# **Efficiently Generating Correlated Sample Paths from Multi-step Time Series Foundation Models**

Ethan Baron\* New York University Boris Oreshkin Amazon Ruijun Ma Amazon Hanyu Zhang Amazon

Kari Torkkola Amazon

Michael W. Mahoney Amazon **Andrew Gordon Wilson** New York University & Amazon Tatiana Konstantinova

### **Abstract**

Many time series applications require access to multi-step forecast trajectories in the form of sample paths. Recently, time series foundation models have leveraged multi-step lookahead predictions to improve the quality and efficiency of multi-step forecasts. However, these models only predict independent marginal distributions for each time step, rather than a full joint predictive distribution. To generate forecast sample paths with realistic correlation structures, one typically resorts to autoregressive sampling, which can be extremely expensive. In this paper, we present a copula-based approach to efficiently generate accurate, correlated sample paths from existing multi-step time series foundation models in one forward pass. Our copula-based approach generates correlated sample paths orders of magnitude faster than autoregressive sampling, and it yields improved sample path quality by mitigating the snowballing error phenomenon.

## 1 Introduction

In recent years, time series foundation models (TSFMs) have received significant attention, in some cases outperforming statistical and deep learning methods on some public time series forecasting benchmarks. One common characteristic of top-performing TSFMs, including TimesFM [Das et al., 2024], Chronos-Bolt [Ansari et al., 2024], and TiRex [Auer et al., 2025], is their ability to make multi-step predictions in a single forward pass. That is, given an input time series  $x_{1:T}$ , these multi-step models produce probabilistic forecasts  $p_{\theta}(x_{T+i} \mid x_{1:T})$  simultaneously for multiple future horizons  $i=1,2,\ldots,H$ . This feature has allowed recent TSFMs to make improved predictions over longer horizons, and it has increased the efficiency of both training and inference.

Multi-step TSFMs produce marginal predictive distributions  $p_{\theta}(x_{T+i} \mid x_{1:T})$ , rather than a full joint predictive distribution  $p(x_{T+1:T+H} \mid x_{1:T})$ . This is an important limitation because in many practical scenarios, having access to a full joint distribution is crucial [Wen et al., 2021]. For instance, a joint predictive distribution allows practitioners to estimate conditional probabilities, such as  $p(x_{T+4} \mid x_{T+2} > c)$ , and to estimate statistics over arbitrary aggregation windows, such as the 90th percentile estimate  $q_{90}(x_{T+5} + x_{T+6} + x_{T+7})$ . Joint distributions can also be represented implicitly via  $sample\ paths$  — possible forecast trajectories  $x_{T+i:T+K}$  drawn from the predictive distribution. These sample paths are used as inputs in diverse downstream applications such as supply chain management [Madeka et al., 2018], reinforcement learning [Moerland et al., 2023], learning dynamic systems [Hewing et al., 2020], and energy systems [Zhang et al., 2025].

There are several possible approaches to extract a joint predictive distribution from multi-step TSFMs. A *naive* approach is to use an independence assumption and set

<sup>\*</sup>Work completed while at Amazon. Correspondence to eb4727@nyu.edu.

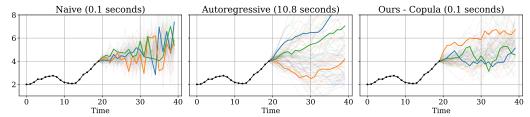


Figure 1: Mechanisms for generating sample paths from multi-step time series models. 100 sample paths are shown for each method, with three randomly-selected sample paths highlighted. The naive approach generates sample paths in  $\mathcal{O}(1)$  time but produces jagged trajectories with unrealistic correlation structures. Autoregressive sampling produces smooth correlated trajectories, but takes  $\mathcal{O}(N \cdot H)$  forward passes, which can be prohibitive, and can result in snowballing errors. Our copula-based approach provides realistic correlated sample paths in  $\mathcal{O}(1)$  time.

 $p(x_{T+1:T+H} \mid x_{1:T}) = \prod_{i=1}^{H} p_{\theta}(x_{T+i} \mid x_{1:T})$ . This approach can be implemented efficiently as it requires only single forward pass of a TSFM. However, the independence assumption is problematic in time series forecasting, where  $x_{T+i}$  typically has significant correlation with previous times. For example, if  $x_{T+1}$  is relatively high, we would expect  $x_{T+2}$  to also be relatively high. Since the naive approach does not account for these correlations, it results in highly unrealistic sample paths.

Alternatively, one could use an *autoregressive* sampling strategy by repeatedly sampling one step from the marginal predictive distribution: first sample  $\hat{x}_{T+1} \sim p_{\theta}(x_{T+1}|x_{1:T})$ , then  $\hat{x}_{T+2} \sim p_{\theta}(x_{T+2}|x_{1:T},\hat{x}_{T+1})$ , and so on. This autoregressive approach allows one to generate realistic sample paths with appropriate correlation structures, but is prohibitively expensive at scale: to generate N sample paths for a horizon length H we must perform  $\mathcal{O}(N \cdot H)$  forward passes of the model. It remains unclear how to efficiently produce realistic sample path forecasts with TSFMs.

In this paper, we present a method to generate realistic correlated sample paths from multi-step time series models in one forward pass. Our method is based on a *copula* decomposition of the joint predictive distribution, allowing us to leverage the accurate marginal predictive distributions of TSFMs and focus our efforts on learning a realistic correlation structure.

The sample paths generated with our method are competitive in quality to those generated autoregressively, but they are obtained significantly more cheaply (see Figure 1). In fact, we find that the copula-based sample paths are often *more* accurate than autoregressive sampling, achieved by mitigating the *snowballing error phenomenon* [Bachmann and Nagarajan, 2024]. Our copula-based approach allows time series practitioners to capitalize on advances in multi-step TSFMs for a wider variety of important use cases, greatly increasing the impact of TSFMs in practice.

#### 2 Decomposing joint predictive distributions with copulas

The core idea of our approach leverages the fact that any joint distribution can be factored into a set of marginal distributions and a *copula* function controlling the correlation structure. Formally, consider a collection of random variables  $X_1, \ldots, X_H$  with marginal CDFs  $F_1, \ldots, F_H$ , such that  $F_i(x) = p(X_i \leq x)$ . Then, by Sklar's Theorem [Sklar, 1959], there exists a copula  $C: [0,1]^H \mapsto [0,1]$  such that  $p(X_1 \leq x_1, \ldots, X_H \leq x_H) = C(F_1(x_1), \ldots, F_d(x_H))$ . Intuitively, the copula allows us to decompose the joint distribution of the random variables, modeling their correlation structure in *quantile space* while separately modeling their marginal distributions via  $F_i$ .

This copula decomposition is particularly compelling for multi-step time series forecasting. We already have access to high-quality marginal distributions  $F_i$  obtained from multi-step TSFMs, so all that remains is to model the correlation structure via a copula C. Moreover, time series have somewhat predictable correlation structures, such as autocorrelation, which we leverage to obtain a simple yet effective parameterization for the copula. Several prior works [Wen and Torkkola, 2019, Drouin et al., 2022, Ashok et al., 2024] use copulas for effective multivariate time series forecasting, whereas our focus is specifically on multi-step forecasting for univariate time series with TSFMs.

## 3 A copula-based approach for sample path generation

We propose a copula-based approach to efficiently generate correlated sample paths from multi-step TSFMs. Specifically, we use a TSFM to construct high-quality marginal predictive distributions for

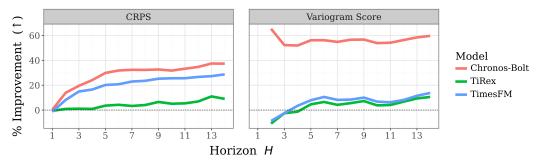


Figure 2: Median percent improvement from copula-based method by horizon for M4 Daily dataset. Switching from autoregressive sampling to our copula-based approach yields high-quality sample paths with realistic correlation structures, and significant improvements in CRPS for longer forecast horizons. Here, we plot the CRPS term at each horizon separately, rather than the cumulative sum.

each time step, and augment these marginals with a copula to impose a realistic correlation structure. To generate sample paths with this approach, we first sample vectors of correlated quantiles from the copula and then map these quantiles to the time series space with the marginal predictive distributions.

For the correlation structure, we use a Gaussian copula. Formally, we have  $C(F_1(x_1),\ldots,F_d(x_H))=\Phi_\Sigma\left(\Phi^{-1}(F_1(x_1)),\ldots,\Phi^{-1}(F_H(x_H))\right)$  where  $\Phi^{-1}$  is the inverse CDF of the standard normal distribution and  $\Phi_\Sigma$  is the joint CDF of the multivariate normal distribution with mean  $\mathbf{0}$  and covariance matrix  $\Sigma$ . In this work, we parameterize  $\Sigma$  as a symmetric Toeplitz matrix with a simple AR(1) correlation structure, such that  $\Sigma_{ij}=\rho^{|i-j|}$ . This parameterization reflects the autocorrelation structure common in time series. We obtain the parameter  $\rho$  by computing the empirical autocorrelation coefficient for each series,  $\mathrm{Corr}(x_{1:T-1},x_{2:T})$ . In Appendix B, we consider instead learning the parameterization of  $\Sigma$  using a pre-trained lightweight neural network. We find that pre-trained neural networks can predict copula parameterizations that further improve the qualify of sample paths, opening a promising direction for further research.

Note that the TSFMs we consider in this paper do not provide full marginal predictive distributions  $F_i$ , instead summarizing their marginal predictions to a pre-defined set of quantile knots such as  $Q = \{0.1, 0.2, \ldots, 0.9\}$ . To convert these predictions to full marginal distributions, we adopt the incremental quantile function approach from Park et al. [2022]. Specifically, we fit a full marginal CDF by linearly interpolating between the predicted quantile knots and using exponential decay extrapolation beyond the minimum/maximum knots. We set the decay towards 0 on the left as we exclusively consider non-negative time series in this paper, but this can be adapted for other settings. See Appendix C for an illustrative example.

### 4 Evaluating multi-step predictions via sample paths

To evaluate joint predictive distributions, we must capture both the accuracy of the marginal predictive distributions and the realism of the correlation structure across time steps. As is standard, we use the continuous ranked probability score (CRPS) to evaluate the marginal distributions. The CRPS for a joint predictive distribution p and observed vector  $\mathbf{x} \in \mathbb{R}^H$  is defined as:

$$CRPS(p, \mathbf{x}) = \sum_{i=1}^{H} \mathbb{E}_{\mathbf{y} \sim p} \left[ |\mathbf{x}_i - \mathbf{y}_i| \right] - \frac{1}{2} \mathbb{E}_{\mathbf{y}, \mathbf{z} \sim p} \left[ |\mathbf{y}_i - \mathbf{z}_i| \right]$$

To evaluate the correlation structure of the joint predictive, we use the variogram score (VS):

$$VS(p, \mathbf{x}) = \sum_{i=1}^{H} \sum_{j=1}^{H} (|\mathbf{x}_i - \mathbf{x}_j|^{0.5} - \mathbb{E}_{\mathbf{y} \sim p} |\mathbf{y}_i - \mathbf{y}_j|^{0.5})^2$$

Realistic multi-step predictions must achieve low scores on both CRPS and VS [Ziel and Berk, 2019].

# 5 Efficiently generating sample paths with multi-step forecasting models

We generate 10 sample paths with each approach for time series in the M1, M3, M4, and Tourism collections from the Monash repository [Godahewa et al., 2021]. These datasets include diverse domains, granularities, and forecast horizons. See Table A1 for more details on these datasets.

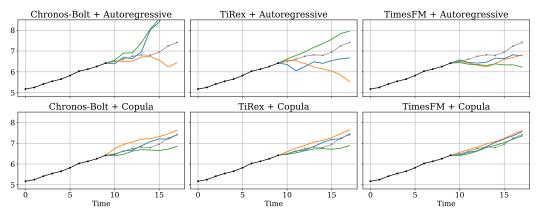


Figure 3: Copula-based sampling mitigates snowballing errors in autoregressive sampling. For each method, we plot 3 random sample paths. The true future trajectory of the time series is shown in gray.

To show that our copula-based approach can be effectively applied to a variety of backbone models, we evaluate performance using several TSFMs with varying architectures: the decoder-only transformer TimesFM [Das et al., 2024], the encoder-decoder transformer Chronos-Bolt Tiny [Ansari et al., 2024], and the xLSTM-based TiRex [Auer et al., 2025]. While M4 data was included in TimesFM's training corpus, this does not affect our conclusions as we primarily contrast copula-based and autoregressive sampling for a given TSFM backbone.

In Figure 2, we show the percent improvement in quality for the copula-based sample baths compared to those generated via autoregressive sampling for the M4 Daily dataset; we show similar patterns hold across other datasets in Appendix A. The copula-based methods produce sample paths with comparable or improved correlation structures compared to autoregressive sampling, as measured by VS. This finding confirms that our simple copula parameterization effectively captures the correlation structures prevalent in time series.

More surprisingly, the copula-based methods also improve the quality of the marginal predictive distributions, as measured by CRPS. We attribute this to the *snowballing error* phenomenon: biases appearing early in generation can amplify biases later on in the forecast horizon. This snowballing error phenomenon is one main argument in favor of multi-token prediction objectives in language modeling [Bachmann and Nagarajan, 2024].

In Figure 3, we contrast autoregressive and copula-based sample paths for an input time series with an approximately linear increasing trend. The autoregressively-generated sample paths frequently deviate from the approximately linear trend of the input time series, such as failing to continue the monotonically increasing pattern, or introducing sudden exponential growth, a result of snowballing errors. In contrast, the copula-generated sample paths realistically approximate the true trajectory.

Lastly, recall that generating sample paths via autoregressive sampling takes  $\mathcal{O}(N \cdot H)$  time, whereas the copula-based approach allows us to generate arbitrarily-many sample paths in  $\mathcal{O}(1)$  time. Unsurprisingly, we therefore observe that the speed of the copula-based methods is orders of magnitude faster than the autoregressive sampling approach. For example, for the M4 Daily dataset, copula-based sampling yielded speedups of 3.7x, 72.4x, and 100.4x for Chronos-Bolt, TiRex, and TimesFM respectively. Crucially, this speedup allows practitioners to adopt our approach for multistep sample path generation in large-scale production settings. For example, generating correlated sample paths for many time series with TabPFN-TS [Hoo et al., 2025] would be prohibitively slow using autoregressive sampling, but becomes practical with our method.

### 6 Conclusion

We present a copula-based technique to generate high-quality sample paths from multi-step TSFMs in one forward pass. We verify that our method produces realistic correlation structures and can even improve the accuracy of marginal predictive distributions compared to autoregressive sampling, at a fraction of the computational cost. This technique has potential to greatly accelerate the impact of TSFMs by extending their powerful capabilities to settings requiring joint predictive distributions.

## References

- Abdul Fatir Ansari, Caner Turkmen, Oleksandr Shchur, and Lorenzo Stella. Fast and accurate zero-shot forecasting with Chronos-Bolt and AutoGluon, 2024. URL https://aws.amazon.com/blogs/machine-learning/fast-and-accurate-zero-shot-forecasting-with-chronos-bolt-and-autogluon/.
- Arjun Ashok, Étienne Marcotte, Valentina Zantedeschi, Nicolas Chapados, and Alexandre Drouin. TACTiS-2: Better, faster, simpler attentional copulas for multivariate time series. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=xt0ydkE1Ku.
- Andreas Auer, Patrick Podest, Daniel Klotz, Sebastian Böck, Günter Klambauer, and Sepp Hochreiter. TiRex: Zero-shot forecasting across long and short horizons. In *1st ICML Workshop on Foundation Models for Structured Data*, 2025. URL https://openreview.net/forum?id=cWrbpOZGRa.
- Gregor Bachmann and Vaishnavh Nagarajan. The pitfalls of next-token prediction. In Forty-first International Conference on Machine Learning, 2024. URL https://openreview.net/forum?id=76zq8Wk16Z.
- Shaojie Bai, J Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv:1803.01271, 2018. URL https://arxiv.org/pdf/1803.01271.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. In Dekai Wu, Marine Carpuat, Xavier Carreras, and Eva Maria Vecchi, editors, *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/W14-4012. URL https://aclanthology.org/W14-4012/.
- Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for timeseries forecasting. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=jn2iTJas6h.
- Alexandre Drouin, Étienne Marcotte, and Nicolas Chapados. TACTiS: Transformer-attentional copulas for time series. In *ICML*, pages 5447-5493, 2022. URL https://proceedings.mlr.press/v162/drouin22a.html.
- Rakshitha Wathsadini Godahewa, Christoph Bergmeir, Geoffrey I. Webb, Rob Hyndman, and Pablo Montero-Manso. Monash time series forecasting archive. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. URL https://openreview.net/forum?id=wEc1mgAjU-.
- Lukas Hewing, Elena Arcari, Lukas P. Fröhlich, and Melanie N. Zeilinger. On simulation and trajectory prediction with Gaussian process dynamics. In Alexandre M. Bayen, Ali Jadbabaie, George Pappas, Pablo A. Parrilo, Benjamin Recht, Claire Tomlin, and Melanie Zeilinger, editors, *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pages 424–434. PMLR, 10–11 Jun 2020. URL https://proceedings.mlr.press/v120/hewing20a.html.
- Shi Bin Hoo, Samuel Müller, David Salinas, and Frank Hutter. From tables to time: How TabPFN-v2 outperforms specialized time series forecasting models. *arXiv:2501.02945*, 2025. URL https://arxiv.org/pdf/25 01.02945.
- Dhruv Madeka, Lucas Swiniarski, Dean Foster, Leo Razoumov, Kari Torkkola, and Ruofeng Wen. Sample path generation for probabilistic demand forecasting. In 4th Workshop on Mining and Learning from Time Series, 2018. URL https://milets18.github.io/papers/milets18\_paper\_14.pdf.
- Thomas M Moerland, Joost Broekens, Aske Plaat, Catholijn M Jonker, et al. Model-based reinforcement learning: A survey. *Foundations and Trends in Machine Learning*, 16(1):1–118, 2023. URL https://www.doi.org/10.1561/2200000086.
- Youngsuk Park, Danielle Maddix, François-Xavier Aubet, Kelvin Kan, Jan Gasthaus, and Yuyang Wang. Learning quantile functions without quantile crossing for distribution-free time series forecasting. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 8127–8150. PMLR, 28–30 Mar 2022. URL https://proceedings.mlr.press/v151/park22a.html.
- M Sklar. Fonctions de répartition à n dimensions et leurs marges. In *Annales de l'ISUP*, volume 8, pages 229–231, 1959.
- Ruofeng Wen and Kari Torkkola. Deep generative quantile-copula models for probabilistic forecasting. arXiv:1907.10697, 2019. URL https://arxiv.org/pdf/1907.10697.

- Zheng Wen, Ian Osband, Chao Qin, Xiuyuan Lu, Morteza Ibrahimi, Vikranth Dwaracherla, Mohammad Asghari, and Benjamin Van Roy. From predictions to decisions: The importance of joint predictive distributions. arXiv:2107.09224, 2021. URL https://arxiv.org/pdf/2107.09224.
- Hanyu Zhang, Reza Zandehshahvar, Mathieu Tanneau, and Pascal Van Hentenryck. Weather-informed probabilistic forecasting and scenario generation in power systems. *Applied Energy*, 384:125369, 2025. ISSN 0306-2619. doi: https://doi.org/10.1016/j.apenergy.2025.125369. URL https://www.sciencedirect.com/science/article/pii/S0306261925000996.
- Florian Ziel and Kevin Berk. Multivariate forecasting evaluation: On sensitive and strictly proper scoring rules. arXiv:1910.07325, 2019. URL https://arxiv.org/pdf/1910.07325.

## A Full results for all datasets

In Figure A1, we contrast the quality of copula-based sample paths to those generated via autoregressive sampling for the M4 Daily dataset, while we show full results for other datasets in Figures A2 and A3. The copula-based methods generally produce sample paths with comparable or improved correlation structures compared to autoregressive sampling, as measured by VS. The copula-based methods also frequently improve upon the marginal predictive distributions from autoregressive sampling, as measured by CRPS. The properties of these datasets are summarized in Table A1.

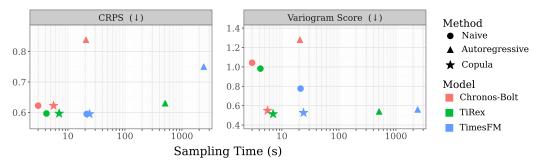


Figure A1: Sample path quality for M4 Daily dataset. Switching from autoregressive sampling to our copula-based approach yields improved sample paths with realistic correlation structures, at a fraction of the time. Points show the median performance across series after dividing by the corresponding score of the SeasonalNaive baseline. The *x*-axis indicates the time taken to generate 10 sample paths per time series on an A100 GPU.

Table A1: Summary of dataset properties.

	Frequency	Seasonality	Horizon	Series	Min Length	Max Length
M1	Monthly	12	18	617	48	150
	Quarterly	4	8	203	18	114
	Yearly	1	6	181	15	58
M3	Other	1	8	174	71	104
	Monthly	12	18	1428	66	144
	Quarterly	4	8	756	24	72
	Yearly	1	6	645	20	47
M4	Hourly	24	48	414	748	1008
	Daily	7	14	4227	107	9933
	Weekly	52	13	359	93	2610
	Monthly	12	18	48000	60	2812
	Quarterly	4	8	24000	24	874
	Yearly	1	6	23000	19	841
Tourism	Monthly	12	24	366	91	333
	Quarterly	4	8	427	30	130
	Yearly	1	4	518	11	47

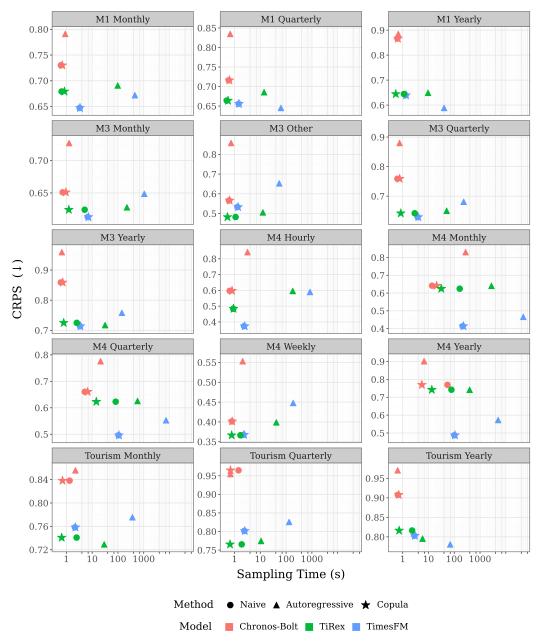


Figure A2: Median CRPS results for other datasets. The copula approach yields sample paths with comparable or higher quality than autoregressive sampling but is significantly faster.

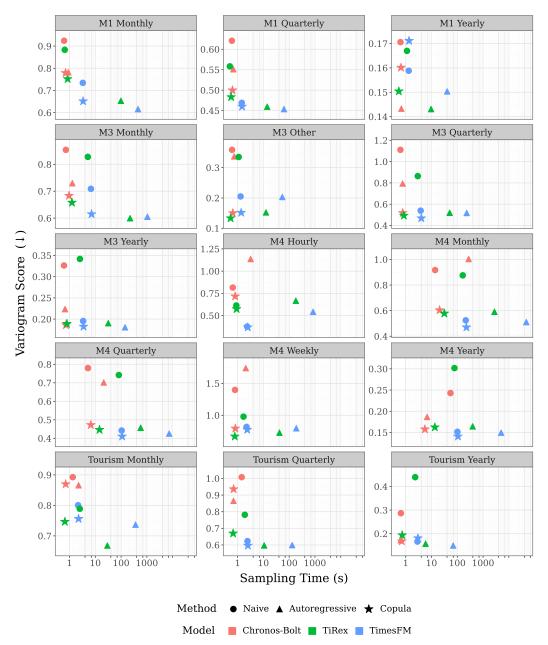


Figure A3: Median variogram score results for other datasets. The copula approach yields sample paths with comparable quality to autoregressive sampling but is significantly faster.

# **B** Pre-trained copula module

Above, we parameterize the Gaussian copula with a single parameter  $\rho$ , and obtain  $\rho$  by computing the autocorrelation coefficient for each time series ("Auto- $\rho$ "). Here, we explore the idea of a *pre-trained copula module*, a lightweight neural network trained to predict the Gaussian copula parameters for a given time series. As the backbone of the neural network, we consider:

- 1. An MLP with hidden sizes 64 and 32. The MLP takes as input the last 30 time steps.
- 2. A Temporal Convolution Network (TCN) [Bai et al., 2018] with hidden state size of 16, kernel size of 2, dilations of [1, 2, 4, 8]. The final hidden state is passed through an MLP with hidden size of 16 to predict the copula parameters. The TCN takes as input the last 30 time steps.
- 3. A Gated Recurrent Unit (GRU) [Cho et al., 2014], with hidden state size of 16. The final hidden state is passed through an MLP with hidden size of 16 to predict the copula parameters. The GRU takes at most the last 128 time steps as input and can handle shorter sequences as well.

For each backbone, we consider two possible parameterizations for the copula covariance matrix  $\Sigma$ :

- 1. Predicting  $\rho$  and setting  $\Sigma_{ij} = \rho^{|i-j|}$  as above.
- 2. Predicting  $\rho$  and  $\beta$ , and setting  $\Sigma_{ij} = (1 \beta)\rho^{|i-j|} + \beta \delta_{ij}$ . This parameterization is indicated with XXX-2.

We train the backbones for 10 epochs on M1 Monthly, M4 Hourly, M4 Daily, M4 Weekly, M4 Monthly, and Tourism Monthly. For each series, we generate 20 sample paths up to a horizon of H=8 and minimize the variogram score using Adam with a learning rate of 0.001. We scale each input time series by dividing by its maximum value. During training, we use the marginal predictive distributions of Chronos-Bolt Tiny.

In Figure A4, we compare the performance of each pre-trained copula module in generating accurate correlation structures on the M3 datasets.

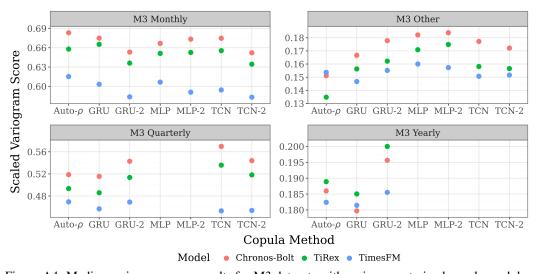


Figure A4: Median variogram score results for M3 datasets with various pre-trained copula modules, normalized by the SeasonalNaive score. Since the MLP and TCN backbones have constraints for the input sequence lengths, their scores are not available for some datasets.

The pre-trained copula modules sometimes yield improved sample paths compared to the "Auto- $\rho$ " approach presented above. In particular, the GRU backbone seems most promising, improving performance on three of the four evaluation datasets, and is also capable of handling inputs with varying lengths. Interestingly, despite being trained solely on the marginal predictions of Chronos-Bolt Tiny, the pre-trained copula modules often yield similar improvements in performance when

applied to the marginal predictive distributions of other models. This observation reinforces the promise of applying a pre-trained copula module to cheaply augment the predictions of any multi-step TSFM and produce correlated sample paths.

Lastly, while switching to the  $\rho,\beta$  parameterization provides more flexibility, in our experiments this does not translate into improved performance in practice. We suspect that more extensive pre-training of the modules, potentially including synthetic training data and longer forecast horizons, could address this shortcoming. Moreover, future research could explore more expressive parameterizations for Gaussian copulas, or even other types of copula structures.

# C Fitting empirical CDF to quantile predictions

The TSFMs we consider in this paper produce marginal predictions in the form of quantile predictions for  $Q = \{0.1, 0.2, \dots, 0.9\}$ . To convert these predictions to full marginal distributions, we adopt the incremental quantile function (IQF) approach from Park et al. [2022]. In Figure A5, we show an example of an empirical CDF fit to quantile predictions from a TSFM.

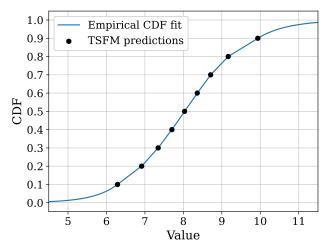


Figure A5: Example of empirical CDF fit to TSFM quantile predictions using IQF.