

## Problem set 10: Time series and market indices

---

**Module 10.** Total: 100 points. Format: no-code sandbox, with an optional API extension in Problem 5 using a shared class API key from your instructor.

### Before you start

This problem set asks you to build a quarterly demand index the way one is actually built in this dataset, by bucketing sold-lot records into calendar quarters and tracking a stable summary statistic across those quarters, and to treat the whole exercise as a lesson in methodology rather than a shortcut to a market call. Choose one category table and one brand, house, or artist within it to use as your primary index, and a second brand or category in the same table for the comparison problem. Confirm the exact spelling of each in the coverage browser before filtering, since exact-match filters are spelling-sensitive.

---

### Problem 1 (15 points): What an index represents

In four to five sentences, explain what a quarterly demand index built from sold-lot transaction records represents, how it differs from a price index built on listed or asking prices, and why a stable statistic such as median realized price, rather than a single large sale, is the right summary to track per quarter.

### Problem 2 (25 points): Build the quarterly index

Write a SQL query against your chosen category table that buckets sold lots for your chosen brand into calendar quarters and computes, for each quarter, both the median realized price and a count of sold lots. Render the result as a line chart. Then state, in one or two sentences, exactly which of the final quarters on your chart you would mark as under-ingested and why.

-- your query here

### Problem 3 (20 points): The recency caveat, explicitly

In three to four sentences, explain precisely why the newest one or two quarters of any index built on this dataset must be treated as under-ingested rather than as a signal of falling demand, and describe where on a chart that caveat should appear so that a reader who only glances at the chart still absorbs it.

### Problem 4 (25 points): Compare two brands as method, not appreciation

Write a second SQL query, structured the same way as Problem 2, for your second brand or category, and describe how you would overlay both lines on one chart. Then, in three to four sentences, describe what the comparison of the two index shapes can support, framed strictly as a methodological observation about relative pattern, and explicitly state one thing the comparison cannot support, such as a claim that one brand is appreciating faster than the other.

-- your query here

### Problem 5 (15 points): The same index via the API (optional code extension)

Using the shared class API key your instructor has provided, write the endpoint and JSON request body you would send to POST `/v1/tables/{name}/query` to retrieve the underlying sold-lot records needed to reproduce your Problem 2 index. Note that the API returns rows, not server-side date bucketing, so describe in one to two sentences how you would paginate on offset and perform the quarterly grouping yourself, in the tutorials

notebook or with the Python client.

```
POST /v1/tables/<table_name>/query
Header: X-API-Key: <class key>
Body:
{
  ...
}
```

---

**Submission.** Turn in this file with your SQL, JSON body, chart descriptions, and written answers filled in.