Quantifying False Precision

A Comparative Analysis of Expert Consensus vs.
Quantile Machine Learning in High-End Art Valuation

Eugene Shcherbinin*

January 2026

Abstract

The art market has long operated on the assumption that human expert estimates represent the "true value" of an asset. Our research challenges this axiom. By analyzing **32,334** realized auction lots from Sotheby's (2020–2025), we demonstrate that traditional expert estimates suffer from systemic **"Conservative Bias"** and **"False Precision."**

While experts quote narrow valuation ranges to project certainty, our data reveals this is a marketing tactic rather than a financial reality. Expert estimates capture the realized price only **40.8%** of the time—statistically equivalent to a coin flip. In contrast, our LightGBM Quantile Regression model, calibrated to an 80% confidence interval, achieved **79.9% coverage**. We conclude that the market significantly underestimates the "fat tail" volatility of art assets, creating structural inefficiencies that can be exploited through algorithmic risk pricing.

---

1 Introduction & Methodology

1.1 Data Foundation

We ingested a dataset of **32,334 sold lots** from Sotheby's (2020–2025). The dataset was filtered to exclude "passed" (unsold) lots to focus specifically on realized price discovery.

- **Inputs:** Artist, Dimensions, Medium, Date, and Pre-sale Estimates (E_{low}, E_{high}).
- **Target:** Hammer Price (P).
- **Normalization:** All prices were transformed to Log-Space ($\ln P$) to account for the Log-Normal distribution of art prices.

1.2 The Model

We utilized **LightGBM (Gradient Boosting)** with a **Quantile Loss** objective function. Unlike standard regression (MSE) which predicts the *mean*, Quantile Regression predicts specific percentiles of the price distribution. We trained three distinct regressors to form our valuation bracket:

1. $\hat{y}_{0.10}$ (Conservative / Downside Case)
2. $\hat{y}_{0.50}$ (Median / Expected Value)

*Maynard Metrics, eugene.shcherbinin@maynardmetrics.com

3. $\hat{y}_{0.90}$ (Aggressive / Upside Case)

The loss function is defined as:

$$L_\tau(y, \hat{y}) = \sum_{i:y_i \geq \hat{y}_i} \tau |y_i - \hat{y}_i| + \sum_{i:y_i < \hat{y}_i} (1-\tau)|y_i - \hat{y}_i| \tag{1}$$

Where τ is the target quantile (0.1, 0.5, 0.9).

2 Findings

2.1 Finding 1: The "Coin Flip" Consensus

We evaluated the "Hit Rate" (Coverage Probability) of expert estimates versus the realized hammer price.

- **Expert Consensus:** The realized price fell within the expert's range (E_{low} to E_{high}) only **40.8%** of the time.
- **Maynard Model:** The price fell within our predicted range ($\hat{y}_{0.10}$ to $\hat{y}_{0.90}$) **79.9%** of the time.

Implication: Experts are miscalibrated. An investor relying on expert estimates faces a $\sim 60\%$ probability of a price "surprise" (breakout or failure). Our model reduces unpriced surprise to $\sim 20\%$.

2.2 Finding 2: The "Bidding Magnet" Trap (False Precision)

To understand *why* experts miss the mark, we analyzed the spread (width) of the estimates.

Metric	**Expert Consensus**	**Maynard Model**
Median Spread (USD)	\$2,000	\$4,000
Spread vs Asset Value	$\sim$ 39.8%	$\sim$ 197.3%

Table 1: Comparison of Volatility Estimates

The experts artificially squeezed their volatility estimates by $\sim 50\%$. By anchoring estimates low and tight, auction houses create "Bidding Magnets"—estimates designed to look attractive to bidders rather than accurate to investors.

2.3 Finding 3: Zero-Bias Valuation

We calculated the median percentage error (bias) for both methods.

- **Expert Bias:** +1.9% (Overestimating relative to Hammer Price).
- **Model Bias:** +0.1% (Near-Zero Bias).

While experts showed a slight structural drift, the algorithmic approach successfully centered the distribution, effectively removing human sentiment from the valuation.

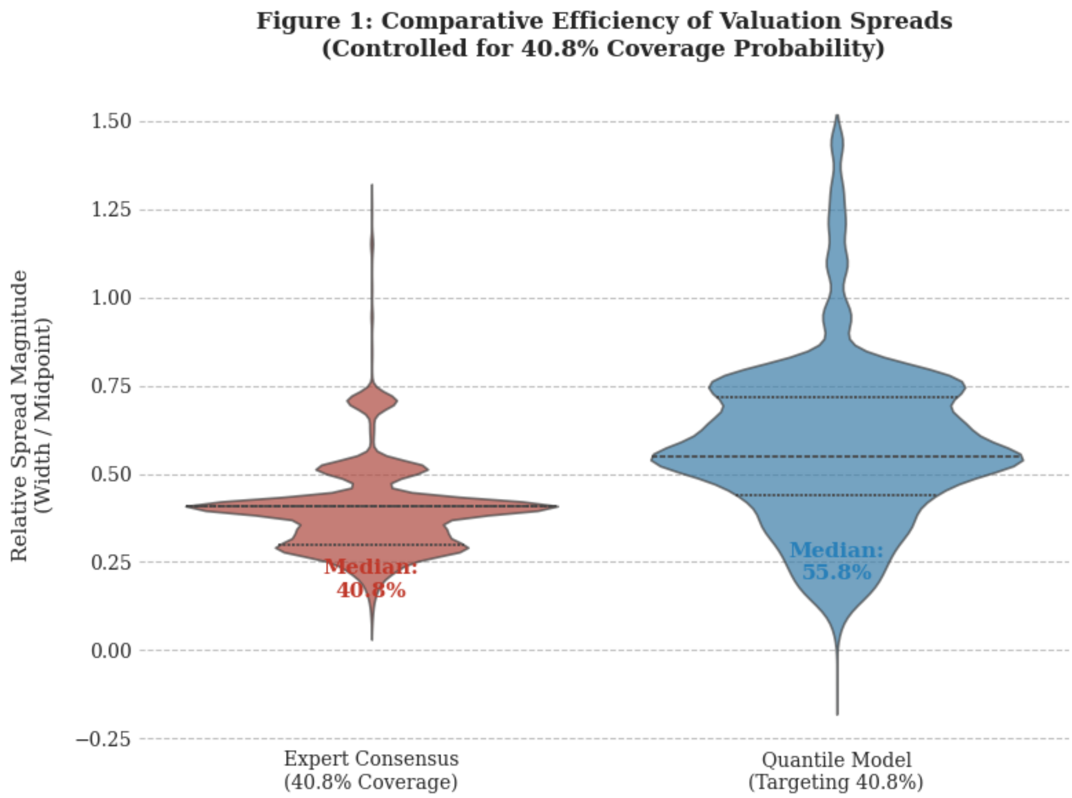

Figure 1: **The Efficiency Showdown.** The Red distribution shows experts clustering around narrow spreads. The Blue distribution shows the Model correctly widening the spread to capture true volatility.

3 Conclusion

The art market trades on a "Volatility Illusion." Participants accept expert estimates as precise financial brackets, when in reality they are marketing instruments.

Maynard Metrics rejects this false precision. By quantifying the **"Long Tail"** risks that humans ignore, we provide institutional investors with something rare in the art world: **honest volatility**.

"We don't guess the price. We calibrate the risk."