

Learning the Bayesian Information Criterion (BIC) with Python

Authored by
Mohammed loot

November 2, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning the Bayesian Information Criterion (BIC) with Python*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=8776>

The [Bayesian Information Criterion](#), universally known by its abbreviation **BIC**, stands as a cornerstone metric in statistical inference. Its primary function is to provide a standardized approach for comparing the goodness of fit among multiple competing [regression models](#) applied to the same dataset. Fundamentally, the utility of **BIC** stems from its unique ability to rigorously quantify the necessary balance between a model's explanatory accuracy (how well it fits the data) and its inherent complexity (the number of parameters used).

In practical data science and advanced [statistical modeling](#), researchers frequently construct and evaluate several candidate models. The ultimate objective is to pinpoint the most **parsimonious model**--the simplest structure that still provides a robust and accurate representation of the underlying data generation process. The principle guiding the use of **BIC** is straightforward: the model yielding the lowest calculated **BIC** value is deemed the superior choice, as this score signifies the optimal trade-off between minimizing prediction error and imposing a penalty for unnecessary complexity.

Crucially, the [Bayesian Information Criterion](#) imposes a significantly higher penalty on model complexity compared to its close relative, the [Akaike Information Criterion \(AIC\)](#), particularly when dealing with large sample sizes. This strong penalty biases the selection process toward simpler, more robust models, solidifying **BIC** as an indispensable tool for reliable [model selection](#) and identification.

The Mathematical Foundation of BIC

To properly leverage the [Bayesian Information Criterion](#) in advanced analyses, it is vital to understand its mathematical derivation. While the most general form of **BIC** relies on the model's [log-likelihood function](#), for the specific context of [Ordinary Least Squares \(OLS\)](#) regression, a highly practical and simplified formula based on the residual error is often employed. This formulation explicitly highlights how the model's complexity is penalized.

The following equation demonstrates how the **BIC** score is calculated utilizing the residual error in an OLS framework, where the penalty term directly incorporates the number of parameters and the sample size:

$$\mathbf{BIC}: (\text{RSS} + \log(n)d\sigma^2) / n$$

A clear understanding of each variable within this formula is paramount for accurate interpretation of the resulting score. Each component meticulously quantifies either the model's structure or its empirical performance:

d: Represents the degrees of freedom or, more simply, the **number of predictors** (parameters) included in the specific [statistical model](#) under evaluation.

n: Denotes the **total number of observations** (sample size) contained within the dataset used to train and fit the model.

σ^2 : An estimate of the variance of the error term (or residual variance) associated with the response measurements in the [regression model](#).

RSS: The **Residual Sum of Squares (RSS)**, which quantifies the unexplained variance--the discrepancy between the observed data points and the values predicted by the model. A lower RSS signifies a superior fit to the training data.

TSS: The Total Sum of Squares, representing the total inherent variation in the dependent response variable. (While primarily used for calculating R-squared, it is listed here to contextualize the relationship between explained and unexplained variance.)

The critical component, $\log(n)d\sigma^2$, functions as the explicit **penalty factor**. As the number of observations (n) increases, the logarithmic weighting ensures that the penalty applied for introducing an additional predictor (d) increases disproportionately. This mathematical structure compels data analysts to prioritize simpler model architectures when abundant data is available, thereby promoting more generalizable and less overfitted solutions.

Implementing BIC Calculation in Python

Fortunately for modern data practitioners, the complexities of manually calculating the **BIC** score using the aforementioned formula are entirely abstracted away by robust statistical libraries in the Python ecosystem. To efficiently compute the **BIC** for various [regression models](#), we rely almost exclusively on the highly respected [statsmodels](#) package, which is specifically designed for rigorous statistical estimation and testing.

Within this package, the `statsmodels.regression.linear_model.OLS()` function is the ideal choice for linear modeling tasks. Once a model object is defined, fitted, and estimated, the resulting fit object possesses an intrinsic property named `.bic`. This property automatically computes and returns the accurate **BIC** value corresponding to the fitted linear model, integrating the log-likelihood and parameter count seamlessly. This streamlined capability allows data scientists to compare dozens of potential models rapidly without engaging in complex manual calculations.

The following practical demonstration illustrates how to harness this built-in functionality within a Python environment. We will proceed to calculate and subsequently interpret the **BIC** scores derived from a set of competing linear [regression models](#), thereby facilitating an objective and quantitative [model selection](#) procedure.

Practical Example: Comparing Regression Models with the MTCars Dataset

We now move to a hands-on application, utilizing the classic automotive dataset, **mtcars**, which

contains thirty-two observations on eleven vehicle characteristics. Our analytical goal is to construct and evaluate two distinct linear regression models, both aiming to accurately predict the dependent variable `mpg` (miles per gallon) but using different combinations of predictor variables from the dataset. This comparison will illustrate the power of [BIC](#) in distinguishing between competing hypotheses.

The first essential step is to prepare the computational environment by importing the necessary libraries and loading the dataset. We specifically employ [statsmodels](#) for robust regression estimation and **pandas** for efficient data manipulation and loading.

The foundational code block below handles the data acquisition and displays the initial structure of the **mtcars** data frame:

```
from sklearn.linear_model import LinearRegression
import statsmodels.api as sm
import pandas as pd

#define URL where dataset is located
url = "https://raw.githubusercontent.com/Statology/Python-Guides/main/mtcars.csv"

#read in data
data = pd.read_csv(url)

#view head of data
data.head()

model mpg cyl disp hp drat wt  qsec vs am gear carb
0 Mazda RX4 21.0 6 160.0 110 3.90 2.620 16.46 0 1 4 4
1 Mazda RX4 Wag 21.0 6 160.0 110 3.90 2.875 17.02 0 1 4 4
2 Datsun 710 22.8 4 108.0 93 3.85 2.320 18.61 1 1 4 1
3 Hornet 4 Drive 21.4 6 258.0 110 3.08 3.215 19.44 1 0 3 1
4 Hornet Sportabout 18.7 8 360.0 175 3.15 3.440 17.02 0 0 3 2
```

With the data successfully loaded, we define and fit our two candidate [regression models](#). Both models share the dependent variable `mpg` (miles per gallon) and utilize two distinct predictors:

Model 1: Hypothesizes that fuel efficiency (`mpg`) is best explained by engine displacement (`disp`) and the time required to complete a quarter-mile run (`qsec`). Formula: $mpg = \beta_0 + \beta_1(disp) + \beta_2(qsec)$

Model 2: Hypothesizes that `mpg` is best explained by engine displacement (`disp`) and vehicle weight (`wt`). Formula: $mpg = \beta_0 + \beta_1(disp) + \beta_2(wt)$

We first fit **Model 1** using the [statsmodels](#) OLS methodology. Immediately after fitting, we access the intrinsic `.bic` property to retrieve the necessary comparative metric:

```
#define response variable
```

```
y = data
```

```
#define predictor variables
```

```
x = data]
```

```
#add constant to predictor variables
```

```
x = sm.add_constant(x)
```

```
#fit regression model
```

```
model = sm.OLS(y, x).fit()
```

```
#view BIC of model
```

```
print(model.bic)
```

```
174.23905634994506
```

The resulting **BIC** score for **Model 1**, which incorporates displacement and quarter-mile time as predictors, is calculated to be approximately **174.239**.

Next, we conduct the identical fitting process for **Model 2**, replacing the `qsec` predictor with `wt` (vehicle weight) to evaluate whether this alternative combination provides a more efficient and parsimonious fit to the data, as judged by the [Bayesian Information Criterion](#):

```
#define response variable
```

```
y = data
```

```
#define predictor variables
```

```
x = data]
```

```
#add constant to predictor variables
```

```
x = sm.add_constant(x)
```

```
#fit regression model
```

```
model = sm.OLS(y, x).fit()
```

```
#view BIC of model
```

```
print(model.bic)
```

```
166.56499196301334
```

The resulting **BIC** value for **Model 2**, incorporating engine displacement and weight, is determined to be **166.565**.

Interpreting BIC Results and Selecting the Optimal Model

The final, and most critical, phase of utilizing the **Bayesian Information Criterion** is the direct comparison and interpretation of the calculated scores. As previously established, the core objective of using **BIC** is to identify the candidate model that yields the minimum score, as this signifies the most favorable balance between minimizing residual error and preventing overfitting due to excessive parameters.

We compare the two values derived from our fitted models:

Model 1 BIC: 174.239

Model 2 BIC: 166.565

Since **Model 2** exhibits a distinctly lower **BIC** value (166.565 is significantly less than 174.239), we confidently conclude that **Model 2** is the superior fitting model. This model offers a more compelling explanation of the variation in miles per gallon while maintaining the exact same level of model complexity (both models included two predictors plus an intercept). The quantitative difference in the scores provides strong evidence that vehicle weight (wt) is a statistically superior and more influential predictor of fuel efficiency, holding engine displacement constant, compared to quarter-mile time ($qsec$).

Once the superior model has been objectively identified through this quantitative [model selection](#) methodology, analysts should proceed with a deeper diagnostic analysis. This typically involves examining critical model summary statistics, such as the [R-squared](#) value, assessing the significance of the individual beta coefficients (p-values), and performing residual diagnostics. These subsequent steps ensure that the model's assumptions are met and fully determine the exact nature and strength of the relationship between the chosen predictors and the response variable.

Alternative Metrics for Model Selection

While the **Bayesian Information Criterion** is recognized for its strong theoretical foundation and its preference for parsimony, it represents only one facet of the complete model assessment toolkit. Analysts should routinely consider other metrics for comparing [statistical models](#) to gain a comprehensive view of performance and robustness. The two most frequently used alternatives or complements to **BIC** are the [Akaike Information Criterion \(AIC\)](#) and the **adjusted R-squared**.

The **AIC** operates on principles similar to **BIC**, penalizing models for increasing the number of

predictors. However, its penalty function is generally less severe than that of **BIC**, especially when the sample size is small to moderate. In contrast, the **adjusted R-squared** metric focuses specifically on the proportion of variance explained by the predictors. It modifies the standard [R-squared](#) downward only when newly added predictors do not significantly contribute to the model's explanatory power, offering a powerful measure of fit that accounts for the degrees of freedom.

By utilizing a combination of these criteria--**BIC** for favoring simplicity, **AIC** for a slightly less strict penalty, and **adjusted R-squared** for assessing explained variance--data scientists can achieve a multi-faceted and highly reliable assessment of their model performance and ensure sound [model selection](#).

For those seeking to expand their knowledge on calculating these related concepts in Python, detailed resources are available explaining how to compute each of these alternative metrics for [regression models](#):