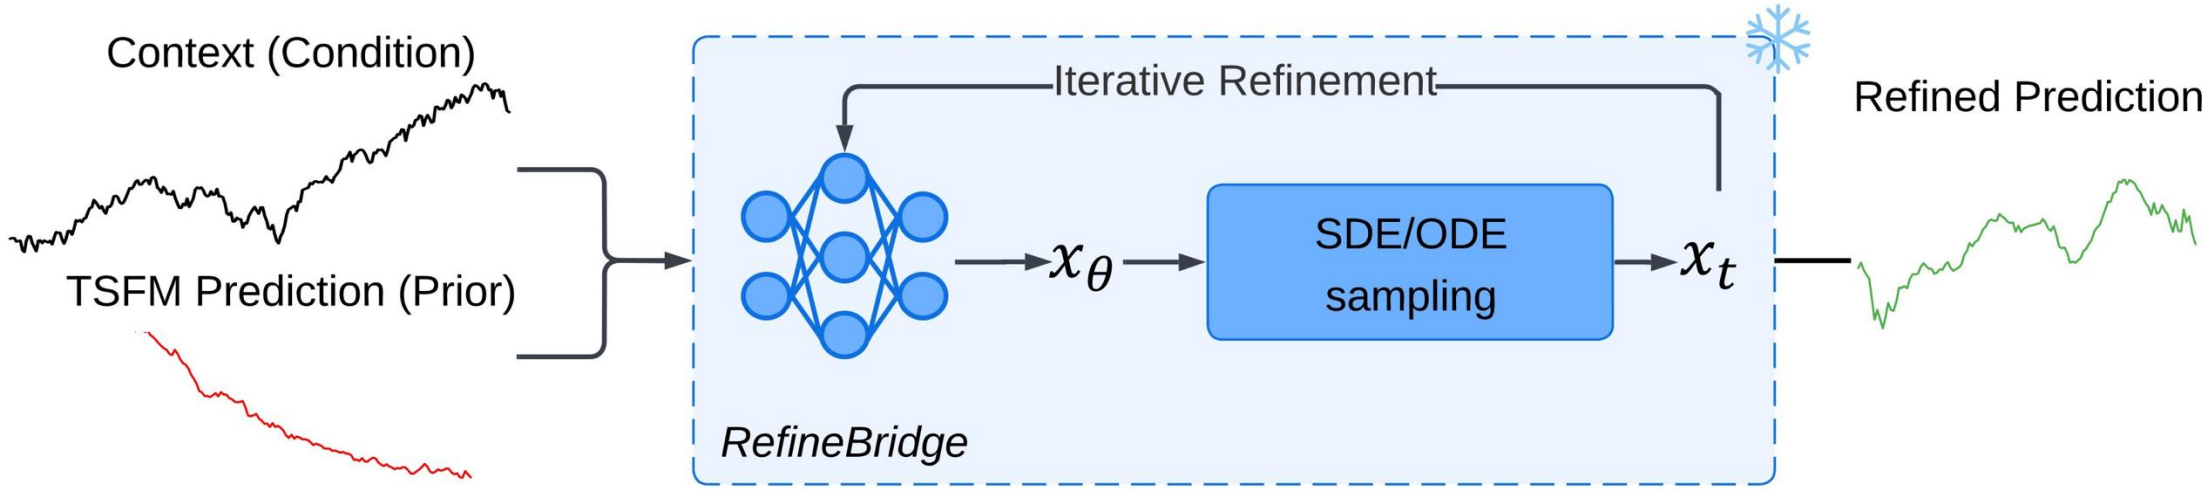
One loss trains both: the U-Net directly, the DLinear encoder implicitly — gradients flow through the same backward pass.

Inference: iterative refinement

IMPERIAL

A single trained bridge — frozen at test time — refines any TSFM prior step by step. Context plus prior in, polished forecast out.

THE INFERENCE LOOP · CONTEXT + PRIOR → BRIDGE → REFINED PREDICTION



01 INITIALISE

Initialise the bridge state with the TSFM prior at s = 1. The encoder produces context z_c from the historical window.

02 STEP

The network predicts the target — the x_θ output. The sampler then takes one ODE or SDE update.

03 EMIT

Iterate the schedule from s = 1 down to s = 0. The final state is the refined forecast.

Two samplers, one frozen bridge

IMPERIAL

Same network, two ways to step. ODE walks deterministically toward the target; SDE adds calibrated noise to escape misleading priors.

PAPER EQUATIONS 7 & 8 · EULER DISCRETISATIONS OF THE BRIDGE SDE

ODE

· deterministic refinement

$$\mathbf{x}_t = \frac{\sigma_t \bar{\sigma}_t}{\sigma_s \bar{\sigma}_s}\mathbf{x}_s + \frac{1}{\sigma_1^2}\left(\sigma_t^2 - \frac{\sigma_s \sigma_t \bar{\sigma}_t}{\bar{\sigma}_s}\right)\mathbf{x}_T + \frac{1}{\sigma_1^2}\left(\bar{\sigma}_t^2 - \frac{\bar{\sigma}_s \sigma_t \bar{\sigma}_t}{\sigma_s}\right)x_\theta(\mathbf{x}_s,\, s,\, \mathbf{x}_T,\, \mathbf{z}_c)$$

Eq. 7 · ODE sampler

At each step the network predicts **where the target lies** (the $x\theta$ term). The update walks **deterministically** toward that prediction, anchored by the prior — same input, same trajectory.

SHORT HORIZONS · H ∈ {5, 10} · 1–10 STEPS

SDE

· stochastic refinement

$$\mathbf{x}_t = \frac{\sigma_t^2}{\sigma_s^2}\mathbf{x}_s + \left(1 - \frac{\sigma_t^2}{\sigma_s^2}\right)x_\theta(\mathbf{x}_s,\, s,\, \mathbf{x}_T,\, \mathbf{z}_c) + \sigma_t\sqrt{1 - \frac{\sigma_t^2}{\sigma_s^2}}\;\varepsilon_\tau$$

Eq. 8 · SDE sampler

Same per-step prediction, **plus Gaussian noise** scaled by σt at every step. The randomness frees the trajectory to **explore** — recovers from priors that point the wrong way.

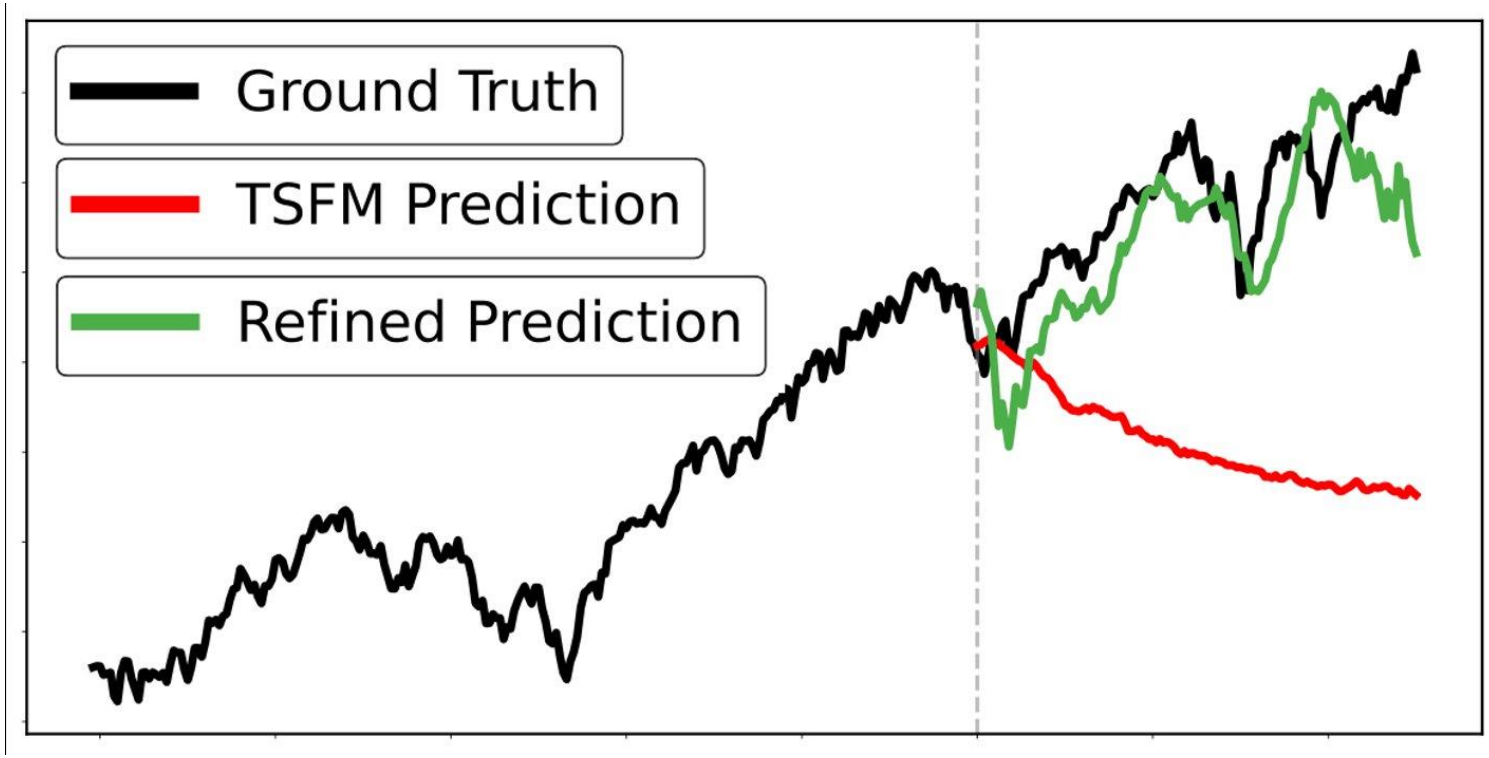
LONG HORIZONS · H ∈ {21, 63, 126} · 10–1000 STEPS · τ ∈ [0.01, 2.0]

Recovering from a misleading prior

IMPERIAL

When the TSFM points the wrong way, the SDE sampler can still recover the trend. Here, Chronos-large predicts a sustained decline; RefineBridge tracks the actual rise.

S&P 500 · CONTEXT 252 DAYS · HORIZON 126 DAYS · CHRONOS-LARGE · SDE SAMPLER



Vertical dashed line marks the context / forecast boundary. Time axis: trading days.

GROUND TRUTH

S&P 500 continues its upward trend through the entire 126-day prediction window — index moves higher across the right-hand panel.

TSFM PRIOR · Chronos-large

Predicts a sustained **decline.** Direction is entirely wrong — effectively an **uninformative prior.**

REFINED · RefineBridge

Recovers the trend within the first few steps and **tracks the actual rise.** The SDE sampler's noise lets the bridge **escape the wrong-pointing prior.**

Quantitative results

IMPERIAL

Forecasting MSE and MAE for three TSFM backbones, three financial assets, and five prediction horizons. Lower is better.

TABLE 1 · CHRONOS / MOIRAI / TIME-MoE · S&P 500, WTI, EUR/USD · H ∈ {5, 10, 21, 63, 126} · AVERAGED OVER 5 RUNS

Metrics		Chronos$_{large}$						Moirai$_{large}$						Time-MoE$_{base}$					
		Original		LoRA		**RefineBridge**		Original		LoRA		**RefineBridge**		Original		LoRA		**RefineBridge**	
Asset	H	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
S&P500	5	0.009	0.075	0.011	0.078	**0.008**	**0.067**	0.016	0.095	0.023	0.110	**0.008**	**0.069**	0.012	0.088	0.011	0.079	**0.009**	**0.073**
	10	0.016	0.101	0.017	0.099	**0.014**	**0.092**	0.021	0.111	0.026	0.122	**0.017**	**0.098**	0.022	0.121	**0.017**	**0.103**	0.018	0.105
	21	0.157	0.324	1.894	1.339	**0.105**	**0.260**	0.115	0.267	0.898	0.329	**0.103**	**0.251**	0.357	0.491	1.330	1.098	**0.279**	**0.426**
	63	1.164	0.894	3.308	1.025	**0.567**	**0.601**	0.388	0.487	3.972	1.906	**0.297**	**0.427**	2.542	1.391	2.293	1.389	**1.328**	**1.007**
	126	5.536	2.026	6.493	2.457	**2.102**	**1.207**	1.566	0.996	5.550	2.188	**0.764**	**0.691**	9.257	2.761	**2.511**	**1.326**	3.338	1.603
WTI	5	0.175	0.319	0.252	0.373	**0.146**	**0.293**	0.351	0.447	0.291	0.412	**0.178**	**0.321**	0.198	0.343	0.202	0.341	**0.146**	**0.295**
	10	0.278	0.410	0.445	0.497	**0.232**	**0.373**	0.554	0.574	0.458	0.522	**0.280**	**0.408**	0.313	0.438	0.319	0.441	**0.237**	**0.378**
	21	0.407	0.503	0.586	0.627	**0.335**	**0.469**	0.409	0.502	1.338	0.965	**0.325**	**0.458**	0.370	0.478	0.395	0.501	**0.316**	**0.457**
	63	0.907	0.771	0.819	**0.715**	**0.774**	0.721	1.338	0.943	3.319	1.505	**0.891**	**0.769**	2.092	1.223	1.897	1.134	**0.928**	**0.798**
	126	1.100	0.844	1.176	0.876	**0.844**	**0.742**	1.813	1.047	2.174	1.190	**1.072**	**0.863**	4.010	1.523	3.630	1.569	**1.160**	**0.904**
EURUSD	5	**0.295**	**0.418**	0.315	0.432	0.333	0.443	0.627	0.629	0.596	0.619	**0.327**	**0.447**	0.419	0.500	0.372	0.476	**0.339**	**0.453**
	10	0.555	0.584	0.631	0.593	**0.509**	**0.557**	1.093	0.845	0.917	0.779	**0.567**	**0.599**	0.786	0.696	0.756	0.685	**0.566**	**0.597**
	21	1.017	0.802	1.347	0.896	**0.888**	**0.756**	1.052	0.801	1.733	1.030	**0.897**	**0.748**	1.224	0.891	**1.018**	**0.773**	1.043	0.806
	63	3.381	1.404	8.423	2.119	**1.596**	**0.966**	3.263	1.298	4.872	1.714	**1.730**	**1.008**	4.626	1.614	4.683	1.801	**1.951**	**1.047**
	126	3.328	1.384	4.833	1.736	**1.702**	**1.004**	2.621	1.267	6.181	2.066	**1.272**	**0.852**	2.956	1.284	6.405	2.067	**1.458**	**0.953**
1st Count		2		1		**27**		0		0		**30**		0		6		**24**	

81 / 90

BEST RESULTS

across all cells (3 TSFMs × 3 assets × 5 horizons × MSE/MAE)

29 / 45

WHERE LoRA FAILS

configurations where LoRA degrades the zero-shot prediction

−11% → −62%

MSE REDUCTION GROWS WITH H

Chronos on S&P 500, prediction horizon H = 5 vs. H = 126

Qualitative refinements

IMPERIAL

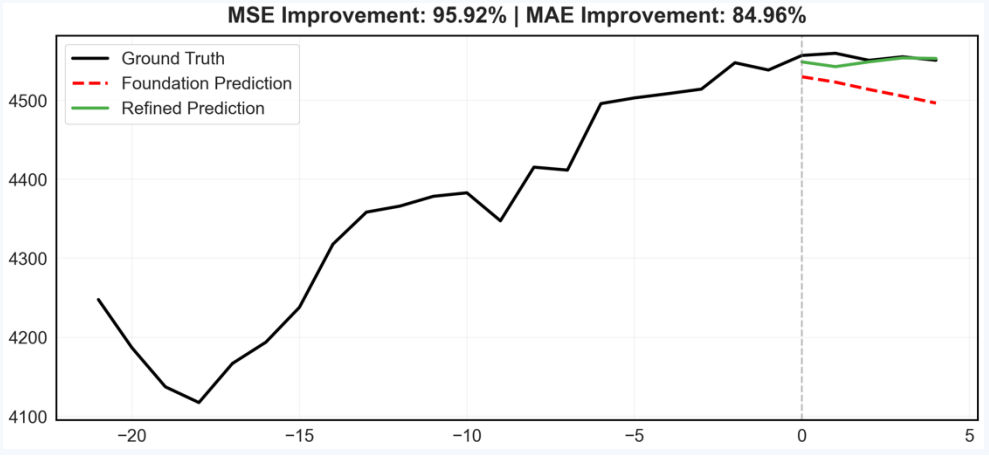
Four S&P 500 examples spanning Chronos & MOIRAI, short to long horizons.

QUALITATIVE EXAMPLES

Ground truth TSFM prior RefineBridge

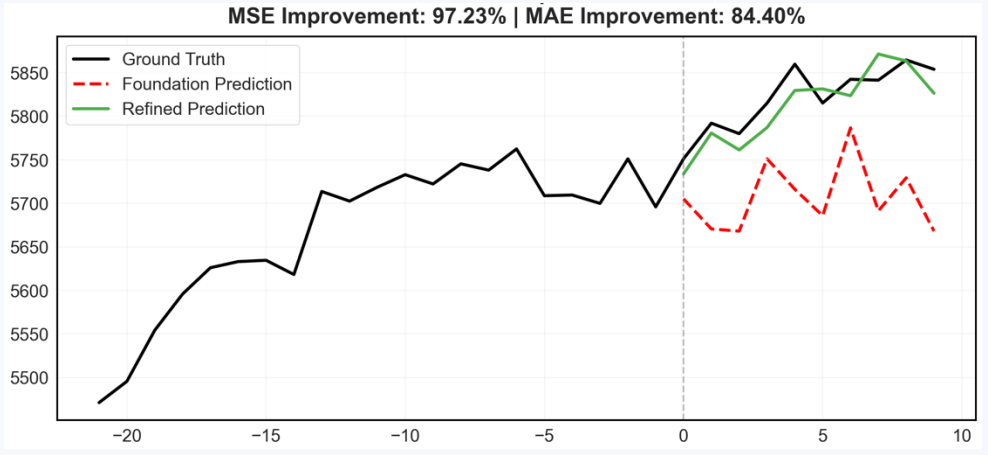
EXAMPLE 01

CHRONOS · 21-day context → 5-day horizon



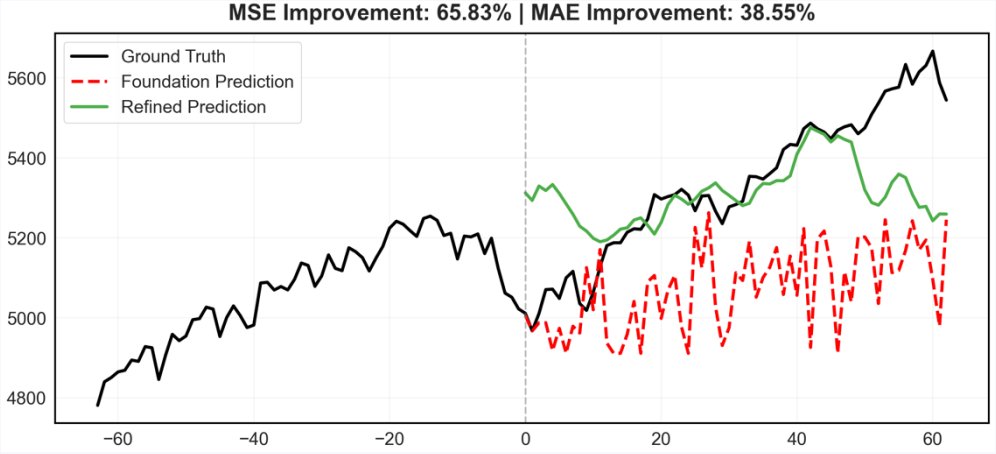
EXAMPLE 02

MOIRAI · 21-day context → 10-day horizon



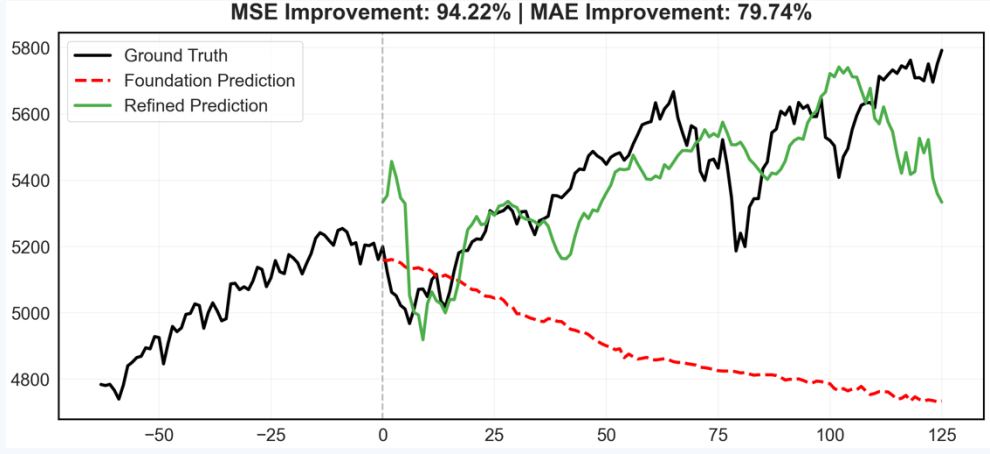
EXAMPLE 03

MOIRAI · 252-day context → 63-day horizon



EXAMPLE 04

CHRONOS · 252-day context → 126-day horizon



Conclusion

IMPERIAL

Bridge-based refinement improves frozen foundation-model forecasts on finance — no fine-tuning of the backbone, no architectural changes.

CONTRIBUTIONS

A new method.

RefineBridge: a Schrödinger Bridge module that refines frozen TSFM predictions on financial time series. Each TSFM prediction is treated as a coarse-but-informative prior — the bridge transports it to the data distribution learned from history.

Distinct from LoRA.

RefineBridge is an independent post-hoc module, not weight-level fine-tuning. The training objective uses the TSFM prediction as conditioning and learns to capture the target distribution, in contrast to LoRA's mean-seeking objective.

Empirical results.

Across three TSFMs (Chronos, Moirai, Time-MoE), three finance assets (S&P 500, WTI, EUR/USD), and five horizons, RefineBridge achieves the best result in 81 of 90 metric cells and improves over LoRA at every backbone.

Refine, don't replace — a small bridge transforms a coarse foundation-model prior into a forecast that tracks the data.

Future work and limitations

IMPERIAL

Where the framework could go next, and the practical caveats of the current design we are most aware of.

FUTURE WORK

→ Schrödinger Bridges across time-series problems.

RefineBridge is one instance of data-to-data optimal transport on time series. The same machinery applies wherever paired distributions exist — imputation, denoising, super-resolution, domain adaptation.

→ A foundation model built on the bridge.

Replace the external TSFM prior with the bridge itself — train a Schrödinger Bridge directly as a generative time-series foundation model rather than as a refiner of one. Conditional generation and asset transfer become natural.

→ Probabilistic refinement.

The SDE sampler already produces samples, not just point predictions. Aggregating these into calibrated predictive distributions — useful for risk, VaR, and option pricing — is a direct extension.

LIMITATIONS

— Fixed prediction window.

The current DLinear / MLP-based encoder fixes the refinement horizon at training time. Adapting to arbitrary horizons at test time would require either retraining or a horizon-conditional encoder.

— Per-asset training.

A separate RefineBridge is trained for each asset. The method does not yet generalise zero-shot to unseen assets — addressing this is what motivates the universal-refiner direction on the left.

— Inference cost at long horizons.

The SDE sampler can use up to 1000 steps for H = 126. Each step is cheap, but total latency is non-trivial relative to a single TSFM forward pass; sampler distillation or fewer-step schedules would help.

Work in progress — the bridge framework opens more questions than it closes.

Thank You

Department of Electrical and Electronic Engineering
Imperial College London, United Kingdom