Weighted averages

Combine group averages using their underlying counts, then compare the answer with an ordinary mean.

Synthetic example · independently invented data

One-page print edition · A3 landscape

01 Inspect the inputs

02 Apply the method

03 Check the result

INVENTED INPUTS

Different group sizes

Group	Mean · min	Count
Group A	4	8
Group B	6	12
Group C	9	20

Three invented groups

Each mean describes its group. The counts determine how much that group contributes to the combined average.

SELECTED PYTHON

The working calculation

Pandas is pd, NumPy is np, and rows is the three-row input table.

```
groups = pd.DataFrame(rows)
weights = groups['count']
if weights.lt(0).any() or weights.sum() == 0:
    raise ValueError('Use non-negative, non-zero weights')
weighted = np.average(
    groups['mean_minutes'], weights=weights
)
ordinary = groups['mean_minutes'].mean()
```

The code shown here is executed to calculate this example's results.

COMPUTED RESULT

Give each observation its weight

7.10

minutes · weighted mean

40 observations represented

284 ÷ 40 = 7.10



The ordinary mean gives each group equal influence. The larger, higher-valued group raises the weighted result.

What this demonstrates

A mean of means can misrepresent the underlying observations when group sizes differ. Weighting preserves their contribution.

Check the boundary

Reject negative weights and an all-zero denominator. This fixture uses complete means and positive counts; missing-data policy would need a separate decision.

Thomas Eapen
nowlabs.net/tangerine-woa

New illustrative material explaining the approach. Values, configurations and results shown here are not historical workplace records.