Decision-Oriented Probabilistic Demand Forecasting for Healthcare Supply Chains Using Quantile Regression Neural Networks

Boutayna Elghomri[1,2], Fayçal Messaoudi[2], Adnen El Amraoui[1], and Hamid Allaoui[1]

Abstract— **Accurate demand forecasting is essential in healthcare and pharmaceutical supply chains, where uncertainty and asymmetric risks may lead to stockouts, overstocking, service disruptions, or product waste. However, many operational forecasting approaches still rely on point predictions, which provide limited information about demand uncertainty. This paper proposes a decision-oriented probabilistic forecasting framework based on a Quantile Regression Neural Network (QRNN) for daily pharmaceutical demand forecasting. The model jointly estimates lower, median, and upper conditional quantiles, enabling adaptive prediction intervals around a robust central forecast. Unlike studies focused mainly on predictive accuracy, this work links quantile forecasts to inventory planning, service-level management, and risk mitigation. Using real-world daily pharmaceutical demand data, the QRNN is evaluated against a persistence-based benchmark. The results show improved point-forecast accuracy, with RMSE reduced by 18.01%, MAE by 19.94%, and sMAPE by 23.24%, while R^2 increases from 0.0266 to 0.3456. The 90% prediction interval achieves 93.39% empirical coverage, indicating slightly conservative but reliable uncertainty estimates. These findings show that QRNN-based probabilistic forecasting can provide actionable uncertainty information for more robust, risk-sensitive, and decision-oriented healthcare supply chain planning.**

Keywords: Probabilistic demand forecasting; Quantile Regression Neural Network; Healthcare supply chains; Pharmaceutical demand; Uncertainty modeling; Inventory planning; Service-level management.

I. Introduction

Demand forecasting is central to supply chain decision-making, as it supports inventory control, production planning, replenishment, service levels, product availability, cost management, and service continuity [1]–[3]. In healthcare and pharmaceutical supply chains, forecast accuracy is particularly critical because of strict service requirements, high stockout costs, and limited tolerance for disruptions [4], [5].

However, pharmaceutical demand is difficult to predict due to volatility, intermittent consumption, seasonality, promotions, unexpected shocks, epidemiological events, policy interventions, and behavioral changes [6]–[8]. Forecasting errors are therefore unavoidable and often asymmetric, since under-forecasting and over-forecasting generate different operational consequences [9]. While many operational methods still rely on point forecasts optimized with symmetric error criteria, such as mean squared error or mean absolute error [10], these forecasts do not explicitly represent uncertainty and provide limited support for risk-sensitive decisions [11]. This limitation is especially important in healthcare supply chains, where under-forecasting may lead to stockouts, whereas over-forecasting may increase holding costs, waste, and obsolescence [12].

Probabilistic forecasting addresses this issue by estimating the conditional distribution of future demand rather than a single expected value [13], [14]. Through prediction intervals or conditional quantiles, it can support safety-stock setting, service-level definition, and risk mitigation [15], [16]. Nevertheless, its operational use remains insufficiently developed in healthcare and pharmaceutical supply chains.

This paper proposes a decision-oriented probabilistic forecasting framework based on a Quantile Regression Neural Network (QRNN). Although QRNNs and quantile forecasting are established methods, the contribution of this study lies in operationalizing quantile forecasts for healthcare supply chain planning. In particular, P50 supports routine demand planning, while upper quantiles inform safety-stock adjustment, service-level management, and stockout-risk mitigation.

The study develops a QRNN-based framework for daily pharmaceutical demand, evaluates it against a persistence benchmark using accuracy and calibration criteria, and links quantile forecasts to inventory planning and service-level management. It addresses two research questions: **RQ1:** To what extent does QRNN-based probabilistic forecasting improve predictive performance and uncertainty calibration compared with persistence-based forecasting? **RQ2:** How can quantile forecasts support inventory planning, service-level management, and risk mitigation?

Using daily pharmaceutical demand data, the results show that the QRNN improves predictive performance while producing calibrated and dynamically adaptive prediction intervals, providing actionable uncertainty information for healthcare supply chain planning.

II. Related Works

A. Demand forecasting in healthcare and pharmaceutical supply chains

Demand forecasting in healthcare and pharmaceutical supply chains is a high-stakes task because demand uncertainty affects product availability, inventory costs, waste reduction, and service continuity [4], [5]. These systems operate under strict service-level requirements, regulatory constraints, product perishability, and limited tolerance for stockouts. Forecasting is therefore a core component of operational planning, especially as digital supply chain research increasingly emphasizes visibility, traceability, and

[1]Univ. Artois, UR 3926, LGI2A, Faculty of Applied Sciences, Béthune 62400, France.
[2]Sidi Mohamed Ben Abdellah University, ENSA Fez, LIASSE, Fez 30000, Morocco.
Corresponding author: boutayna.elghomri@usmba.ac.ma

decision support as drivers of resilience and transparency [17], [18].

Classical forecasting approaches, including time-series models and intermittent-demand methods, are widely used in pharmaceutical demand planning [6], [7]. Although interpretable and useful for stable demand patterns, they often struggle with abrupt demand shifts, nonlinear dynamics, product heterogeneity, seasonality, epidemiological events, policy changes, promotions, and unexpected disruptions.

Forecasting errors also generate asymmetric consequences. Under-forecasting may cause stockouts, service failures, and patient-related risks, whereas over-forecasting may increase holding costs, expiration risk, and financial losses [9]. This motivates forecasting approaches that account not only for average error reduction, but also for uncertainty and risk exposure.

B. From point forecasting to probabilistic forecasting

Traditional forecasting models mainly produce point estimates optimized under symmetric loss functions, such as mean squared error or mean absolute error [10]. Although useful for baseline planning, point forecasts provide limited information about future uncertainty and may be insufficient when overestimation and underestimation have different costs [11].

Probabilistic forecasting addresses this limitation by estimating the conditional distribution of future demand rather than a single expected value [13]. Through prediction intervals, quantiles, or full predictive distributions, it supports supply chain decisions where inventory policies, safety stocks, and service levels depend on both expected demand and demand uncertainty.

In supply chain management, probabilistic forecasts support inventory control, service-level optimization, and risk management [15], [16]. Recent studies emphasize uncertainty-aware forecasting as a basis for robust decision-making [19], [20]. However, probabilistic forecasts are still often evaluated methodologically, while their operational interpretation remains underdeveloped, especially in pharmaceutical supply chains.

C. Quantile regression and neural network-based approaches

Quantile regression provides a flexible, distribution-free framework for estimating conditional quantiles under asymmetric loss functions [21]. By estimating different parts of the demand distribution, it is relevant for supply chain decisions where lower, median, and upper quantiles correspond to different operational concerns.

Combined with neural networks, Quantile Regression Neural Networks (QRNNs) enable nonlinear estimation of multiple conditional quantiles within a unified learning architecture [22]. QRNNs are suitable for demand forecasting problems involving nonlinear relationships, heteroskedasticity, and changing variability. Unlike post-hoc uncertainty methods, they integrate uncertainty estimation directly into the learning process, improving coherence between central forecasts and prediction intervals.

QRNNs and related neural quantile approaches have shown potential in volatile and nonlinear environments [23]. In pharmaceutical supply chains, the median quantile can support routine replenishment, while upper quantiles can inform safety-stock adjustment and service-level protection. This decision-oriented view is consistent with recent work on digital twins and multi-agent systems, which highlights the role of simulation, decentralized intelligence, and data-driven coordination in sustainable and resilient supply chain planning [24], [25].

Data scarcity remains a challenge when applying deep learning models in healthcare and pharmaceutical supply chains. Limited observations, product-level sparsity, intermittent demand, and restricted access to operational data may reduce robustness and generalizability. Recent surveys identify mitigation strategies such as transfer learning, data augmentation, semi-supervised learning, self-supervised learning, few-shot learning, and regularization [26]. In this study, this issue is partially addressed through a compact feed-forward QRNN, causal lagged features, dropout, L_2 regularization, and early stopping.

D. Research gap and positioning of the present study

Despite methodological progress, probabilistic forecasting remains insufficiently integrated into operational decision-making, especially in healthcare and pharmaceutical supply chains [27]. Existing studies often focus on accuracy, calibration, or model comparison, with less attention to how probabilistic outputs can support concrete decisions. This limitation echoes broader digital supply chain research, where the bridge between technological innovation and practical adoption remains fragmented across AI accountability, digital twins, IoT-enabled supply chains, and multi-agent circular systems [17], [18], [24], [25]. Moreover, some approaches estimate uncertainty separately from the forecasting model, which may weaken coherence between predicted demand and uncertainty information [28], [29].

This study addresses this gap by positioning probabilistic demand forecasting as a decision-oriented framework for pharmaceutical supply chain planning. The contribution is not to propose a new neural architecture, but to show how QRNN-based probabilistic forecasts can support uncertainty-aware decision-making by jointly estimating demand levels and conditional quantiles.

Specifically, lower, median, and upper quantiles are interpreted as complementary decision signals. The median quantile supports routine demand planning, while upper quantiles inform safety-stock adjustment, service-level protection, and stockout-risk mitigation. This explicitly aligns QRNN-based probabilistic forecasting with decision-making needs under uncertainty.

III. Methodology

A. Probabilistic forecasting framework

Let $y_t \in \mathbb{R}^+$ denote daily demand at time t, and $\mathbf{x}_t \in \mathbb{R}^d$ the features available before y_t. The objective is to estimate the conditional demand distribution, rather than a single expected value, since point forecasts provide limited support under uncertainty [11], [13].

We adopt a quantile-based framework that estimates selected conditional quantiles without assuming a parametric distribution. This approach is flexible, distribution-free, and suited to asymmetric operational contexts [12], [21], especially in healthcare supply chains where under-forecasting may cause stockouts, while over-forecasting may generate excess inventory, waste, and financial losses [15], [16]. The median quantile supports routine planning, while lower and upper quantiles inform inventory planning, service-level management, and risk mitigation [19], [20].

B. Data splitting and sliding-window feature construction

The analysis focuses on the daily pharmaceutical demand series of the N02BE category. Repeated observations are aggregated daily, and the task is formulated as one-step-ahead forecasting, where y_t is predicted using only information available up to $t-1$.

To avoid temporal leakage, the data are chronologically split into training, validation, and test subsets, representing 70%, 10%, and 20% of the observations after feature construction. The validation set is used for early stopping and learning-rate adjustment, while the test set is reserved for out-of-sample evaluation.

The input vector $\mathbf{x}_t$ includes causal demand lags at 1, 7, 14, 28, 56, and 365 days; rolling means, standard deviations, and medians over 7-, 14-, 28-, and 56-day windows; and calendar features including day-of-week, month, day-of-year, and weekly and yearly Fourier terms. All rolling statistics are shifted by one day to prevent leakage.

To improve numerical stability and reduce the effect of extreme peaks, the target is transformed as:

$$z_t = \log(1 + y_t). \tag{1}$$

The QRNN estimates conditional quantiles of z_t, which are back-transformed to the original scale as:

$$\widehat{Q}_q^y(\mathbf{x}_t) = \exp\left(\widehat{Q}_q^z(\mathbf{x}_t)\right) - 1. \tag{2}$$

All metrics are computed after back-transformation. Since the study focuses on one product category, the findings represent case-based validation rather than universal generalization.

C. Quantile regression neural network formulation

For $q \in (0,1)$, the QRNN estimates $\widehat{Q}_q^z(\mathbf{x}_t)$ by minimizing the pinball loss [21]:

$$L_q = \frac{1}{T}\sum_{t=1}^{T} \rho_q\left(z_t - \widehat{Q}_q^z(\mathbf{x}_t)\right), \tag{3}$$

where

$$\rho_q(u) = \begin{cases} q\,u, & u \geq 0, \\ (q-1)\,u, & u < 0. \end{cases} \tag{4}$$

This asymmetric loss is suitable when overestimation and underestimation have different operational consequences [12]. Embedded in a neural network, it enables nonlinear estimation of conditional demand distributions [22], [23].

Three quantiles are jointly estimated:

$$q \in \{0.05, 0.50, 0.95\}. \tag{5}$$

They define a central 90% prediction interval, $[P05, P95]$, around $P50$. The joint objective is:

$$L_{\text{QRNN}} = \sum_{q\in\{0.05,0.50,0.95\}} L_q. \tag{6}$$

The model uses fully connected hidden layers with nonlinear activations and dropout. The output layer contains one neuron per quantile, enabling simultaneous quantile estimation and avoiding post-hoc interval construction [11].

D. Parameter choice and expected influence

The QRNN parameters were selected to balance accuracy, calibration, and stability. The quantiles $q \in \{0.05, 0.50, 0.95\}$ provide a 90% interval around P50, where P50 supports routine planning and P05, P95 represent demand-risk bounds.

The architecture includes three hidden layers of 256, 128, and 64 neurons. Dropout rates of 0.20 and 0.10, with an L_2 penalty of 10^{-5}, limit overfitting and stabilize quantile estimates. Training uses Adam with a learning rate of 5×10^{-4}, batch size 64, and up to 400 epochs, with early stopping and learning-rate reduction based on validation loss. The selected lags and rolling windows capture recent persistence, weekly seasonality, medium-term dynamics, and annual patterns.

E. Benchmark and evaluation metrics

The QRNN is evaluated against a persistence benchmark:

$$\widehat{y}_t = y_{t-1}. \tag{7}$$

This benchmark is transparent and widely used in short-horizon forecasting [30]. Point accuracy is evaluated using P50 after back-transformation, with MSE, RMSE, MAE, R^2, and sMAPE [7]. Probabilistic performance is assessed using quantile-specific pinball losses [13] and empirical coverage of $[P05, P95]$ to evaluate calibration [11]. The persistence comparison is conservative but limited; broader baselines should be considered in future work.

F. Computational complexity

The QRNN is implemented as a feed-forward multilayer perceptron. Let T be the number of observations, d the number of features, E the number of epochs, and P the number of trainable parameters. The training cost is:

$$\mathcal{O}(ETP). \tag{8}$$

The input vector contains 29 causal features, and the network has three hidden layers with 256, 128, and 64 neurons followed by three quantile outputs, totaling approximately 49,000 trainable parameters. Since lags and rolling windows are fixed, feature construction scales linearly:

$$\mathcal{O}(Td). \tag{9}$$

At inference time, each one-step-ahead forecast requires a single forward pass:

$$\mathcal{O}(P). \tag{10}$$

Thus, the QRNN remains suitable for regular operational updates while producing probabilistic demand forecasts.

IV. Results and Discussion

A. Point-forecast performance: median quantile (P50)

Although the QRNN is designed for probabilistic forecasting, its median quantile (P50) can also serve as a robust point forecast. Table I compares the QRNN median forecast with the NAIVE persistence benchmark.

TABLE I

Test-set point-forecast performance for daily demand (N02BE).

Model	MSE	RMSE	MAE	R^2	sMAPE (%)
NAIVE (y_{t-1})	256.3330	16.0104	12.0080	0.0266	43.977
QRNN (P50)	172.3105	13.1267	9.6138	0.3456	33.757
Improvement	32.78%	18.01%	19.94%	+0.319	23.24%

The QRNN outperforms the persistence benchmark across all point-forecasting metrics. MSE decreases by 32.78%, RMSE by 18.01%, MAE by 19.94%, and sMAPE by 23.24%. The increase in R^2, from 0.0266 to 0.3456, shows that the QRNN captures demand dynamics beyond simple temporal persistence. Thus, with respect to RQ1, the model improves point-forecast accuracy while preserving its probabilistic structure.

B. Temporal behavior and visual analysis

Figure 1 shows the NAIVE benchmark. Although it may appear acceptable during stable periods, it only reproduces the previous observation and therefore reacts after demand changes occur. This explains its lag during abrupt variations and peak events.

Figure 2 reports the QRNN median forecast. Compared with persistence, the P50 forecast follows the demand trajectory more consistently and responds better to regime changes, supporting the quantitative gains reported in Table I.

C. Probabilistic performance and uncertainty calibration

Beyond point accuracy, the main contribution of the QRNN is its ability to estimate predictive uncertainty through conditional quantiles. Table II reports the quantile-specific pinball losses and the empirical coverage of the $[P05, P95]$ interval.

TABLE II

Probabilistic forecasting performance on the test set.

Pinball (q=0.05)	Pinball (q=0.50)	Pinball (q=0.95)	Coverage (%)
1.1324	4.8069	1.9132	93.39

The empirical coverage reaches 93.39%, close to the nominal 90% level. This indicates reasonably calibrated and slightly conservative uncertainty estimates, which can be acceptable in healthcare supply chains where underestimating uncertainty may increase stockout risk.

Figure 3 shows the QRNN median forecast with its uncertainty band. The interval widens during unstable periods and narrows when demand is more stable, suggesting that the model produces time-varying uncertainty estimates rather than fixed dispersion bands.

These results further answer RQ1 by showing that the QRNN improves point accuracy while providing calibrated probabilistic information.

D. Decision-oriented interpretation of quantile forecasts

The estimated quantiles can be interpreted as complementary decision signals. The median quantile (P50) supports routine replenishment and baseline inventory planning. The upper quantile (P95) provides a risk-sensitive estimate of high-demand scenarios, supporting safety-stock adjustment, service-level protection, and stockout-risk mitigation. Conversely, the lower quantile (P05) helps identify low-demand scenarios associated with overstocking, holding costs, expiration risk, and waste.

With respect to RQ2, quantile-based forecasts can therefore be aligned with distinct supply chain decisions: P50 for routine planning, P95 for service-level and stockout-risk management, and P05 for overstock-risk monitoring. The interval $[P05, P95]$ provides a structured range of plausible demand outcomes rather than a single isolated forecast.

E. Discussion of findings

Overall, the results support the relevance of QRNN-based probabilistic forecasting for daily pharmaceutical demand. The model improves point-forecast accuracy relative to the NAIVE benchmark and provides prediction intervals with coverage close to the nominal level. Point accuracy supports routine planning, while calibrated uncertainty intervals support risk-sensitive decisions.

The findings also confirm the limitations of relying only on point forecasts in healthcare supply chains, since they do not describe the range of plausible future demand values. By contrast, the QRNN estimates multiple conditional quantiles within a unified model, producing both demand levels and uncertainty ranges.

However, the results should be interpreted with caution. The analysis focuses on one product category, N02BE, and therefore provides case-based validation rather than universal evidence. In addition, the benchmark is simple and transparent, but broader comparisons with statistical, machine learning, and probabilistic baselines would strengthen the assessment. Future work should also examine interval sharpness and additional probabilistic scores to assess the trade-off between reliability and informativeness.

Despite these limitations, the study shows that QRNN-based probabilistic forecasting can support uncertainty-aware demand planning by linking conditional quantiles to inventory planning, service-level management, and risk mitigation.

V. Conclusion

This paper proposed a decision-oriented probabilistic forecasting framework based on a Quantile Regression Neural Network for daily pharmaceutical demand forecasting. Instead of producing only point estimates, the model estimates multiple conditional quantiles, providing both a median

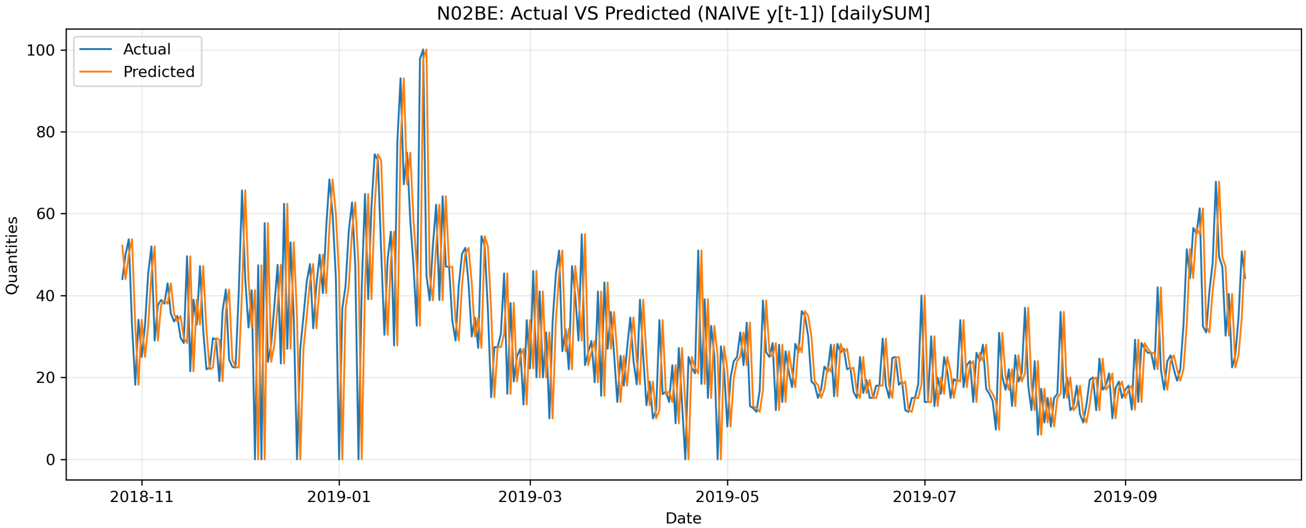

Fig. 1. N02BE daily demand: actual vs predicted using the NAIVE persistence baseline ($\widehat{y}_t = y_{t-1}$).

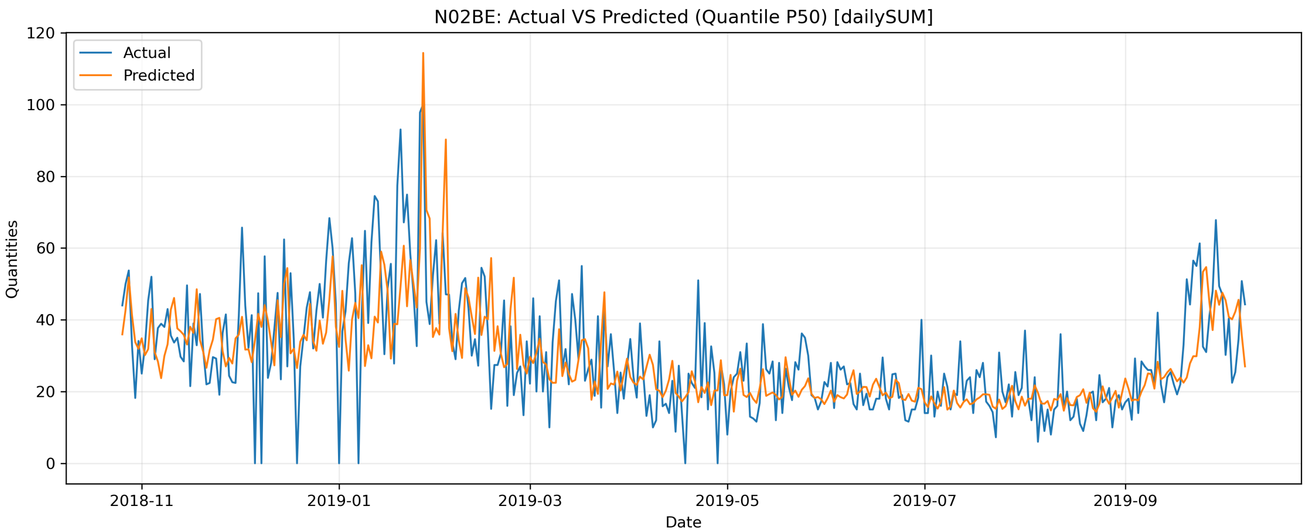

Fig. 2. N02BE daily demand: actual vs QRNN median quantile forecast (P50).

forecast and an uncertainty interval for healthcare supply chain planning.

With respect to RQ1, the QRNN improves predictive performance over the persistence benchmark while providing reasonably calibrated uncertainty estimates. The P50 forecast reduces MSE, RMSE, MAE, and sMAPE, increases R^2, and the $[P05, P95]$ interval achieves 93.39% empirical coverage, close to the nominal 90% level. These results indicate that the model captures both demand dynamics and time-varying uncertainty more effectively than simple persistence.

With respect to RQ2, the estimated quantiles can be linked to operational decisions. P50 supports routine demand planning, while P05 and P95 help manage asymmetric risks related to overstocking, stockouts, safety-stock adjustment, and service-level control. Thus, the framework positions probabilistic forecasting as a decision-support tool rather than a purely statistical extension of point forecasting.

The study remains limited by its focus on a single pharmaceutical product category and by comparison with a persistence benchmark only. Future work should extend the evaluation to multiple products, include statistical, machine learning, recurrent, and probabilistic baselines, and integrate QRNN forecasts into inventory optimization models. Overall, the results support the value of quantile-based forecasting for robust, risk-sensitive, and decision-oriented pharmaceutical supply chain planning under uncertainty.

Acknowledgment

The authors gratefully acknowledge the partial support of the RITMEA project, Research and Innovation on Ecological Responsibility and Autonomous Transport and Mobility. This project is co-financed by the European Union through the European Regional Development Fund, the French State, and the Hauts-de-France Region Council.

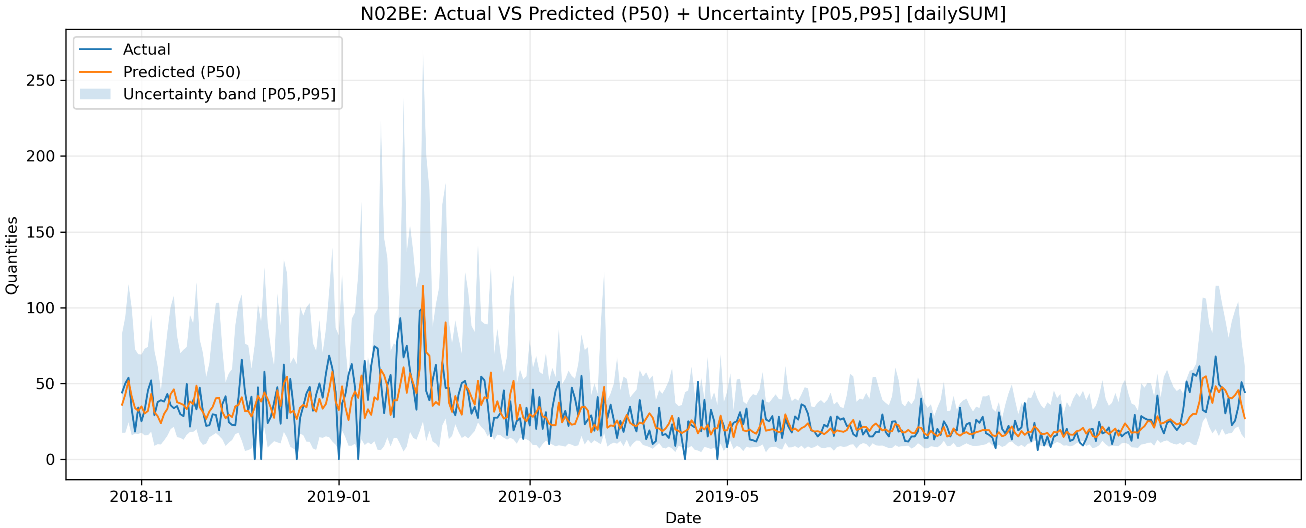

Fig. 3. N02BE daily demand: QRNN probabilistic forecast with median (P50) and uncertainty band $[P05, P95]$.

REFERENCES

[1] E. A. Silver, D. F. Pyke, R. Peterson *et al.*, *Inventory management and production planning and scheduling*. Wiley New York, 1998, vol. 3.

[2] D. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi, *Designing and managing the supply chain: Concepts, strategies, and cases*. McGraw-hill New York, 1999.

[3] S. Chopra and P. Meindl, "Strategy, planning, and operation," *Supply Chain Management*, vol. 15, no. 5, pp. 71–85, 2001.

[4] J. De Vries and R. Huijsman, "Supply chain management in health services: an overview," *Supply chain management: An international journal*, vol. 16, no. 3, pp. 159–165, 2011.

[5] J. Volland, A. Fügener, J. Schoenfelder, and J. O. Brunner, "Material logistics in hospitals: A literature review," *Omega*, vol. 69, pp. 82–101, 2017.

[6] A. A. Syntetos and J. E. Boylan, "The accuracy of intermittent demand estimates," *International Journal of forecasting*, vol. 21, no. 2, pp. 303–314, 2005.

[7] R. J. Hyndman and G. Athanasopoulos, *Forecasting: principles and practice*. OTexts, 2018.

[8] C. Chatfield and H. Xing, *The analysis of time series: an introduction with R*. Chapman and hall/CRC, 2019.

[9] F. Petropoulos, R. Fildes, and P. Goodwin, "Do 'big losses' in judgmental adjustments to statistical forecasts affect experts' behaviour?" *European Journal of Operational Research*, vol. 249, no. 3, pp. 842–852, 2016.

[10] J. S. Armstrong, *Principles of forecasting: a handbook for researchers and practitioners*. Springer Science & Business Media, 2001, vol. 30.

[11] T. Gneiting and M. Katzfuss, "Probabilistic forecasting," *Annual Review of Statistics and Its Application*, vol. 1, no. 1, pp. 125–151, 2014.

[12] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *International journal of forecasting*, vol. 22, no. 4, pp. 679–688, 2006.

[13] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *Journal of the American statistical Association*, vol. 102, no. 477, pp. 359–378, 2007.

[14] T. Hong, P. Pinson, S. Fan, H. Zareipour, A. Troccoli, and R. J. Hyndman, "Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond," pp. 896–913, 2016.

[15] P. Zipkin, "Foundations of inventory management mcgraw-hill," *Irwin, New York, USA*, 2000.

[16] L. V. Snyder, Z. Atan, P. Peng, Y. Rong, A. J. Schmitt, and B. Sinsoysal, "Or/ms models for supply chain disruptions: A review," *Iie transactions*, vol. 48, no. 2, pp. 89–109, 2016.

[17] B. Elghomri, F. Messaoudi, and N. Touti, "The role of ai in driving accountability and transparency in global supply chains: The fragmented bridge between research and practice," *Operations and Supply Chain Management An International Journal*, pp. 275–287, 2025.

[18] B. Elghomri and F. Messaoudi, "Mapping current trends in iot and industry 4.0 supply chains: A bibliometric approach," *Grenze International Journal of Engineering & Technology (GIJET)*, vol. 11, 2025.

[19] T. Boone, J. E. Boylan, R. Fildes, R. Ganeshan, and N. Sanders, "Perspectives on supply chain forecasting," pp. 121–127, 2019.

[20] M. Z. Babai, M. M. Ali, J. E. Boylan, and A. A. Syntetos, "Forecasting and inventory performance in a two-stage supply chain with arima (0, 1, 1) demand: Theory and empirical analysis," *International Journal of Production Economics*, vol. 143, no. 2, pp. 463–471, 2013.

[21] R. Koenker, *Quantile regression*. Cambridge university press, 2005, vol. 38.

[22] L. M. De Menezes, D. W. Bunn, and J. W. Taylor, "Review of guidelines for the use of combined forecasts," *European Journal of Operational Research*, vol. 120, no. 1, pp. 190–204, 2000.

[23] S. Ben Taieb and R. Hyndman, "Recursive and direct multi-step forecasting: the best of both worlds," Monash University, Department of Econometrics and Business Statistics, Tech. Rep., 2012.

[24] B. Elghomri and F. Messaoudi, "Digital twins for sustainable and low-carbon supply chains in industry x. 0: Crossing the fragmented bridge between academic research and industrial adoption," *Operations and Supply Chain Management: An International Journal*, vol. 19, no. 1, pp. 13–29, 2025.

[25] B. Elghomri, F. Messaoudi, H. Allaoui, and A. El AMRAOUI, "From agents to sustainability: Charting the evolution and impact of multi-agent systems in circular supply chain management," in *2025 International Conference on Circuit, Systems and Communication (ICCSC)*. IEEE, 2025, pp. 1–9.

[26] L. Alzubaidi, J. Bai, A. Al-Sabaawi, J. Santamaría, A. S. Albahri, B. S. N. Al-Dabbagh, M. A. Fadhel, M. Manoufali, J. Zhang, A. H. Al-Timemy *et al.*, "A survey on deep learning tools dealing with data scarcity: definitions, challenges, solutions, tips, and applications," *Journal of Big Data*, vol. 10, no. 1, p. 46, 2023.

[27] D. Koutsandreas, E. Spiliotis, F. Petropoulos, and V. Assimakopoulos, "On the selection of forecasting accuracy measures," *Journal of the Operational Research Society*, vol. 73, no. 5, pp. 937–954, 2022.

[28] J. Gasthaus, K. Benidis, Y. Wang, S. S. Rangapuram, D. Salinas, V. Flunkert, and T. Januschowski, "Probabilistic forecasting with spline quantile function rnns," in *The 22nd international conference on artificial intelligence and statistics*. PMLR, 2019, pp. 1901–1910.

[29] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, "Deepar: Probabilistic forecasting with autoregressive recurrent networks," *International journal of forecasting*, vol. 36, no. 3, pp. 1181–1191, 2020.

[30] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "Statistical and machine learning forecasting methods: Concerns and ways forward," *PloS one*, vol. 13, no. 3, p. e0194889, 2018.