

Time series and market indices

This session teaches students to build a market index the way it is actually built, by bucketing cleared prices into time periods and tracking a summary statistic across those periods, and to treat the exercise as a lesson in methodology rather than a shortcut to a market call. Students construct a quarterly demand index for a brand or category from sold-lot records, using date bucketing in SQL to group transactions and a stable statistic such as median realized price or count of sold lots to summarize each quarter, then compare the shape of two brands' indices side by side.

What today looks like

- Length: 90 minutes. A 60 minute and a 120 minute variant are noted under Teaching notes.
- Level: advanced undergraduates or graduate students in a time series analysis, applied econometrics, or market research methods course, who have taken or are concurrently taking an introductory statistics or econometrics course.
- Access: the no-code sandbox at sandbox.altfndata.com is the primary path, and is all that is needed for the core session. No API key is required. An optional code extension, retrieving the underlying sold-lot records via the production API and reproducing the quarterly bucketing in Python, requires a class API key from info@altfndata.com.
- Goal: students leave able to build a quarterly demand index from sold-lot records in the sandbox, state precisely where an index stops being trustworthy because of recency under-ingestion, and describe a two-brand comparison as a methodological exercise in relative pattern rather than a market call.

OBJECTIVES

By the end of class students can

- Explain what a market index built from transaction data represents, and distinguish it from a price index built from listed or asking prices.
- Bucket sold-lot records into calendar quarters using SQL date functions and compute a quarterly summary statistic for each bucket.
- Build a quarterly demand index for a single brand or category and identify the reference period against which later quarters would be compared.
- Explain, in specific and concrete terms, why the newest quarters of any such index must be treated as under-ingested rather than as evidence of a market decline, and identify where in a chart that caveat should appear.
- Compare the index shapes of two brands or categories and describe the comparison as a methodological exercise in relative pattern, not as a claim about which one is appreciating.
- (Extension) Retrieve the underlying sold-lot records programmatically via the production API and reproduce the quarterly bucketing in Python.

How the session runs

- 0 to 15 min: Introduce the idea of a transaction-based market index, contrast it with a listed-price index, and preview the recency caveat as the session's central discipline.
- 15 to 25 min: Sandbox orientation. Confirm students can open a category table and locate `sale_date`, `usd_price_decimal`, and `designer` in the data dictionary.
- 25 to 45 min: Guided demo, build a quarterly demand index for a single brand using date bucketing and a stable summary statistic.
- 45 to 60 min: Guided demo, build the same index for a second brand and overlay the two, discussing what the comparison can and cannot support as a conclusion.
- 60 to 75 min: Small-group exercise, each group builds a quarterly index for a different brand or category and marks, on their own chart, the exact quarter after which the recency caveat applies.
- 75 to 85 min: Class discussion, groups present their index shapes and defend where they drew the recency cutoff.
- 85 to 90 min: Wrap-up and homework assignment.

The in-class walkthrough

- Open sandbox.altfndata.com, sign in, and select the fine art data table (or another category the instructor prefers) from the SQL editor dropdown.
- Open the coverage browser tab and confirm how a chosen artist or brand is written in the designer field before filtering.
- Return to the SQL editor and run the first guided query below, which buckets sold lots for one artist into calendar quarters and computes a median realized price per quarter.
- Open the pre-built charts tab and render the result as a line chart. Explain that this line is a quarterly demand index in miniature, a stable statistic tracked over time, not a single headline number.
- Point to the final one or two quarters and state the recency caveat plainly: these quarters are still being ingested, additional records will continue to arrive for them, and any apparent drop at the tail of the line should be treated as incomplete data, not as falling demand. Mark that cutoff visually on the chart, for example with a vertical line or shaded region.
- Run the second guided query, the same bucketing and statistic for a second artist or brand, and overlay both lines on one chart.

The core query

- Run this in the sandbox SQL editor:
 - `SELECT date_trunc('quarter', CAST(sale_date AS date)) AS sale_quarter,`
 - `approx_percentile(usd_price_decimal, 0.5) AS median_price_usd,`
 - `COUNT(*) AS sold_lots`
 - `FROM all_fine_art_data`
 - `WHERE status = 'sold'`
 - `AND designer LIKE '%Warhol%'`
 - `GROUP BY date_trunc('quarter', CAST(sale_date AS date))`
 - `ORDER BY sale_quarter;`

DISCUSS

Questions for the room

- What does a quarterly median-price index built from auction transactions represent that a listed asking-price index would not, and what does it miss that a listed-price index would capture?
- Why is median realized price a more stable summary statistic for an index than a single record sale or a raw average, given how concentrated value can be among a small number of lots?
- Precisely where would you draw the line on an index chart between quarters you trust and quarters you flag as under-ingested, and how would you defend that choice to someone reading the chart?
- Two brands' index lines diverge sharply in their final two quarters. What is the more likely explanation given the recency caveat, and what would you need to check before ruling out an ingestion artifact?
- What is lost when a demand index is built on price alone, without also tracking sold-lot count or sell-through alongside it?

This week's assignment

- Each student selects two brands or two categories represented in the same data table and builds a quarterly demand index for each, using either median realized price or sold-lot count as the summary statistic, consistently applied to both.
- The submission is a short report, no more than two pages, that includes the SQL queries used, an overlaid chart of the two index lines, an explicit marking of the quarter after which the recency caveat applies, and a written comparison of the two index shapes framed strictly as a methodological observation about relative pattern, with an explicit statement that the comparison is not a claim about which brand is appreciating.
- Grading criteria:
- correct quarterly bucketing and choice of a stable summary statistic (30 percent), explicit and correctly placed recency caveat (30 percent), quality and restraint of the comparative interpretation (25 percent), and clarity of the report (15 percent).

NEXT

Wrapping up

- Restate the one-sentence takeaway before students leave.
- Point to the facilitator guide for discussion guidance and teaching notes.
- Questions or a class API key: info@altfndata.com