

Auction theory and behavioral economics

This session uses real auction transaction data to test the ideas that auction theory and behavioral economics predict about how buyers and sellers behave in a saleroom. Rather than treating the estimate, the reserve, and the hammer as abstractions from a textbook, students compute pricing power and sell-through directly from more than 10 million transaction records spanning more than 100 auction houses, and use those figures to examine three classic ideas together:

What today looks like

- Length: 90 minutes. A 60 minute and a 120 minute variant are noted under Teaching notes.
- Level: upper-level undergraduates and graduate students who have completed an introductory microeconomics or game theory course. No finance industry background is assumed.
- Access: the no-code sandbox at sandbox.altfndata.com only. No API key is needed. Students self-register with a work or school email and are approved automatically.
- Goal of the session: students leave able to compute pricing power and sell-through for a brand or category and argue, using the figures alone, which of anchoring, the winner's curse, or the reserve price the data most clearly supports.

OBJECTIVES

By the end of class students can

- Define anchoring in the context of a pre-sale estimate and explain why measuring pricing power relative to the high estimate is a direct test of an anchoring effect.
- Compute pricing power, the median of realized price over the high estimate, for a category or brand, and interpret a value above or below 1.0.
- Explain the winner's curse in a competitive-bidding setting and describe what pattern in the pricing power data would be consistent with it.
- Describe the economic function of a reserve price and explain why an unsold lot is a rational outcome for a seller rather than a failure of the auction.
- Compute sell-through for a category or house and interpret it as a demand signal that is shaped jointly by buyer appetite and consignment quality.
- Identify the limits of using realized-price data alone to test behavioral predictions, including the recency caveat on the newest quarters of data.

How the session runs

- 0 to 15 min: Introduce the three ideas the session will test: the estimate as an anchor, the winner's curse, and the reserve price as the mechanism behind an unsold lot. Preview pricing power and sell-through as the two measures students will compute.
- 15 to 25 min: Sandbox orientation. Confirm students can open a data table and locate designer, item_title, sale_date, usd_price_decimal, sale_estimates_high_usd_price, status, and vendor in the data dictionary.
- 25 to 45 min: Guided demo, compute pricing power for a single category or brand and discuss what a value above 1.0 implies about the estimate as an anchor.
- 45 to 60 min: Guided demo, compute sell-through for the same category and discuss the reserve price as the mechanism separating a sold lot from an unsold one.
- 60 to 75 min: Small-group exercise, each group selects a different category or brand, computes both measures, and prepares a short argument about whether the pattern is consistent with anchoring, the winner's curse, or neither.
- 75 to 85 min: Class discussion, groups report their findings and the class debates where the textbook predictions hold and where the data complicates them.
- 85 to 90 min: Wrap-up and homework assignment.

The in-class walkthrough

- Open sandbox.altfndata.com, sign in, and select a data table, for example the watches data, from the SQL editor dropdown.
- Open the coverage browser tab and confirm how a well-known brand appears in the designer field. Spelling matters for exact filters.
- Return to the SQL editor and run the first guided query, pricing power for that brand, and read the result aloud. Ask students what a value above 1.0 implies about the pre-sale estimate as an anchor for bidding behavior.
- Introduce the winner's curse and explain that it predicts the eventual buyer in a competitive bid, on average, pays a price that reflects some overestimation of value relative to the field of bidders. Ask students what pattern in pricing power, sustained above 1.0 across many lots rather than a single outlier, would be more consistent with a systematic anchoring or overbidding effect than with noise.
- Run the second guided query, sell-through for the same brand, and explain the reserve price as the seller's protection against selling below an acceptable floor. Point out that an unsold lot is retained by the consignor and marked unsold, not discarded from the record.
- Explain the recency caveat before anyone reads a single recent period as evidence of a trend: the newest quarters of data are still being ingested, so any time-bucketed comparison should rely on the stable, longer-run pattern rather than the most recent few months.

The core query

- Run this in the sandbox SQL editor:

```
- SELECT designer,  
-         approx_percentile(usd_price_decimal / sale_estimates_high_usd_price, 0.5) AS pricing_power,  
-         COUNT(*) AS sold_lots  
- FROM all_watches_data  
- WHERE status = 'sold'  
-       AND sale_estimates_high_usd_price > 0  
-       AND designer LIKE '%Rolex%'  
- GROUP BY designer;
```

Questions for the room

- If pricing power for a brand sits well above 1.0 across hundreds of lots, is that better explained by anchoring on the estimate, by a winner's curse dynamic among competitive bidders, or by the house deliberately setting a conservative estimate? How would you tell these apart with this data alone?
- Why is an unsold lot a rational outcome of a reserve price, rather than evidence that the auction mechanism failed?
- Sell-through blends buyer demand and consignment quality into a single figure. What additional information would help you separate the two?
- The winner's curse is typically described in common-value settings, where bidders share uncertainty about a single true value. How well does that setting describe a luxury auction, where taste and private value plausibly differ across bidders?
- If a house consistently sets high estimates relative to eventual realized prices, what incentive might explain that choice, and how would it show up in the pricing power figure?

This week's assignment

- Each student selects one brand or category represented in any of the production tables and writes a two-page analytical memo applying auction theory to the data.
- The memo must compute pricing power and sell-through for the chosen brand or category, state explicitly which of the three theoretical ideas (anchoring, winner's curse, reserve price and sell-through) the figures most support, and address at least one alternative explanation for the pattern observed.
- The memo must include the SQL queries used as an appendix, the exported result set, and an explicit limitations section addressing the recency caveat and the fact that realized-price data alone cannot fully distinguish anchoring from an accurate estimate.
- Grading criteria:
- correct computation of pricing power and sell-through (30 percent), soundness of the theoretical argument and consideration of an alternative explanation (35 percent), honest treatment of limitations (20 percent), and clarity of the memo (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