

Comparing indices without fooling yourself

REBASING

Setting two series to a common starting value, usually 100, so their subsequent moves can be compared as percentages. Index levels are unitless and arbitrary, so comparing them directly - or plotting them on one axis - compares the accidents of their construction rather than their performance.

One index is at 8,200 and another at 40,000. The second is not larger, higher, better or more expensive. The numbers are not in the same units, because they are not in any units.

Where the levels came from

An index level is whatever the starting value was, multiplied by everything since. A base of 100 in 1984 and a base of 1,000 in 1896 produce wildly different numbers for identical performance.

So the level carries information about when the index started and what its divisor has done, and almost none about the market it measures.

Rebasing is the whole technique

Set both series to 100 at the same date and plot from there. Every subsequent value is then a percentage of that starting point, and the two lines can be read against each other.

Choose the date deliberately, because it does all the work. Rebase from a trough and everything looks strong; rebase from a peak and everything looks weak. Comparisons drawn from a convenient start date are the most common way honest-looking charts mislead.

Two axes is not a solution

Plotting two indices on separate scales makes the lines fit the same picture, and that is exactly the problem: the axes can be chosen to make two unrelated series appear to move together, or to hide that they do.

If a comparison needs two axes to look meaningful, it is usually the axes doing the work.

What is still not comparable after rebasing

Different weighting rules - a price-weighted and a cap-weighted index are measuring differently before you start. Different sector mixes, so you may be comparing technology against banks and calling it two countries. Different currencies, which move independently of the businesses. Total return against price return: one includes dividends, the other does not, and the gap compounds.

That last one catches people out over long horizons more than any of the others, because the two versions of the same index share a name and diverge quietly for years.

The honest comparison

Same start date, chosen for a reason you can state. Percentages, not levels. Same return basis. And a sentence acknowledging what still differs, because two indices are never really measuring the same thing - that was lesson 2 of module 1.

Check your understanding

1. Why is comparing index levels directly meaningless?

- a. Levels are updated at different frequencies
- b. A level reflects the index's starting value and divisor history, not the market's size or performance
- c. Levels exclude dividends

Answer: b. The number is whatever the base was, compounded by everything since. Two indices with identical performance can sit at completely different levels.

2. Which difference quietly compounds over long horizons and shares a name?

- a. Price return against total return
- b. Cap weighting against price weighting
- c. Cash against futures

Answer: a. One version includes dividends and the other does not. They carry the same index name and diverge steadily for years, which makes it the easiest mismatch to miss.