

ALT/FNDATA ACADEMY · MODULE 14

The auction business model

This session extends Module 7 by turning from the art market itself to the business that runs it: the auction house.

90 minutes · sandbox.altfndata.com

What today looks like

- Length: 90 minutes.
- Level: art business, auction studies, or arts administration course, and a natural extension for a business or finance elective examining intermediary and marketplace business models. Suitable for graduate students in art market or arts management programs, advanced undergraduates pairing art history or economics with business, and professional or continuing education students entering the trade.
- Access: sandbox-only, with an optional API/code extension for retrieving one house's sold lots programmatically and computing figures client-side.
- Goal: students leave able to describe the buyer's premium and seller's commission as the two revenue streams that fund an auction house, and to compute estimate accuracy, average lot value versus lot volume, and value concentration from documented fields on the fine art table.

OBJECTIVES

By the end of class students can

- Explain the auction house's role as a two-sided intermediary and describe the buyer's premium and the seller's commission as the two revenue streams that fund the business.
- Describe how a pre-sale estimate is set, why estimate accuracy matters commercially to a house, and compute the ratio of realized price to high estimate from documented fields.
- Explain guarantees and irrevocable bids as mechanisms that shift financial risk between house, consignor, and a third-party guarantor, and how each changes the incentives around a sale.
- Compute average lot value by house and contrast it with lot volume, and explain why these describe two different kinds of auction businesses.
- Measure value concentration among a house's top lots and explain why a single lot or a single consignor can dominate a season's results.
- Identify the limits of auction data for studying the auction business itself, including that it captures the secondary market only, that no field records buyer's premium, hammer price, or commission, and that the newest periods are still being ingested and should not be read as trends.

How the session runs

- 0 to 15 min: The auction house as a business. Introduce the buyer's premium and seller's commission as the two revenue streams, and explain why winning consignments, not running the sale itself, is the real competitive contest among houses.
- 15 to 25 min: Estimates, guarantees, and irrevocable bids. Explain how a pre-sale estimate is set, why the house has a commercial stake in getting it right, and how a guarantee or an irrevocable bid shifts risk before the sale even opens.
- 25 to 45 min: Sandbox orientation and guided demo, estimate accuracy. Confirm the documented fields on the fine art table and build the ratio of realized price to high estimate for a chosen house.
- 45 to 60 min: Guided demo, average lot value versus lot volume, contrasting a high-value house against a high-volume house.
- 60 to 75 min: Small-group exercise, students pick a different house and compute its value concentration among its own top lots.
- 75 to 85 min: Class discussion, groups report what estimate accuracy, average lot value, and concentration together suggest about the kind of business their house runs.
- 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 from the SQL editor dropdown.
- Open the coverage browser tab and confirm how a major auction house is written in the vendor field. Spelling matters for exact filters.
- Before running any query, walk the class through the two revenue streams conceptually: the buyer's premium, a fee the winning bidder pays on top of the hammer price, and the seller's commission, a fee the house deducts from what the consignor receives. Note that neither figure appears in the data students are about to query, only the final realized price the buyer paid.
- Run the first guided query, estimate accuracy for a chosen house, and explain that the ratio of `usd_price_decimal` to `sale_estimates_high_usd_price` tells students how often the house's own pre-sale number tracked what buyers actually paid, which is a direct read on the skill of the specialists who set that estimate.
- Explain guarantees and irrevocable bids conceptually at this point: a guarantee promises the consignor a minimum price regardless of the room's bidding, and an irrevocable bid commits a third party to buy the lot at a set price if no one else bids higher. Both remove downside risk from the consignor before the sale, in exchange for giving up some upside, and both are a house's tool for winning a prized consignment away from a competitor.
- Run the second guided query, average lot value by house, and explain that dividing total realized value by the count of sold lots separates a house that competes on a handful of very expensive lots from a house that competes on a large volume of moderately priced ones. Ask the class which model looks more resilient to a single lot failing to sell.

The core query

- Run this in the sandbox SQL editor:

```
- SELECT vendor,  
-         approx_percentile(usd_price_decimal / sale_estimates_high_usd_price, 0.5) AS estimate_accuracy,  
-         COUNT(*) AS sold_lots  
- FROM all_fine_art_data  
- WHERE status = 'sold'  
-       AND sale_estimates_high_usd_price > 0  
-       AND vendor LIKE '%Christie%'  
- GROUP BY vendor;
```

DISCUSS

Questions for the room

- The buyer's premium and the seller's commission are the two revenue streams that fund an auction house, yet neither figure appears in the transaction data. What can an outside analyst still infer about a house's revenue without ever seeing its fee schedule?
- Why does estimate accuracy matter commercially to a house, beyond simply making the specialists look good on sale night?
- A guarantee removes the consignor's downside risk before the sale. Who absorbs that risk instead, and what does the house gain by offering one?
- An irrevocable bid and a guarantee both shift risk before a sale opens. How do they differ in who is taking on the exposure, and why might a house prefer one over the other for a given lot?
- Two houses show similar total realized value in a season, but one has a much higher average lot value and the other a much higher lot count. What different businesses are these two houses actually running?

This week's assignment

- Each student selects one auction house represented in the fine art data and writes a two-page business brief on that house, as if preparing a competitive analysis for a rival saleroom's management.
- The brief must report the house's estimate accuracy, its average lot value against its lot volume, and its value concentration among its own top lots, and must argue, using those three figures together, what kind of auction business the house runs and where its consignment strategy appears to be working or under strain.
- The brief must include the SQL queries used as an appendix, at least one exported table or chart, and an explicit limitations section that states the data is secondary market only, notes that no field records buyer's premium, hammer price, commission, or guarantees so none of those can be computed, and does not draw conclusions from the most recent quarter given ongoing ingestion.
- Grading criteria:
 - correct computation of estimate accuracy with the required `sale_estimates_high_usd_price > 0` guard (25 percent), correct computation and contrast of average lot value and concentration (25 percent), quality of the business argument connecting the three figures to a coherent read of the house's strategy (30 percent), and honest treatment of the analysis's limitations (20 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