mathematics

Article

Quantifying the Uncertainty of Reservoir Computing: Confidence Intervals for Time-Series Forecasting

Laia Domingo [1,2], Mar Grande [1,3], Florentino Borondo [4] and Javier Borondo [3,5,*]

1 Grupo de Sistemas Complejos, Universidad Politécnica de Madrid, 28040 Madrid, Spain; laia@ingenii.dev (L.D.); mar.grande@alumnos.upm.es (M.G.)
2 Ingenii Inc., New York, NY 10013, USA
3 AGrowingData, 04001 Almería, Spain
4 Departamento de Química, Universidad Autónoma de Madrid, Cantoblanco, 28049 Madrid, Spain; f.borondo@uam.es
5 ICAI Engineering School, Universidad Pontificia de Comillas, Alberto Aguilera 23, 28015 Madrid, Spain
* Corresponding: jborondo@comillas.edu

Abstract: Recently, reservoir computing (RC) has emerged as one of the most effective algorithms to model and forecast volatile and chaotic time series. In this paper, we aim to contribute to the understanding of the uncertainty associated with the predictions made by RC models and to propose a methodology to generate RC prediction intervals. As an illustration, we analyze the error distribution for the RC model when predicting the price time series of several agri-commodities. Results show that the error distributions are best modeled using a Normal Inverse Gaussian (NIG). In fact, NIG outperforms the Gaussian distribution, as the latter tends to overestimate the width of the confidence intervals. Hence, we propose a methodology where, in the first step, the RC generates a forecast for the time series and, in the second step, the confidence intervals are generated by combining the prediction and the fitted NIG distribution of the RC forecasting errors. Thus, by providing confidence intervals rather than single-point estimates, our approach offers a more comprehensive understanding of forecast uncertainty, enabling better risk assessment and more informed decision-making in business planning based on forecasted prices.

Keywords: reservoir computing; uncertainty; confidence intervals; time series; market; prices

MSC: 37N40; 91B84; 03H10; 68T07

check for updates

Citation: Domingo, L.; Grande, M.; Borondo, F.; Borondo, J. Quantifying the Uncertainty of Reservoir Computing: Confidence Intervals for Time-Series Forecasting. *Mathematics* **2024**, *12*, 3078. https://doi.org/10.3390/math12193078

Academic Editor: Miquel Montero

Received: 28 August 2024
Revised: 25 September 2024
Accepted: 25 September 2024
Published: 1 October 2024

1. Introduction

Quantifying the uncertainties associated with the deterministic forecasting of financial time series is essential to enhance the reliability of many business decisions, as it provides deeper insight into the performance of machine learning models [1–4]. Deterministic forecasting provides a single-point estimate of future prices, lacking information about the uncertainty associated with the forecasts. While deterministic forecasting can be useful for providing a simple estimate of future values, it does not capture the inherent variability and uncertainty present in real-world situations. It is limited in its ability to account for unforeseen events, fluctuations in market conditions, or other factors that may impact the accuracy of the forecast.

In contrast, probabilistic forecasting approaches [5,6], such as confidence interval prediction [5,7], explicitly consider uncertainty by providing a range within which a future value is likely to fall, based on a given level of confidence. Such methods give the predicted lower and upper limits at a particular nominal confidence level. High-quality prediction intervals are defined as those that cover a specified proportion of the observations while still being as narrow as possible. Indeed, under the same level of confidence, narrower prediction intervals can provide more accurate information with fewer uncertainties.

Regarding financial markets, confidence intervals provide a measure of uncertainty and variability in the predictions, allowing for a more comprehensive understanding of the volatility of the market and the performance of the forecasting model. This information is valuable for making informed decisions, assessing risks, and planning appropriate strategies [8–10]. Simply relying on single-point predictions is not enough to make accurate business decisions based on forecasted prices. Indeed, it is essential to know the range within which future prices are likely to fall. To illustrate this phenomenon, imagine a trader who trades in the stock market and receives a trading alert indicating that the price of a certain asset is expected to increase by $10. If there is no extra information regarding the uncertainty in that alert, the trader will probably buy the asset. However, if the confidence intervals of the prediction underlying that alert are of [$−15, $25], the trader would be aware that the asset price can increase or decrease and pass up the trading alert. As this example highlights, the additional information provided by confidence intervals empowers us to design more informed strategies.

In this paper, we characterize the uncertainty of reservoir computing (RC) [11,12] in the agri-commodities market and propose a methodology to generate confidence intervals for time-series forecasting using RC. As we have already argued, predictions with confidence intervals are very important when forecasting markets. However, despite the existence of some works limited to the energy industry [13,14], there is still no standard methodology to generate time-series confidence intervals for RC models. The RC algorithm was originally proposed by Jaeger [11] to compute the time evolution of dynamical systems. It uses a fixed randomly initialized network—the reservoir—to generate many high-dimensional complex representations of the input data, which are then used to train a linear readout layer. The readout layer learns to read the relevant information extracted by the reservoir and transform it to generate accurate forecasts. One of the main advantages of RC over neural networks is that the recurrent weights of the reservoir are randomly generated and remain fixed during the training process, which significantly simplifies the training process and reduces the risk of overfitting. Hence, RC stands out as one of the most effective machine learning algorithms for time-series forecasting [15–18], especially in highly volatile scenarios [19,20]. Accordingly, RC has proven to be an optimal solution [21,22] to model the complex behavior of financial markets [22], as it is capable of anticipating the future direction of the markets [22].

The remainder of the paper is organized as follows. Section 2 is devoted to the description of the dataset used in this work. In Section 3, we explain the methodology followed to characterize the errors in the RC and build the confidence intervals. Next, in Section 4, we present our results. Finally, in Section 5, we present our conclusions and discuss the importance and limitations of the results.

2. Dataset

The dataset of this work consists of the weekly prices of three different agri-commodities: zucchini, aubergine, and tomato. Originally, the prices have a daily resolution and refer to the price per kilogram. However, for the present paper, we have aggregated the prices of each product weekly by averaging the prices of the week. We chose this approach because weekly predictions are more representative and therefore more valuable than daily ones for the agri-food market. The period covered by our dataset covers from March 2013 to March 2022 and has been temporally split into training, validation, and test sets. The training set contains the prices until December 2019, the validation set contains the prices in 2020, and the test set contains the prices from January 2021 onwards. Additionally, the time series have been standarized to the $[-1, 1]$ domain using a linear scale. All prices are given in euros per kilogram throughout the paper. In addition to the time series, the RC models employed in this work also incorporate 16 regressor variables to predict the evolution of prices over time. These variables include data on production volumes and international trade, offering complementary insights for price prediction.

3. Methods

In this paper, we present a method based on modeling the distribution of forecasting errors made by the RC to generate its confidence intervals. We will illustrate the method by applying it to the decomposition-RC (D-RC) that was presented in [22]. The D-RC approach first decomposes the time series into the trend, seasonal, and residuals components and then models each component with a separate ensemble RC. We use an ensemble RC instead of a classic RC since the final configuration of RC models may be conditioned by the initial configuration of the reservoir. Thus, by building an ensemble, we are able to mitigate this effect and achieve more robust results. As shown in [22], the D-RC is the most adequate RC variant to forecast volatile time series. More specifically, the D-RC, which predicts each component in the time series separately, is significantly more effective in preventing the model from mimicking the previous value of the series, which is a common failure when modeling financial time series. The hyperparameters used to train the D-RC models are shown in Table 1.

Table 1. Optimal hyperparameters used to train the different reservoir computing models: α is the leaking rate, N the number of neurons of the reservoir, γ is the L^2 regularization parameter for the ridge regression, $\rho(W)$ and $D(W)$ are the spectral radius and density of the reservoir matrix W, and N_e is the number of reservoirs used in the ensemble.

Model	α	N	γ	$\rho(W)$	$D(W)$	N_e
D-RC (trend)	0.58	60	0.011	1.1	0.016	80
D-RC (seasonality)	0.74	20	0.011	0.32	0.023	100
D-RC (residuals)	0.91	60	7.24	0.85	0.022	120

The method begins by studying the error $\epsilon(t)$ of the D-RC predictions in the training set, defined as $\epsilon(t) = y_{\text{teach}}(t) - y(t)$, where $y_{\text{teach}}(t)$ is the actual price at time t and $y(t)$ is the prediction of the D-RC model. The agricultural prices are very volatile, presenting abrupt changes of tendency. As a result, we expect the error distribution to exhibit different behavior depending on whether or not there is a significant change of tendency in the predictions. For this reason, the training data have been separated into three sets according to the relative increase in the prices. The first set contains the data with a relative increase higher than +10% (*Increase*). The second class contains the data with a relative increase lower than −10% (*Decrease*). Finally, the third set contains the data with a relative increase between −10% and +10% (*Constant*):

$$\begin{aligned} &Increase \rightarrow \quad \frac{y_{\text{teach}}(t-1) - y(t)}{y(t)} \geq 0.1, \\ &Decrease \rightarrow \quad \frac{y_{\text{teach}}(t-1) - y(t)}{y(t)} \leq -0.1, \\ &Constant \rightarrow \quad -0.1 < \frac{y_{\text{teach}}(t-1) - y(t)}{y(t)} < 0.1. \end{aligned} \tag{1}$$

Notice that the classification in Equation (1) only depends on the *predicted* prices at time t, $y(t)$, and not on the true prices at time t, $y_{\text{teach}}(t)$. In other words, the *Increase* set contains data points where the RC model predicted a price increase compared to the previous true price $y_{\text{teach}}(t-1)$. Similarly, the *Decrease* set comprises points where the RC model predicted a price decrease relative to the previous known price.

Once the training data have been divided into these three sets, the error distribution $\epsilon(t) = y(t) - y_{\text{teach}}(t)$ is studied. Then, we find the best fit for the probability distribution. In this case, all distributions are modeled using a Normal Inverse Gaussian (NIG) distribution [23], whose probability density function is given by

$$f(x,a,b) = \frac{aK_1(a\sqrt{1+y^2})}{\pi\sqrt{1+y^2}} e^{\sqrt{a^2-b^2}+by}, \quad y = \frac{x - x_0}{\sigma}, \tag{2}$$

where x is a real number, the parameter a is the tail heaviness, and b is the asymmetry parameter satisfying $a > 0$ and $|b| \leq a$, x_0 and σ are the location and scale factors, and K_1 is the modified Bessel function of second kind [24]. Then, the resulting probability distribution is used to calculate the lower and upper bounds of the prediction interval at a given prediction interval coverage (PINC):

$$U_t = y(t) + \mathrm{NIG}^{a,b}_{x_0,\sigma}(1-\alpha/2)(\frac{\mathrm{PINC}}{2}),$$
$$L_t = y(t) + \mathrm{NIG}^{a,b}_{x_0,\sigma}(1-\alpha/2)(1-\frac{\mathrm{PIC}}{2}) \tag{3}$$

This process is illustrated in Figure 1. Notice that, since each of the three training sets (*Constant*, *Increase*, *Decrease*) is modeled with a different distribution, the width of the interval varies depending on the specific set to which a given prediction belongs. After calculating the confidence intervals $[L_t, U_t]$ for the validation/testing steps $t = 1, \cdots, T$, two metrics are used to evaluate the quality of such intervals. The first one is the Average Coverage Error (ACE), which measures how much the predicted intervals can be trusted:

$$\mathrm{ACE} = \mathrm{PICP} - \mathrm{PINC}, \quad \mathrm{PICP} = \frac{1}{T}\sum_{t=1}^{T} c_t, \tag{4}$$
$$c_t = \begin{cases} 1 & \text{if } y_{\text{teach}}(t) \in [L_t, U_t] \\ 0 & \text{otherwise,} \end{cases}$$

where PICP is the Prediction Interval Coverage Probability. The second metric is the Prediction Interval Average Width (PIAW), which measures the average width of the predicted intervals:

$$\mathrm{PIAW} = \frac{1}{T}\sum_{t=1}^{T}(U_t - L_t). \tag{5}$$

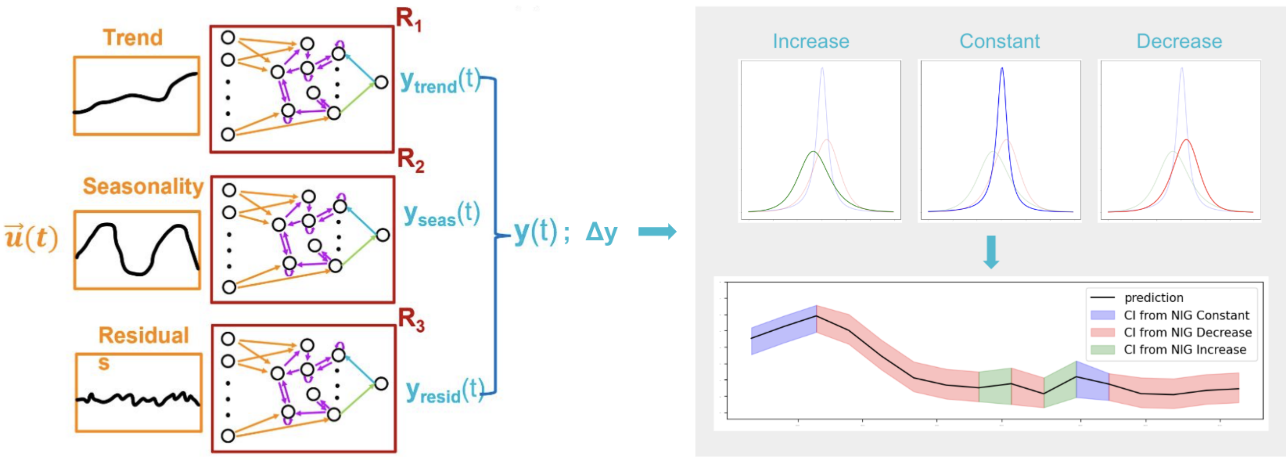

Figure 1. Methodology followed to generate the confidence intervals for the D-RC predictions. First, the RC generates a forecast for the time series. Then the forecast is compared to the last observed value of the series to obtain the predicted trend. Next, the distribution from which to generate the confidence intervals according to the predicted trend is chosen. Finally, the lower and upper confidence intervals for the prediction are computed.

The predicted intervals should be as narrow as possible, meaning that they should minimize the PIAW, while following the condition ACE ≥ 0. In the following sections we present and evaluate the predicted intervals for the three time series in the study: zucchini, aubergine, and tomato.

4. Results

4.1. Zucchini

The relative error distributions for the training set are shown in Figure 2A. As can be seen, the relative error is significantly smaller for the set containing the *Constant* distribution (see Equation (1)). Moreover, the *Increase* distribution has positive skewness, while the *Decrease* distribution has negative skewness. This skewness agrees with the definition of the three sets and means that the predictions tend to underestimate big changes in the tendency in the data. The distributions in Figure 2A are modeled using an NIG with parameters $a = 0.12$, $b = -0.048$, $x_0 = 0.0069$, $\sigma = 0.051$ for the *Constant* set, $a = 1.36$, $b = 0.22$, $x_0 = -0.083$, $\sigma = 0.20$ for the *Increasing* set, and $a = 1.66$, $b = -0.25$, $x_0 = 0.056$, $\sigma = 0.18$ for the *Decreasing* set. The NIG distribution is chosen among other distributions since it minimizes the BIC (Bayesian Information Criterion) indicator and the squared error, which are given in Table 2. The BIC indicator measures the quality of a statistical model while taking into account the balance between the goodness of fit and model complexity [25]. Figure 2A shows the shape of the fitted distribution together with the best fit of the Gaussian distribution. As can be seen, the difference of both distributions is specially relevant for the *Constant* set, where the Gaussian distribution does not capture the peak around $\epsilon = 0$.

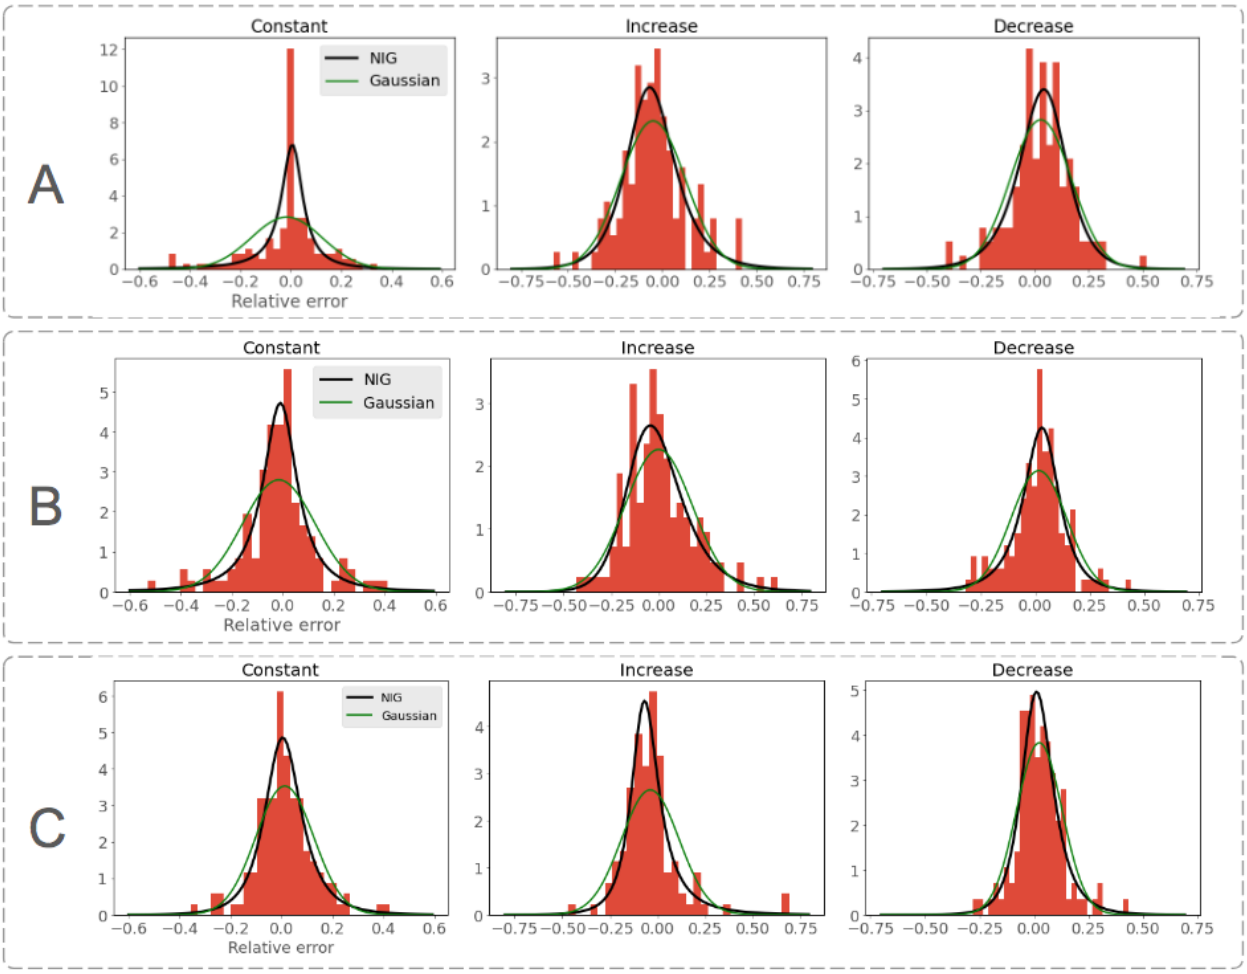

Figure 2. Distribution of the relative error in the training set for the *Constant* set (**left panels**), the *Increase* set (**middle panels**), and the *Decrease* set (**right panels**). (**A**) shows the results for zucchini, (**B**) for aubergine, and (**C**) for tomato. See Equation (1) for the definition of the three sets.

Table 2. Squared error and BIC indicators for the Normal Inverse Gaussian (NIG) distribution for the three time series and relative increase sets, compared with the Gaussian distribution.

Time Series	Set	Distribution	Squared Error	BIC
Zucchini	*Constant*	NIG Gaussian	94 109	276 286
	Increase	NIG Gaussian	88 98	−8.3 −7.0
	Decrease	NIG Gaussian	74 78	−44 −42
Aubergine	*Constant*	NIG Gaussian	80 134	−23 18
	Increase	NIG Gaussian	74 80	−40 −40
	Decrease	NIG Gaussian	134 159	23 35
Tomato	*Constant*	NIG Gaussian	98 122	−18 1.7
	Increase	NIG Gaussian	55 83	−65 −28
	Decrease	NIG Gaussian	141 155	37 39

The fitted distributions have been used to construct the predicted confidence intervals for four different values of PINC = 90%, 80%, 70%, and 60%. The intervals for the validation and test set are displayed in Figure 3, where it can be observed that most of the target values are enclosed by these constructed prediction intervals. The PICP, ACE, and PIAW for both the NIG and Gaussian distributions are given in Table 3. As can be seen, the condition of $ACE \gtrapprox 0$ is almost always accomplished, except for PINC = 60%, where the ACE for the NIG distribution is slightly smaller than 0. The Gaussian distribution tends to have higher values of ACE but also higher values of PIAW, which means that such distribution is overestimating the size of the confidence intervals. Thus, the predicted Gaussian intervals are less informative than the predicted NIG intervals, verifying the superiority of the NIG distribution over the Gaussian distribution.

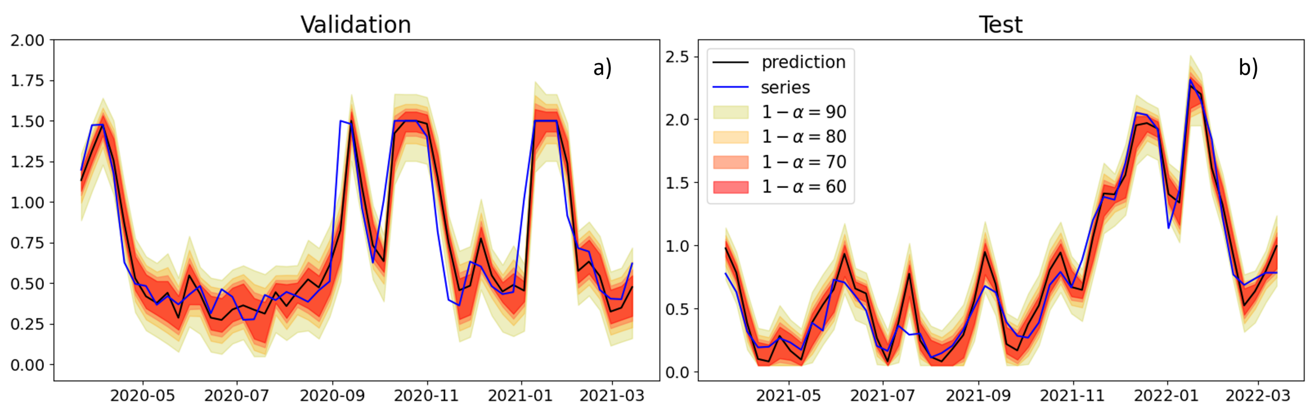

Figure 3. Predicted confidence interval for the zucchini time series for the validation (**a**) and test (**b**) sets.

Table 3. Prediction Interval Nominal Coverage (PINC), Prediction Interval Coverage Probability (PICP), Average Coverage Error (ACE), and Prediction Interval Average Width (PIAW) for the interval prediction for the validation and test sets for the zucchini time series, compared with the Gaussian distribution.

Zucchini							
		Validation			Test		
PINC	**Distribution**	**PICP**	**ACE**	**PIAW**	**PICP**	**ACE**	**PIAW**
90%	NIG	0.865	−0.035	0.469	0.942	0.042	0.458
	Gaussian	0.895	−0.005	0.489	0.942	0.042	0.475
80%	NIG	0.846	0.046	0.324	0.846	0.046	0.328
	Gaussian	0.865	0.065	0.383	0.904	0.104	0.376
70%	NIG	0.731	0.031	0.246	0.712	0.012	0.254
	Gaussian	0.769	0.069	0.310	0.846	0.146	0.308
60%	NIG	0.577	−0.023	0.191	0.577	−0.023	0.201
	Gaussian	0.731	0.131	0.251	0.635	0.035	0.252

4.2. Aubergine

The previous analysis is repeated for the aubergine time series. The relative error distributions are shown in Figure 2B. In this case, the distributions are modeled using an NIG with parameters $a = 0.31$, $b = -0.025$, $x_0 = -0.0097$, $\sigma = 0.084$ for the *Constant* set, $a = 2.68$, $b = 1.05$, $x_0 = -0.11$, $\sigma = 0.257$ for the *Increasing* set, and $a = 0.77$, $b = -0.14$, $x_0 = 0.037$, $\sigma = 0.11$ for the *Decreasing* set. The BIC and squared error for these distributions, compared with the normal distributions, are given in Table 2. In this case, the distribution for the *Constant* set also has smaller variance. The *Increase* distribution has positive skewness and the *Decrease* distribution has negative skewness, which are captured by the NIG distribution but not by the Gaussian distribution. The resulting predicted intervals are shown in Figure 4. Again, most of the predictions lie within the confidence intervals. The metrics to evaluate such intervals are given in Table 4. The results are qualitatively equivalent to those from the zucchini time series. That is, the predicted Gaussian intervals are less informative than the predicted NIG intervals because the Gaussian distribution produces confidence intervals that are too wide, while the NIG distribution is more suitable for estimating the relative error.

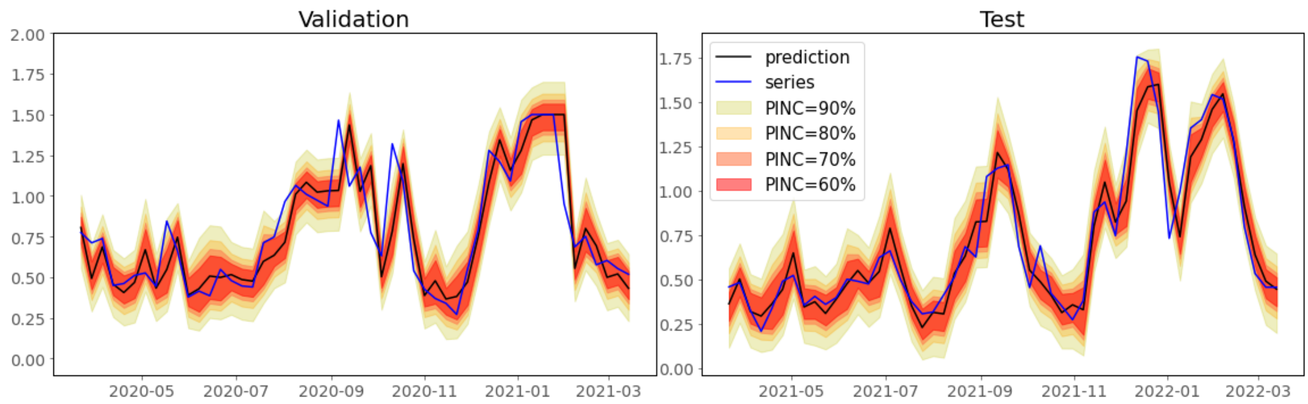

Figure 4. Predicted confidence interval for the aubergine time series for the validation and test sets.

Table 4. Same as Table 3 for the aubergine time series.

		Aubergine					
		Validation			Test		
PINC	Distribution	PICP	ACE	PIAW	PICP	ACE	PIAW
90%	NIG	0.865	−0.035	0.470	0.942	0.042	0.475
	Gaussian	0.846	−0.054	0.485	0.942	0.042	0.488
80%	NIG	0.807	0.007	0.327	0.866	0.065	0.335
	Gaussian	0.826	0.027	0.378	0.9865	0.065	0.380
70%	NIG	0.692	−0.008	0.249	0.692	−0.008	0.256
	Gaussian	0.788	0.088	0.305	0.846	0.146	0.308
60%	NIG	0.615	0.015	0.194	0.596	−0.004	0.200
	Gaussian	0.711	0.111	0.248	0.692	0.0092	0.202

4.3. Tomato

Finally, the analysis is repeated for the tomato time series. The relative error distributions are shown in Figure 2C. In this case, the distributions are modeled using an NIG with parameters $a = 0.699$, $b = 0.074$, $x_0 = 0.0013$, $\sigma = 0.094$ for the *Constant* set, $a = 0.508$, $b = 0.192$, $x_0 = -0.077$, $\sigma = 0.093$ for the *Increasing* set, and $a = 1.174$, $b = 0.229$, $x_0 = 0.0022$, $\sigma = 0.11$ for the *Decreasing* set. The BIC and squared errors for this distributions, compared with the normal distributions, are given in Table 2. In this case, the distribution for the *Constant* set also has smaller variance, but this effect is less pronounced when compared with the other two time series. The predicted intervals resulting from this analysis are shown in Figure 5. Most of the predictions lie within the confidence intervals. The metrics to evaluate such intervals are given in Table 5. Again, the results are qualitatively equivalent to those from the zucchini and aubergine time series, confirming the hypothesis of the NIG distribution being more suitable than the Gaussian distribution in modeling the error of the D-RC model.

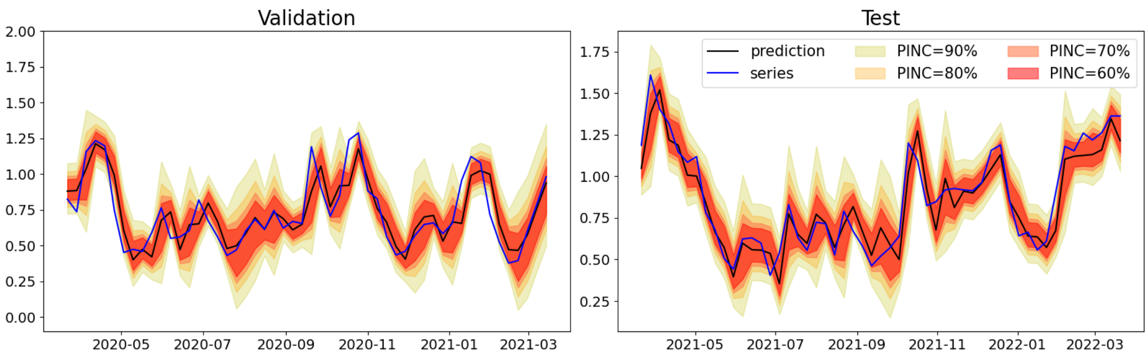

Figure 5. Predicted confidence interval for the tomato time series for the validation and test set.

Table 5. Same as Table 3 for the tomato time series.

		Tomato					
		Validation			Test		
PINC	Distribution	PICP	ACE	PIAW	PICP	ACE	PIAW
90%	NIG	0.904	0.0038	0.537	0.962	0.062	0.492
	Gaussian	0.885	−0.015	0.512	1.0	0.1	0.481
80%	NIG	0.769	−0.0307	0.374	0.905	0.105	0.346
	Gaussian	0.846	0.046	0.399	0.925	0.124	0.374
70%	NIG	0.673	-0.027	0.284	0.698	−0.002	0.263
	Gaussian	0.788	0.088	0.323	0.792	0.092	0.303
60%	NIG	0.596	−0.004	0.220	0.622	0.023	0.205
	Gaussian	0.673	0.073	0.262	0.755	0.155	0.246

4.4. Confidence Intervals Based on the Predicted Trend Outperform Generic Confidence Intervals

Finally, we compare the performance of the confidence intervals proposed in this paper—which depend on the predicted trend—with the more classic approach, in which the confidence intervals are independent from the predicted trend. We begin by building the relative error distribution for the full training sample for the three agri-commodities. Again, the distribution is best modeled using an NIG, as it minimizes the BIC and squared error. For example, in this case, the parameters of the NIG distribution for zucchini are $a = 0.9$, $b = -0.137$, $x_0 = -0.0109$, $\sigma = 0.147$. In Figure 6A–C, we compare the fitted distributions for the full training samples of zucchini, aubergine, and tomato with the fitted distributions obtained at the beginning of the section for each trend (increase, decrease, constant) separately. An important difference between them is that the fitted distributions for the constant case are significantly narrower than the general one. This fact means that, when the model predicts variations of less than 10% in absolute value, the confidence intervals will be significantly smaller for the trend-dependent version.

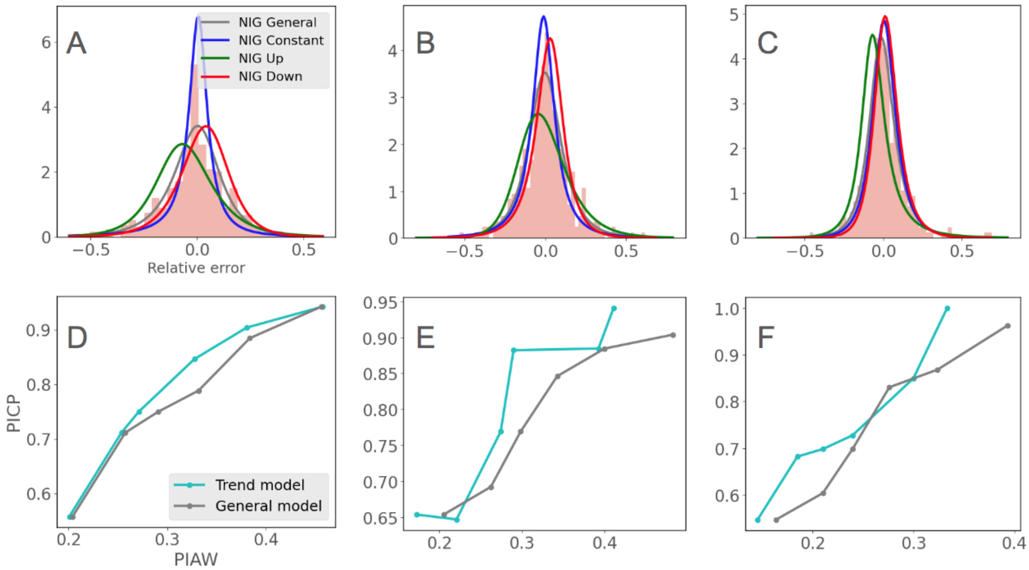

Figure 6. The top row compares the NIG fitted distribution of the relative error for the full training sample (General), constant predictions, increase predictions, and decrease predictions for zucchini (**A**), aubergine (**B**), and tomato (**C**). The bottom row shows the relation between the PICP and PIAW of the predicted confidence intervals in the test set for PINC levels of 0.6, 0.7, 0.75, 0.85, 0.9 for zucchini (**D**), aubergine (**E**), and tomato (**F**).

Next, we use the fitted general distribution to generate confidence intervals for the predictions in the test sets and compute the PICP, ACE, and PIAW. In Figure 6D–F, we illustrate the relation between PICP and PIAW for the two cases: generic confidence intervals and trend-dependent confidence intervals for the three time series. As already explained, we ideally aim to generate confidence intervals that are as narrow as possible (low PIAW) that, at the same time, cover as many times as possible the real value (high PICP). As the figure shows, with the trend-dependent confidence intervals, we are able to capture inside the intervals a higher percentage of observations with narrower intervals.

5. Conclusions

The RC method has recently emerged as one of the most promising frameworks to model volatile or chaotic time series, especially when the time series are conditioned by external factors [22]. Despite several studies having taken advantage of RC to model a wide variety of time series, we still lack a standard methodology to generate confidence intervals for RC models. In actuality, confidence intervals are critical in order to make informed decisions from predictive models. Confidence intervals empower us to make better decisions by taking into account the uncertainty of the predictions. In this paper, we have presented a method for estimating confidence intervals for RC by examining the statistical distribution of prediction errors. Our method consists of first generating predictions with an RC model and then estimating the confidence intervals from a distribution that depends on the predicted trend.

We have analyzed the errors that result from the RC model when predicting the price time series of three agri-commodities. More specifically, we have analyzed the overall errors of the model and the errors separated according to the predicted trend of the series (whether the series is expected to increase, decrease, or stay constant). The analysis reveals that the NIG distribution is the most appropriate for modeling the error distribution for all the studied cases. In fact, the NIG distribution represents a superior choice over the Gaussian distribution, as the latter tends to overestimate the width of the confidence intervals. In addition, our results support the choice of separating the errors according to the predicted trend of the series, as the three fitted distributions are significantly different. Finally, we show that, by generating the confidence intervals from specific distributions that depend on the predicted trend, we are able to generate narrower intervals without reducing the percentage of observations that fall inside them. In other words, this strategy enables us to achieve a better trade-off between PIAW and PICP.

To conclude, we discuss some limitations of our approach and some future lines of development. One of the limitations of our approach is that the confidence intervals only depend on the predicted value of the time series and not on regressor variables. In some scenarios, external variables are not necessary to model the time series; however, in other situations, they can have a big impact on the performance of RC. Thus, in the future, we aim to further develop our approach so that it will not only account for the predicted trend of the time series, but also for the value of additional regressor variables that influence the future values of the time series. We believe that including regressor variables will further improve our results by providing narrower intervals in scenarios where all the external variables point in the same direction. Additionally, we also plan to explore how the errors of the model evolve over longer time horizons. The insights derived from such analyses will be relevant, as long-term changes in time series can impact the distribution of errors. Finally, we plan to extend this framework to multivariate price forecasting, by considering the influence of multiple products in the forecast of individual prices. Such extensions hold great potential for delivering more precise and insightful forecasts, thereby enabling improved decision-making processes in the agricultural and food industries.

Author Contributions: Conceptualization, L.D., F.B. and J.B.; Validation, L.D., M.G. and J.B.; Formal analysis, L.D. and J.B.; Investigation, L.D., J.B. and F.B.; Data curation, M.G.; Writing—original draft, L.D. and J.B.; Writing—review & editing, L.D., M.G., F.B. and J.B.; Visualization, L.D., M.G. and J.B.; Supervision, F.B. and J.B. All authors have read and agreed to the published version of the manuscript.

Funding: The project that gave rise to these results received the support of a fellowship from "la Caixa" Foundation (ID 100010434). The fellowship code is LCF/BQ/DR20/11790028. This work has also been partially supported by the Spanish Ministry of Science, Innovation and Universities, Gobierno de España, under Contract No. PID2021-122711NB-C21 and by DG of Research and Technological Innovation of the Community of Madrid (Spain) under Contract No. IND2022/TIC-23716.

Data Availability Statement: The raw data used for this study are publicly available from https://www.agroprecios.com/es/precios-subasta/ (accessed on 25 September 2022).

Conflicts of Interest: Author Laia Domingo was employed by the company Ingenii Inc., Authors Mar Grande and Javier Borondo were employed by the company AGrowingData. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Heskes, T. Practical confidence and prediction intervals. *Adv. Neural Inf. Process. Syst.* **1996**, *9*, 177–182 .
2. Allen, M.R.; Stott, P.A.; Mitchell, J.F.; Schnur, R.; Delworth, T.L. Quantifying the uncertainty in forecasts of anthropogenic climate change. *Nature* **2000**, *407*, 617–620. [CrossRef] [PubMed]
3. Chatfield, C. Model uncertainty and forecast accuracy. *J. Forecast.* **1996**, *15*, 495–508. [CrossRef]
4. Palmer, T.N. Predicting uncertainty in forecasts of weather and climate. *Rep. Prog. Phys.* **2000**, *63*, 71. [CrossRef]
5. Gneiting, T.; Katzfuss, M. Probabilistic forecasting. *Annu. Rev. Stat. Its Appl.* **2014**, *1*, 125–151. [CrossRef]
6. Gneiting, T. Editorial: Probabilistic forecasting. *J. R. Statist. Soc. Ser. A* **2008,** *171*, 319–321. [CrossRef]
7. Chryssolouris, G.; Lee, M.; Ramsey, A. Confidence interval prediction for neural network models. *IEEE Trans. Neural Netw.* **1996**, *7*, 229–232. [CrossRef]
8. Poole, C. Beyond the confidence interval. *Am. J. Public Health* **1987**, *77*, 195–199. [CrossRef]
9. Di Stefano, J. A confidence interval approach to data analysis. *For. Ecol. Manag.* **2004**, *187*, 173–183. [CrossRef]
10. Zhao, C.; Wan, C.; Song, Y. Cost-oriented prediction intervals: On bridging the gap between forecasting and decision. *IEEE Trans. Power Syst.* **2021**, *37*, 3048–3062. [CrossRef]
11. Jaeger, H. The "echo state" approach to analysing and training recurrent neural networks-with an erratum note. *Ger. Natl. Res. Cent. Inf. Technol. GMD Tech. Rep.* **2001**, *148*, 13.
12. Jaeger, H. Echo state network. *Scholarpedia* **2007**, *2*, 2330. [CrossRef]
13. Gao, L.D.; Li, Z.H.; Wu, M.Y.; Fan, Q.L.; Xu, L.; Zhang, Z.M.; Zhang, Y.P.; Liu, Y.Y. Interval reservoir computing: Theory and case studies. *Front. Energy Res.* **2024**, *11*, 1239973. [CrossRef]
14. Guerra, M.; Scardapane, S.; Bianchi, F.M. Probabilistic load forecasting with Reservoir Computing. *IEEE Access* **2023**, *11*, 145989–146002. [CrossRef]
15. Gauthier, D.; Bollt, E.; Griffith, A.; Barbosa, W. Next Generation Reservoir Computing. *Nat. Commun.* **2021**, *12*, 5564. [CrossRef]
16. Chen, P.; Liu, R.; Aihara, K.; Chen, L. Autoreservoir computing for multistep ahead prediction based on the spatiotemporal information transformation. *Nat. Commun.* **2020**, *11*, 4568. [CrossRef]
17. Jaurigue, L.; Lüdge, K. Connecting reservoir computing with statistical forecasting and deep neural networks. *Nat. Commun.* **2022**, *13*, 227. [CrossRef]
18. Gao, R.; Du, L.; Duru, O.; Yuen, K.F. Time series forecasting based on echo state network and empirical wavelet transformation. *Appl. Soft Comput.* **2021**, *102*, 107111. [CrossRef]
19. Pathak, J.; Lu, Z.; Hunt, B.; Girvan, M.; Ott, E. Using Machine Learning to Replicate Chaotic Attractors and Calculate Lyapunov Exponents from Data. *Chaos* **2017**, *27*, 121102. [CrossRef]
20. Pathak, J.; Hunt, B.; Girvan, M.; Lu, Z.; Ott, E. Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach. *Phys. Rev. Lett.* **2018**, *120*, 024102. [CrossRef]
21. Wang, W.J.; Tang, Y.; Xiong, J.; Zhang, Y.C. Stock market index prediction based on reservoir computing models. *Expert Syst. Appl.* **2021**, *178*, 115022. [CrossRef]
22. Domingo, L.; Grande, M.; Borondo, F.; Borondo, J. Anticipating food price crises by reservoir computing. *Chaos Solitons Fractals* **2023**, *174*, 113854. [CrossRef]
23. Barndorff-Nielsen, O. Hyperbolic Distributions and Distributions on Hyperbolae. *Scand. J. Stat.* **1978**, *5*, 151–157.
24. Abramowitz, M.; Stegun, I.A. *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*; Government Printing Office: Washington, DC, USA, 1965; Volume 55, pp. 374–396.
25. Chakrabarti, A.; Ghosh, J.K. AIC, BIC and Recent Advances in Model Selection. In *Philosophy of Statistics*; Bandyopadhyay, P.S., Forster, M.R., Eds.; Handbook of the Philosophy of Science; North-Holland: Amsterdam, The Netherlands, 2011; Volume 7, pp. 583–605. [CrossRef]