







```
3 require File.expand_path("../config/environment", __FILE__)
4 # Prevent database truncation if the environment is production
5 abort("The Rails environment is running in production mode!")
6 require 'spec_helper'
7 require 'rspec/rails'
8
9 require 'capybara/rspec'
10 require 'capybara/rails'
11
12 Capybara.javascript_driver = :webkit
13 Category.delete_all; Category.create
14 Shoulda::Matchers.configure do |config|
15   config.integrate do |with|
16     with.test_framework :rspec
17     with.library :rails
18   end
19 end
20
21 # Add additional requires below this line. This will be added to the
22 # Requires supporting ruby files with support/ directory.
23 # spec/support/ and its subdirectories. These files will be loaded
24 # run as spec files by default. This will be added to the
25 # in _spec.rb will both be required and run as spec files.
26 # run twice. It is recommended that you do not use the same file
```



MODELING MUNICIPAL FISHERIES PRODUCTION IN THE PHILIPPINES: A COMPARATIVE TIME SERIES ANALYSIS OF TRADITIONAL, MACHINE LEARNING, AND DEEP LEARNING APPROACHES

AN ENTRY TO THE 2026 PSAI ANNUAL CONFERENCE BEST STUDENT PAPER COMPETITION (BSPC)

Presented by: Fernandez, S.R., Andres, M.A., Tribiana, K

University of the Philippines Visayas

RESEARCH OBJECTIVES



1

Describe trend, seasonality, and structure of both series

2

Build univariate models (statistical, ML, DL)

3

Build multivariate models using commercial production as an external regressor

4

Identify the best-performing models using out-of-sample accuracy

5

Forecast production for the next eight quarters

DATA SOURCE & DESCRIPTION

Source

OpenSTAT Database
(from PSA)

Analyzing Period

2002 Q1 → 2025 Q3
(95 observations)

Two quarterly series

- Municipal (target)
- Commercial (external regressor)


Software Used

R
Version 4.5.1

EXPLORATORY ANALYSIS

Municipal Fisheries Production

from Q1 2002 → Q3 2025




Seasonality

Peaks in Q1 to Q2

(Summer)

Lows in Q3-Q4

(Rainy and Typhoon Season)



Rising Municipal Production


2002 → 2010




Declining Municipal Production

2011 → 2025

Shared Seasonality and a Long-term Decline



Data Cleaning and Transformation



Outliers

2 detected and replaced using `tsoutliers()`

(Chen and Liu, 1993)

Municipal Series

Log transformation to stabilize variance (the reciprocal suggested by Box-Cox, $\lambda = -0.97$, gave invalid ARIMA models)

Commercial Series


No transformation needed ($\lambda = 0.93$)

Both became stationary after differencing


Series	Transformation	Test	Statistic	p-value
Municipal	Log transformed	ADF	-1.449	0.80
		KPSS	1.227	<.01*
	First differenced (log)	ADF	-4.693	<.01*
		KPSS	0.784	<.01*
	First + seasonal differenced (log)	ADF	-6.854	<.01*
		KPSS	0.035	>0.10
Commercial	Original scale	ADF	-3.038	0.15
		KPSS	1.779	<.01*
	First differenced	ADF	-5.002	<.01*
		KPSS	0.172	>0.10

*Significant at 5%. A significant ADF test and a nonsignificant KPSS test indicate stationarity.
Shaded rows: stationary.

Municipal, first + seasonal differenced (log scale)



Commercial, first differenced



Train-Test Split

Training: 2002 Q1 – 2020 Q4 (76 quarters, 80%)

**Test: 2021 Q1 – 2025 Q3
(19 quarters, 20%)**

2002

2021

2025 Q3

The same 19-quarter test set was used for every model.

METHODOLOGY



SARIMA

Trend and seasonality through differencing plus seasonal AR and MA terms


SARIMAX

Regression with SARIMA errors, using lagged commercial production as regressors

Grid search: $p, q, P, Q \leq 2$; $d = D = 1$
Keep models with Ljung-Box $p > 0.05$
Pick the lowest test RMSE

Choosing the commercial lags

CCF of the pre-whitened commercial and municipal series



Both series were filtered with the SARIMA model fitted to commercial production before computing the CCF.

Significant lags

Lag -5 +0.2484

Lag -11 +0.2984

Lag -12 -0.2971

Negative lag = commercial production **leads** municipal production. All 7 lag combinations were tested.

METHODOLOGY

Random Forest

Averages 500 decision trees grown on bootstrap samples. Tuned: mtry, node size.

XGBoost

Adds trees one at a time, each correcting the previous errors. Tuned: depth, learning rate, rounds.

Municipal lags 1–5

+ Commercial lags 1–5
(multivariate only)

Quarter as sine and
cosine

Model
RF or XGBoost

Predicts $Y_t - Y_{t-4}$
change from same quarter
last year

Add back Y_{t-4}
= forecast

Tuning: rolling-origin cross-validation inside the training set (4 folds, 8-quarter validation windows)

METHODOLOGY

NNAR

Feedforward network on lagged log values; each model was the average of 20 networks.

Multivariate: up to 3 commercial lags from 1–8.

LSTM

Recurrent network with memory gates, fed 5-quarter sequences; units and dropout tuned, with early stopping.



Metrics and Evaluation

RMSE

MAE

MAPE

All on the test set, in metric tons (original scale). Lowest test RMSE = best model.

One-step-ahead



Recursive multi-step



SARIMA, SARIMAX, and NNAR forecast the whole test period from 2020 Q4, so they are compared with the **recursive** results of the ML and LSTM models.

RESULTS AND DISCUSSION


Best univariate traditional model: a seasonal random walk

Model	AIC	LB p	RMSE	MAE	MAPE
SARIMA(0,1,0)×(0,1,0)₄	-250.52	0.1005	15900.54	11980.22	6.06
SARIMA(0,1,0)×(1,1,0) ₄	-252.95	0.1232	17421.77	12360.22	6.32
SARIMA(0,1,0)×(0,1,1) ₄	-253.39	0.1490	19667.58	13954.66	7.14

No estimated parameters: each forecast = the same quarter last year, adjusted by the latest year-over-year growth.

Extra seasonal terms lowered AIC slightly but raised test errors.

Residuals of the final SARIMA model



Ljung-Box p = 0.1005: white noise residuals

Best multivariate traditional model: SARIMAX with lags 5 and 11

Model	Regressors	RMSE	MAE	MAPE
SARIMA(0,1,0)×(1,1,2)₄	lag 5, lag 11	12385.23	10016.27	4.86
SARIMA(0,1,0)×(1,1,2) ₄	lag 5, 11, 12	13024.00	10132.85	5.01
SARIMA(0,1,0)×(1,1,2) ₄	lag 11	13576.52	10243.94	5.15


Lag 5 alone: RMSE 16301.94, the worst of the seven combinations.

Likelihood ratio tests

Commercial lag(s) added	LRT	df	p-value
+ lag 5 alone	3.75	1	0.0527
+ lag 11 alone	7.10	1	0.0077*
+ lag 5, given lag 11	8.05	1	0.0046*
+ lag 11, given lag 5	11.39	1	0.0007*
+ lags 5 and 11 jointly	15.14	2	0.0005*

*Significant at 5%. All models share SARIMA(0,1,0)×(1,1,2)₄ errors and the same 65 training quarters.

Residuals of the final SARIMAX model



Ljung-Box p = 0.1261: white noise residuals

Compared with
SARIMA

RMSE -22.1%

MAE -16.4%

MAPE 6.06% → 4.86%

Higher commercial catch is followed by higher municipal catch

$$Z_t = \beta_5 \ddot{X}_{t-5} + \beta_{11} \ddot{X}_{t-11} + \eta_t$$
$$(1 - \Phi_1 B^4)(1 - B)(1 - B^4)\eta_t = (1 + \Theta_1 B^4 + \Theta_2 B^8)a_t$$

Z_t = log municipal production; $\ddot{X}_t$ = commercial production in thousand metric tons


+1,000 metric tons of commercial production is associated with about **+0.05%** municipal production 5 quarters later and **+0.07%** 11 quarters later.

Note: Associations, not causal effects. Seasonal parameters are read together as the error structure.

Term		Param.	Estimate	SE	p-value
Seasonal	AR(1)	Φ_1	-0.9984	0.0185	<.001*
Seasonal	MA(1)	Θ_1	0.8497	0.2340	<.001*
Seasonal	MA(2)	Θ_2	-0.1219	0.1963	0.534
comm_lag5		β_5	0.0005	0.0002	0.023*
comm_lag11		β_{11}	0.0007	0.0002	0.003*

**Significant at 5%; commercial production in thousand metric tons.*


Tree-based models barely beat naive forecasts




Test RMSEs were about twice the cross-validation RMSEs (13,292–14,407). Trees average training observations, so they cannot extrapolate the recent decline.

Last year's same quarter mattered most


Random Forest, univariate




Random Forest, Multivariate



XGBoost, univariate



XGBoost, multivariate



y_lag4 (same quarter last year) ranks at the top in both RF models.
Commercial lags = **43.1%** of total gain in multivariate XGBoost (x_lag1, x_lag3 in the top five)

Deep learning: mixed gains from commercial production

Model	Specification	Test RMSE
NNAR(4,2,2) ₄	Univariate	19,711.02
NNAR(4,0,2) ₄ + commercial lags 3, 8	Multivariate	19,145.71
LSTM	Univariate	28,978.06
LSTM	Multivariate	22,737.05

NNAR: commercial lags improved RMSE by only 2.9%; MAE and MAPE were slightly worse.


LSTM: best one-step-ahead RMSE; adding commercial production cut its recursive RMSE by 21.5%

Comparative Performance

Specification	Candidate model	RMSE	MAE	MAPE
Univariate	SARIMA	15900.54	11980.22	6.06
	Random Forest	30356.43	25706.44	12.66
	XGBoost	34443.68	29018.52	14.36
	NNAR	19711.02	13892.31	7.14
	LSTM	28978.06	22389.68	11.11
Multivariate	SARIMAX	12385.23	10016.27	4.86
	Random Forest	30257.55	24912.48	12.35
	XGBoost	29383.15	24577.06	12.13
	NNAR	19145.71	14044.44	7.19
	LSTM	22737.05	17962.71	8.66

Random Forest, XGBoost, and LSTM results are based on recursive multi-step forecasts for comparability.


SARIMA-based forecasts tracked the 2024–2025 decline best



Why traditional models performed better

- Short series (76 training observations)
- Dominated by persistent seasonality and a shifting trend, which differencing captures directly
- Tree-based models cannot extrapolate beyond the training range, so they miss the continued decline
- ML/DL test errors were roughly twice their cross-validation errors, so learned patterns did not transfer
- Recursive forecasting accumulates errors

Future Forecasts: Municipal fisheries production volume is projected to keep declining



Peak: Q2 each year

183,668 t in 2026

163,838 t in 2027

Lowest: 2027 Q3

133,141 t

Trend

Each 2027 quarter is 7–12% below the same quarter of 2026

CONCLUSION

- Traditional models beat ML and DL: **SARIMAX** was the most accurate (RMSE 12,385; 22% below SARIMA)
- Commercial catch **leads** municipal catch by about 1 and 3 years (lags 5 and 11), consistent with shared fish stocks
- Short series and bounded tree predictions held back ML and DL
- Production is projected to **keep falling** through 2027, lowest in Q3–Q4

Associations are predictive, not causal; based on national aggregate data.




Implications and future work

For policy and planning

- Use commercial landings as a ~1-year early warning
- Time support for fishers to the lean Q3–Q4 months
- Protect shared stocks: spawning-season closures and no commercial fishing in municipal waters

For future research

- Species- and region-level catch data
- Environmental drivers (SST, ENSO) to explain the 5- and 11-quarter lags
- Hybrid and regularized models for short series



END

