

# How to Easily Sort Values in Pandas Crosstabs

Authored by  
**stats writer**

November 20, 2025

## RECOMMENDED CITATION

stats writer (2025). *How to Easily Sort Values in Pandas Crosstabs*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=98390>

The `Pandas crosstab` function is an indispensable tool for performing frequency analysis, allowing users to build a simple cross-tabulation table from data stored within a `DataFrame`. While extremely useful for summarizing categorical data, the resulting table often requires specific reordering for better interpretability, particularly when dealing with large datasets or complex reporting requirements. Understanding how to control the order of rows and columns--which are essentially the index and column headers of the resulting table--is a fundamental skill for advanced data manipulation using `Pandas`.

By default, when creating a contingency table using `pd.crosstab`, both the row and column indices are sorted alphabetically. However, real-world data analysis frequently necessitates sorting these indices based on custom criteria, such as descending order (Z-A) or based on the magnitude of counts contained within the body of the table. Because the output of `pd.crosstab` is a standard `DataFrame`, we leverage standard Pandas methods, most notably `sort_index()`, coupled with the appropriate `axis` parameter, to achieve precise control over the presentation of our cross-tabulation.

This comprehensive guide details the precise methods for controlling the sorting order of the indices (rows) and columns of a Pandas cross-tabulation table. We will explore the critical difference between sorting by rows versus sorting by columns, relying on the robust `sort_index()` function to achieve our desired output structure. This method allows analysts to override the default alphabetical ordering, ensuring that the visual presentation of the data aligns perfectly with the analytical objective.

To achieve the desired rearrangement of your contingency table, you primarily employ one function combined with different axis specifications:

### The Core Mechanism: Using `sort_index()`

The standard way to rearrange the rows or columns based on their labels (index values) in a Pandas object is by utilizing the `sort_index()` method. This method is highly flexible and accepts several crucial parameters, most importantly the `axis` parameter, which dictates whether the sorting operation applies horizontally (columns) or vertically (rows), and the `ascending` parameter, which controls the direction of the sort (A-Z or Z-A). It is critical to apply this method directly to the output of the `crosstab` function to ensure that the resulting table is immediately presented in the required order.

When chaining this function directly after the `pd.crosstab()` call, the results are immediately sorted before being displayed. This practice ensures that the data presentation is instantly optimized for the required analysis. The power of this method lies in its efficiency, as it operates directly on the index objects rather than requiring complex multi-level sorting logic applied to the underlying data values themselves. Understanding the index structure of the cross-tabulation is

key to successfully applying this sorting technique.

We typically define the sort order using the `ascending` parameter. Setting `ascending=True` (the default) results in alphabetical or numerical ascending order, while setting `ascending=False` yields descending order. This simple binary switch provides complete control over the presentation of the index labels, whether they represent time periods, categories, or other unique identifiers used in the cross-tabulation. For custom reporting, setting `ascending=False` is frequently used to highlight the most recent or numerically largest categories first.

## Method 1: Sorting Crosstab by Row Index Labels (axis=0)

To specifically sort the rows of the resulting cross-tabulation, which correspond to the primary index of the table, we must apply the `sort_index()` function and set the `axis` parameter to 0. In Pandas conventions, `axis=0` refers exclusively to the index (rows). This is often necessary when the row categories, such as team names or dates, need to be displayed in a reverse alphabetical or reverse chronological order, deviating from the inherent default ascending sort.

The syntax below illustrates how to generate a crosstab and immediately sort the resulting row index in descending order. Note the placement of the method call directly after the creation of the cross-tabulation table, ensuring streamlined execution and output generation. This demonstrates the power of method chaining in Pandas for concise data manipulation.

```
pd.crosstab(df.col1, df.col2).sort_index(axis=0, ascending=False)
```

Using `axis=0` ensures that the entire row--including all the corresponding counts across the columns--moves together during the sorting process, maintaining the integrity of the contingency data. This approach keeps the analysis accurate while optimizing the visual flow for the end-user or report reader. This command only rearranges the rows based on their labels; the column arrangement remains unaffected.

## Method 2: Sorting Crosstab by Column Index Labels (axis=1)

Conversely, if the objective is to rearrange the columns of the cross-tabulation table, we utilize the same `sort_index()` method but set the `axis` parameter to 1. In Pandas, `axis=1` always refers to the columns. Sorting columns is crucial when dealing with variables that inherently possess a natural order (like size categories: small, medium, large) or when column categories need a specific logical ordering that overrides the default alphabetical sort for better comparison.

The following syntax demonstrates sorting the column labels in descending order. While the data within the cells remains the same, their horizontal arrangement is altered, which can significantly impact how quickly patterns are identified by the observer. This control is vital in data visualization

preparation, as the sequence of columns often guides the eye during comparative analysis.

### **pd.crosstab(df.col1, df.col2).sort\_index(axis=1, ascending=False)**

By specifying `axis=1`, we instruct Pandas to treat the column labels (the secondary variables used in the cross-tabulation) as the primary keys for the sorting operation. This level of granular control is one of the key reasons Pandas remains the dominant library for data manipulation in Python, offering flexibility far beyond standard spreadsheet applications.

## **Practical Example Setup: Creating the Data and Baseline Crosstab**

To demonstrate these sorting methods effectively, we will first establish a sample dataset representing player statistics, categorized by team and position. This dataset, defined as a Pandas DataFrame, contains three columns: `team`, `position`, and `points`. The cross-tabulation will summarize the frequency of players, showing how many players from each `team` occupy each `position`. This initial step is necessary to provide the structured data upon which our sorting functions will operate.

The fundamental step is importing the necessary library and initializing the data structure. Observe how the `pd.crosstab` function is used to aggregate the counts based on the intersection of the categorical variables `df.team` (rows) and `df.position` (columns). This initial output serves as our baseline, where the indices are sorted alphabetically by default (A, B, C for teams; F, G for positions). We use the `print()` function to visualize this baseline table before modification.

### **import pandas as pd**

```
#create DataFrame
```

```
df = pd.DataFrame({'team': ,  
'position':,  
'points': })
```

```
#create crosstab to display count of players by team and position
```

```
my_crosstab = pd.crosstab(df.team, df.position)
```

```
#view crosstab
```

```
print(my_crosstab)
```

```
position F G
```

```
team
```

```
A 1 2
```

```
B 3 1
```

## C 2 2

As clearly visible in the output, the teams (rows) are ordered A, B, C, and the positions (columns) are ordered F, G. The examples that follow will demonstrate how to reverse these orders using our targeted sorting techniques, showing the immediate visual impact of the `sort_index()` method.

### Example 1: Sorting Rows in Descending Order (Z to A)

A common requirement in reporting is to display categories in reverse order, perhaps to emphasize the latest or most recent entries if the row index represents time. In our current example, we want to sort the teams (rows) in reverse alphabetical order: C, B, then A. This is achieved by combining the cross-tabulation creation with the `sort_index()` method, explicitly setting `axis=0` and `ascending=False`. This combination signals to Pandas that the sorting operation should target the row index and execute in reverse order.

The use of `ascending=False` is the key differentiator here, as it reverses the default sorting behavior which is always ascending (A-Z or 1-10). When `axis=0` is specified, the sorting operation examines the row index labels (in this case, 'A', 'B', 'C') and reorders the entire table vertically based on those labels. This ensures that the structural integrity of the cell counts remains perfectly preserved, as the counts move along with their corresponding team labels.

**#create crosstab with rows sorted from Z to A**

```
pd.crosstab(df.team, df.position).sort_index(axis=0, ascending=False)
```

```
position F G
team
C 2 2
B 3 1
A 1 2
```

Upon reviewing the output, it is clear that the rows of the cross-tabulation are now structured according to the desired descending alphabetical order (C, B, A). It is crucial to remember that the `crosstab()` function, when used without explicit sorting, always defaults to arranging both the row and column indices in ascending alphabetical or numerical order. Therefore, methods like `sort_index()` are necessary whenever a non-default presentation order is required for the index labels, particularly for client-facing reports.

### Example 2: Sorting Columns in Descending Order (Z to A)

Sorting columns requires the same methodology used for sorting rows, with the critical adjustment

of the `axis` parameter. Here, we set `axis=1` to target the column headers, which represent the different positions in our dataset (F and G). By maintaining `ascending=False`, we instruct Pandas to sort these position categories in reverse alphabetical order, meaning 'G' will appear before 'F'. This adjustment is commonly used to place critical or high-frequency columns closer to the beginning of the table.

The rearrangement of columns is often necessary when certain columns need visual prioritization, perhaps placing the totals or a specific key performance indicator column prominently on the left side of the table. Although our example uses simple alphabetical labels, this technique applies equally well to numerical or temporal column labels, allowing for flexible data presentation regardless of the data type used for the column index.

**#create crosstab with columns sorted from Z to A**

```
pd.crosstab(df.team, df.position).sort_index(axis=1, ascending=False)
```

```
position G F
team
A 2 1
B 1 3
C 2 2
```

The resulting cross-tabulation demonstrates the successful vertical rearrangement of column indices. The row labels (teams A, B, C) remain in their default alphabetical order, but the column labels (positions) are now sorted G then F. This highlights the independent nature of index sorting: altering `axis=1` has no effect on the row order, and vice versa, providing precise control over table structure.

It is important to reiterate the default behavior: if `sort_index()` is omitted, the `crosstab()` function ensures that column values are displayed in ascending alphabetical order. Mastering the control over the `axis` parameter is paramount for controlling the visual layout of complex data summaries and ensuring immediate data readability.

## Advanced Sorting: Ordering by Calculated Values, Not Index Labels

While `sort_index()` is perfect for rearranging the rows or columns based on their labels (e.g., team names or position types), data analysis often requires sorting based on the actual cell counts or calculated summary values within the body of the cross-tabulation. For instance, we might want to list teams from the one with the highest total player count in a specific position to the one with the lowest, regardless of their alphabetical name.

To achieve value-based sorting on the resulting contingency table, we must switch from using

`sort_index()` to using the `sort_values()` method. This method requires specifying the column or row based on which the sort should occur. When sorting rows based on a column's value, you must specify the column name. For example, to sort the teams based on the total count of players in position 'F', you would use `.sort_values(by='F', ascending=False)`. This operation will reorder the rows based on the magnitude of the counts within the designated column.

A key consideration here is that sorting by column values implicitly sorts the rows. For example, if we sort the rows based on the counts in the 'G' column, the row labels (teams A, B, C) will be reordered according to the counts in the 'G' column. This technique is extremely powerful for identifying leading categories immediately, rather than requiring the user to scan the entire table for the maximum values, thereby focusing attention on the most significant data points.

## Handling Multi-Index Sorting in Crosstab Results

In more complex scenarios, the `pd.crosstab()` function might be called with multiple variables for rows or columns, resulting in a hierarchical or multi-index structure. Sorting these complex structures requires careful management of the `level` parameter within the `sort_index()` function. The `level` parameter allows you to specify exactly which level of the hierarchical index should be used as the key for sorting, ensuring that the hierarchy is respected during rearrangement.

For a multi-index row structure, setting `level=0` sorts based on the highest level index (the primary category), while `level=1` sorts within the groups defined by the primary category. For instance, if rows are defined by `(region, city)`, setting `level=0` sorts by region. If two regions are identical, the sorting of their respective cities is determined by the next level up, unless `level=1` is explicitly used to sort the cities. This fine-grained control is indispensable for structured reports that use multiple dimensions for categorization.

When working with multi-level columns (`axis=1`), the same principle applies. Defining the correct `level` is essential to ensuring that the hierarchical structure is maintained and sorted logically. Failing to specify the level when dealing with a multi-index will often lead to a cryptic error or unexpected sorting behavior, underscoring the importance of understanding index structure when performing advanced manipulations in Pandas. Always inspect the index names and structure before attempting multi-level sorting.

## Summary of Sorting Parameters and Best Practices

To summarize the fundamental differences and best practices for sorting the output of `pd.crosstab`, we rely primarily on two parameters: `axis` and `ascending`. Properly combining these parameters gives us complete control over the layout of our contingency table. Here is a quick reference guide demonstrating the typical use cases:

**Row Sorting (Index Labels):** Use `.sort_index(axis=0, ascending=...)`. This sorts the primary categorical variables defining the rows based on their label value.

**Column Sorting (Column Labels):** Use `.sort_index(axis=1, ascending=...)`. This sorts the secondary categorical variables defining the columns based on their label value.

**Descending Order:** Set `ascending=False` for reverse alphabetical (Z-A) or highest-to-lowest numerical order, often used to highlight extremes.

**Ascending Order (Default):** Set `ascending=True` or omit the parameter for standard alphabetical (A-Z) or lowest-to-highest numerical order.

**Sorting by Value:** Use `.sort_values(by='ColumnName', axis=0)` to reorder rows based on the magnitudes within a specific column.

For analysts, prioritizing clarity and consistency in reports is crucial. Always choose a sorting method that directly supports the narrative or hypothesis being tested. If the goal is comparison, sorting rows by a key metric using `sort_values()` is preferable. If the goal is chronological display or ensuring a specific categorical order, sorting the index using `sort_index()` with an appropriate `ascending` setting is the correct and most efficient choice.

Ultimately, the ability to generate a clean, custom-sorted contingency table is a cornerstone of effective exploratory data analysis and reporting using the Pandas library. By mastering `sort_index()` and its critical parameters, users can transform raw frequency counts into highly polished, readable data summaries, ready for presentation or further visualization.