A Work Project, presented as part of the requirements for the Award of a Master's degree in Finance from the Nova School of Business and Economics.

THE VALUE OF ARTWORKS:
A QUANTITATIVE ANALYSIS OF BUYER BEHAVIOR AND EXPERT VALUATION IN AUCTIONS

MATTIA SAVIOLI

Work project carried out under the supervision of:

Miguel Lebre Freitas

17/12/2024

Abstract

This study explores the mechanisms of art valuation in auction markets, focusing on buyer behavior and expert appraisal. Using a hedonic classification framework and machine learning, artworks are analyzed based on hammer prices relative to pre-sale estimates. Key drivers, including textual descriptions, technical attributes, and market seasonality, are identified. The findings highlight the complexities of valuation, the impact of auction practices, and buyer behavior, providing a quantitative perspective on auction dynamics. This research enhances understanding of valuation processes and their implications for art markets.

Keywords: Art Valuation, Hedonic Classification Framework, Machine Learning for Auction Prediction, Buyer Behavior in Art Markets, Auction Market Dynamics, Art as a Durable Commodity

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209).

1. Introduction

Art valuation is a complex process that combines cultural significance with financial analysis. Unlike traditional financial instruments, the value of artworks is shaped by a mix of tangible attributes, such as size and medium, and intangible factors, including historical significance and collective perceptions of cultural value. This interaction makes art valuation inherently subjective and a unique area of study in financial and economic research.

The global art market, particularly auction houses, serves as the focal point for this valuation process, where expert appraisals intersect with buyer behavior to establish prices. Auctions function as performative events, shaping cultural narratives and market hierarchies of value. However, the opacity of valuation, compounded by practices like reserve pricing and selective cataloging, creates challenges in understanding the mechanisms that drive price outcomes.

This thesis investigates the fundamental question: "What drives the valuation of artworks in auction markets, and how can we understand the interaction of buyer behavior with these dynamics using modern analytical tools?" The study aims to identify the factors influencing auction outcomes, evaluate the role of expert appraisals in shaping pre-sale estimates, and examine how buyer behavior interacts with market conditions such as seasonality and auction prices. Buyer behavior is central to understanding art valuation, as auctions often highlight emotional psychological dynamics alongside financial considerations. Competitive bidding, cultural preferences, and speculative tendencies frequently drive buyers to exceed rational valuations, particularly for rare or culturally significant pieces.

To address these questions, this research applies a quantitative approach, leveraging machine learning within a hedonic classification framework to analyze data from Sotheby's auctions. By categorizing artworks based on their performance relative to pre-sale estimates, the study identifies key drivers of valuation, including textual descriptions, technical attributes, and market context. This approach provides a deeper understanding of the dynamics influencing art valuation and buyer strategies.

Section 2 explores existing research on art valuation and its multifaceted challenges. Section 3 presents the data, methodology, and development of a predictive model to classify artworks based on their performance relative to pre-sale estimates. The results are presented and interpreted in Section 4, offering insights into the drivers of valuation and the role of descriptive, technical, and contextual factors. Finally, Section 5 discusses the analysis in the context of its implications for the art market, and Section 6 draws conclusions with suggestions for future research.

By bridging cultural economics with quantitative analysis, this research contributes to a growing discourse on the value of art, offering insights into the mechanisms that shape its financial and cultural worth.

2. Literature Review

2.1. Art as an Emerging Asset Class

The art market has increasingly become a focus of both academic studies and investments (Guo, Li, and Wei 2024). This is driven in part by the search for alternative assets during financial crises (David, Oosterlinck, and Szafarz 2013). Art’s unique position as a cultural and financial asset distinguishes it from conventional markets. Unlike traditional financial instruments, art derives its worth not only from tangible attributes such as size or medium but also from its intrinsic cultural significance, as acknowledged by a collective, and its extrinsic financial potential within the market (Lamont 2012). It provides a mix of financial returns and "psychic returns," the latter being the pleasure derived from owning and experiencing artworks, which often outweigh purely monetary considerations in the behavior of collectors (Frey and Eichenberger 1995). This blend of cultural and financial value highlights art’s dual appeal as both a consumer good and an investment opportunity. While art's financial appeal is undeniable, an overemphasis on marketability risks sidelining culturally significant works that do not align with prevailing trends, reducing the diversity of represented artists and styles.

Art is recognized not just as a vehicle for wealth preservation or speculation but also to hedge against inflation and diversify investment portfolios (Campos and Barbosa 2008). This has been particularly evident during times of economic uncertainty when investors seek alternative assets to stabilize their portfolios.

Furthermore, art markets exhibit a mix of public and private participation. Public institutions, such as museums, acquire art as part of a broader socio-cultural mandate, contributing to the preservation and celebration of cultural heritage. On the other hand, private investors navigate their dynamic valuation influenced by factors like rarity, provenance, and aesthetic appeal, which collectively shape the demand and pricing dynamics of the market (Frey and Eichenberger 1995). Moreover, fundraising for acquisitions by public institutions often increases during periods of economic upturn and for high-profile contemporary artists (Singer and Lynch 1994).

The behavioral tendencies of art collectors further enhance its appeal as an asset class. Unlike strict financial investments, art's valuation is often influenced by emotional attachment and cultural recognition. Behavioral patterns, such as the endowment effect, where owners place a higher value on their possessions, and the sunk cost effect, where personal attachment intensifies perceived value, reveal the unique psychological dynamics of art buyers (Frey and Eichenberger 1995). These patterns reinforce art's role as a socially constructed asset where collective acknowledgment of cultural significance plays a critical role in shaping its market value (Plante, Free, and Andon 2021).

The role of the state adds another layer of complexity to this evolving market. Governments contribute to art valuation through funding, governance, and regulatory frameworks, shaping the art market while reinforcing art's position as a socially constructed and culturally significant asset (Cassese and Casini 2012; Crepaz, Huber, and Scheytt 2016). State interventions, combined with private and institutional participation, underscore the intricate chemistry of cultural and economic factors that influence valuation and demand.

These factors collectively highlight the evolving significance of art as a multifaceted asset class, offering rich potential for economic and cultural analysis. This duality, where personal passion intersects with public recognition, continues to position art as a compelling subject for academic inquiry and a dynamic vehicle for investment.

2.2. The Challenges of Art Valuation

Valuing art is an inherently subjective and intricate process that challenges even the most experienced experts. Unlike traditional financial assets, the worth of an artwork depends not only on tangible factors, such as size, medium, or condition, but also on intangible and often abstract elements like cultural significance, historical context, and the reputation of the artist (Agnello 2002; Edwards 2004). This duality makes it difficult to establish consistent criteria or methodologies for valuation. Art experts, who are often entrusted with appraising value, face significant hurdles in predicting an artwork's performance in the marketplace. Research shows that experts frequently struggle to forecast whether a piece will sell at all, let alone estimate its final hammer price with accuracy (Campos and Barbosa 2008). The inability to reliable predict even fundamental outcomes underscores the complexity of art valuation.

The art world's reliance on multi-factor valuation compounds these difficulties. Experts must assess factors ranging from provenance, rarity, and condition to less quantifiable attributes like the artwork's cultural relevance or its ability to evoke emotional resonance in buyers. In this context, historical prices, while useful, are often of limited relevance due to the unique nature of each artwork and shifts in market demand over time (Coslor 2016). Furthermore, the lack of transparency in valuation processes complicates matters. Auction prices are among the few publicly accessible benchmarks, yet these data points fail to account for the broader, nuanced determinants of value. This over-reliance on thin, auction-focused datasets can lead to distorted estimates (Coslor 2016).

Art critics experts also contend with challenges to their credibility. Past valuation failures have significantly undermined trust in their assessments. For example, works once dismissed as insignificant have later achieved remarkable critical and financial success, raising questions about the reliability of expert opinion (Plattner 2000). The subjective nature of art itself further complicates these judgements, as experts must balance aesthetic interpretations with market-oriented appraisals.

Art valuation remains a fraught process where cultural and financial considerations intersect in unpredictable ways. The necessity of combining qualitative judgements with quantitative benchmarks places a heavy burden on experts, who must navigate the fine line between subjective interpretation and objective appraisal. To address these challenges, the art world must push for greater transparency and accountability, enabling experts to better align their valuations with the intricate realities of the art market.

2.3. How Art is Valued

Art valuation is a complex social construct rooted in cultural norms and conventions, structured through three core modes: interpreting, credentialing, and projecting. Interpreting involves assessing an artwork's cultural and historical significance, such as its aesthetic style, provenance, or connection to broader narratives. Credentialing focuses on establishing the reputation of an artist or artwork, often through endorsements like gallery representation, awards, or institutional exhibitions. Projecting combines qualitative and quantitative methods, such as auction records and price trends, to estimate an artwork's future monetary value (Plante, Free, and Andon 2021).

Valuation is not a neutral or passive process; it actively participates in the creation of value. This performative aspect is evident in the strategies employed by auction houses, with a particular emphasis on serving the priorities of sellers. In many cases, the valuation process prioritizes these interests over the intrinsic value of the artwork itself, highlighting how market dynamics often rely on strategic considerations rather than purely cultural or aesthetic judgements (Plante, Free, and Andon 2021).

The process remains opaque, relying heavily on the nuanced, contextual expertise of art professionals, which cannot be fully replicated by algorithmic models. Human judgement is essential in weighing subjective factors like aesthetic appeal, rarity, and cultural resonance, all of which are critical to determining value (Plante, Free and Andon 2021).

In sum, art valuation is a performative and strategic process driven by social practices, expert judgment, and market behavior. This chemistry continuously shapes how value is constructed in the art world, reflecting the intricate relationship between cultural significance and financial worth (Plante, Free, and Andon 2021).

2.4. Auction Houses and Their Influence on Valuation

Auction houses are pivotal in shaping the perceived value of artworks, primarily by acting as intermediaries between sellers and buyers and structuring how consumer preferences translate into market prices (Campos and Barbosa 2008). Through catalogs, they provide detailed descriptions of artworks, which are not only informative but also serve as marketing tools to create a narrative around the pieces and strategically set estimates to influence buyer behavior and maximize hammer prices (Plante, Free, and Andon 2021). Undisclosed reserve prices underscore the auction houses' prioritization of sellers' interests, often taking precedence over ensuring a successful sale. Typically, the reserve price is set at 80% of the minimum estimate provided by experts, serving as a safeguard to prevent artworks from being undervalued during the auction process. If bids fall below this reserve, the artwork is "bought-in," remaining unsold to protect its perceived value and prevent reputational damage to the artist or to the seller (Ashenfelter and Graddy 2003).

The auction system, however, is not without its quirks and inefficiencies. The declining price anomaly or "afternoon effect" reflects the psychological fatigue of bidders, resulting in lower prices for lots sold later in the session (Beggs and Graddy 1997). Furthermore, private buyers frequently drive-up prices during auctions, creating valuation bubbles that do not always align with the intrinsic value of the artwork (Coslor 2016). While auction prices remain a critical input in valuation practices, they are only one component of the broader valuation landscape, which incorporates both quantitative data and qualitative judgments by experts (Karpik 2010).

Auctions exemplify the unpredictability of the art market, where the final hammer price often diverges from any pre-sale valuation or true intrinsic value. This underscores the performative nature of valuation, where strategic positioning and market dynamics significantly influence outcomes beyond mere assessments of cultural or aesthetic worth. The reliance on auction results as benchmarks illustrates the complex interaction between market mechanisms, consumer psychology, and expert valuations in shaping the art world's financial narratives.

2.5. Limitations of Repeat-Sales Analysis

Repeat-sales analysis, a commonly used methodology in art valuation, faces several significant limitations. High transaction costs, which can vary significantly across different auctions for the same artworks, make calculating returns using repeat-sales methodologies potentially misleading (Agnello 2002). Additionally, selection bias arises because unsold artworks are excluded from analyses, leading to a skewed representation of the market (Edwards 2004). Moreover, the dependence of art prices on political and administrative interventions, such as government-imposed restrictions, further complicates the study of average returns on art investments. These factors redirect analytical focus from returns to understanding buyer behavior and the institutional determinants of art market dynamics, such as museum organization and public administration policies (Frey and Eichenberger 1995).

The uniqueness of artworks and the limited availability of comparable reference data have led some scholars to criticize the repeat-sales approach altogether. Many pieces are traded only once or twice in a decade, complicating the calculation of risk and volatility for investment purposes (Collier 2005). This critique emphasizes that the repeat-sales method often oversimplifies the complex and multi-dimensional nature of art valuation, neglecting qualitative factors such as cultural significance or aesthetic appeal, which are pivotal in determining value (Coslor 2016).

As an alternative, emerging valuation frameworks such as hedonic pricing models emphasize the inclusion of a broader set of variables, including artist reputation, work characteristics, and sales context, to address the limitations of traditional repeat-sales analysis (Guo, Li, and Wei 2024). These approaches aim to reduce biases and provide a more holistic understanding of art as a socially constructed and market-driven commodity.

This study builds upon the previously cited literature by addressing some of the limitations observed in traditional art valuation methods. While previous research highlights the challenges of capturing the multi-dimensional nature of art prices, particularly with approaches like repeat-sales analysis, this study explores a hedonic pricing perspective to incorporate both intrinsic and extrinsic factors. By focusing on deviations from pre-sale estimates, a key benchmark in auction dynamics, it provides an alternative lens for understanding how various attributes influence auction outcomes.

At the same time, this work aligns with existing theoretical perspectives on the role of descriptive, technical, and contextual features in art market behavior. By applying a classification approach, it examines auction results in a structured manner, offering a data-driven exploration that aims to complement prior research. In doing so, this work takes a step toward integrating insights from valuation theory with practical analysis, offering a framework that can be adapted for further exploration of art market behavior.

3. Model Development

3.1. Choice of the Model

The XGBoost (Extreme Gradient Boosting) model was selected due to its ability to efficiently handle large datasets, multi-class classification tasks, and complex relationships between features. XGBoost builds on the gradient boosting algorithm by incorporating enhancements such as advanced regularization and optimized computational speed. These attributes make it ideal for predicting whether an artwork's final hammer price falls below, within, or above the estimated range.

The model was implemented within a pipeline that integrated Recursive Feature Elimination (RFE) for feature selection. RFE systematically removes less important features, enhancing model interpretability and efficiency selecting only the top ten most impactful features.

The XGBoost classifier then leveraged gradient-boosting decision trees to classify artworks into the three price categories. Gradient boosting works by building trees sequentially, each correcting the errors of its predecessor. The final model aggregates the predictions of all trees to make robust classification.

3.2. Hedonic Model Setup

The hedonic pricing framework underpins this analysis, modeling artwork prices as a function of intrinsic and extrinsic attributes. Intrinsic attributes include features like the artwork's dimensions and medium, while extrinsic attributes involve contextual factors such as the artist's reputation and sale timing. The target variable Y, representing price categories, is modeled as

$$\hat{Y} = f(X; \Theta) \qquad [1]$$

where $\hat{Y}$ is the predicted category, X represents the feature matrix, Θ are the model parameters, and f is the ensemble of decision trees. XGBoost minimizes the multi-class logarithmic loss:

$$\mathcal{L}(\Theta) = -\frac{1}{N} \sum_{i=1}^{N} \sum_{k=1}^{K} y_{ik} \log(\hat{\rho}_{ik}) \qquad [2]$$

where N is the total number of observations, $K = 3$ is the number of classes, y_{ik} is a binary indicator of whether artwork i being in class k, and $\hat{\rho}_{ik}$ is the predicted probability of artwork i being in class k. Regularization prevents overfitting and ensures better generalization, expressed as:

$$\Omega(f_t) = \gamma T + \frac{1}{2} \lambda \sum_{j=1}^{T} w_j^2 \qquad [3]$$

where T is the number of leaves, w_j is the weight of leaf j, γ controls the number of leaves and λ applies an L2 regularization penalty on leaf weights.

3.3. Parameter Tuning

Hyperparameter optimization was conducted using a randomized search approach to identify the best configuration of the model's parameters. Key parameters were adjusted to enhance performance and balance model complexity:

- The number of boosting rounds was set to 300 rounds, ensuring the model captures sufficient patterns without overfitting.

- The maximum depth of each tree was optimized to 10 steps, controlling the complexity of individual trees, and balancing the model's ability to learn detailed relationships with the risk of overfitting.
- The learning rate was set to 0.1, scaling the contribution of each tree during training to achieve a balance between faster convergence and stable learning.
- To introduce diversity and prevent overfitting, 80% of the features were sampled for each tree.
- All training data were used in each boosting round, maximizing the model's generalization ability by incorporating a comprehensive range of data points in the learning process.

3.4. Cross-Validation

Cross-validation is a technique used to evaluate the generalization ability and robustness of a model by splitting the dataset into multiple subsets (folds) for training and testing. In this study, 10-fold cross-validation was employed, where the dataset was divided into ten equally sized parts. Each fold was used as a test set while the remaining nine were used for training, ensuring that every instance in the dataset was tested exactly once. This iterative process provides a comprehensive evaluation of the model's performance across different data splits and minimizes the risk of overfitting.

The primary metric for evaluation was the weighted F1 score, calculated as follows:

$$F1_k = \frac{2\, Precision_k\, Recall_k}{Precision_k + Recall_k} \quad [4]$$

$$F1_{weighted} = \sum_{k=1}^{K} w_k F1_k \quad [5]$$

Here, $F1_k$ combines precision and recall for each class *k*. Precision is the fraction of true positives among all predicted positives for a given class, while recall is the fraction of true positives correctly identified out of all actual positives for the class:

$$Precision_k = \frac{True\ Positiveses_k}{True\ Positives_k + False\ Positives_k} \quad [6]$$

$$Recall_k = \frac{True\ Positiveses_k}{True\ Positives_k + False\ Negatives_k} \quad [7]$$

The weighted F1 score averages the F1 scores of all classes, adjusting for their relative proportions (w_k) in the dataset.

3.5. Data and Methodology

The development of the predictive model required a robust and carefully prepared dataset to ensure reliable performance and accurate analysis. The dataset used in this study consisted of 4,126 observations and 29 variables, capturing transactions of artworks sold at Sotheby's over the past five years. This dataset focuses on artists whose total sales across platforms exceeded $10 million in 2023. It was collected through a custom Python web scraping algorithm and cleaned to ensure quality, with no missing values or inconsistencies.

The preprocessing phase began with the creation of the target variable, "Price Category," by comparing each artwork's final hammer price to Sotheby's low and high pre-sale estimates. This categorization formed three distinct classes: "Below," "Within," and "Above." The multi-class classification problem provided a nuanced framework for understanding how artworks perform relative to auction house expectations.

An early challenge in this phase was addressing class imbalance, as the "Below" category was significantly underrepresented compared to the "Within" and "Above" classes (*Figure 1*). To tackle this, the Adaptive Synthetic Sampling (ADASYN) method was employed, generating synthetic data points by interpolating existing instances in the minority class. This process ensured that the dataset was balanced while maintaining its representativeness, thereby improving the model's ability to learn equitably across all categories (*Figure 2*).

Feature engineering played a key role in enhancing the dataset's utility. Textual descriptions were target-encoded to capture their impact on price variations. Categorical variables, such as medium and surface, were one-hot encoded to preserve interpretability, while temporal features, including auction season, were retained to capture time-based patterns. Continuous variables like dimensions and year of creation were standardized to ensure comparability and reduce skewness. Robust scaling techniques were used to minimize the influence of outliers while retaining their contextual significance in valuation.

Figure 1. Class Imbalance before ADASYN

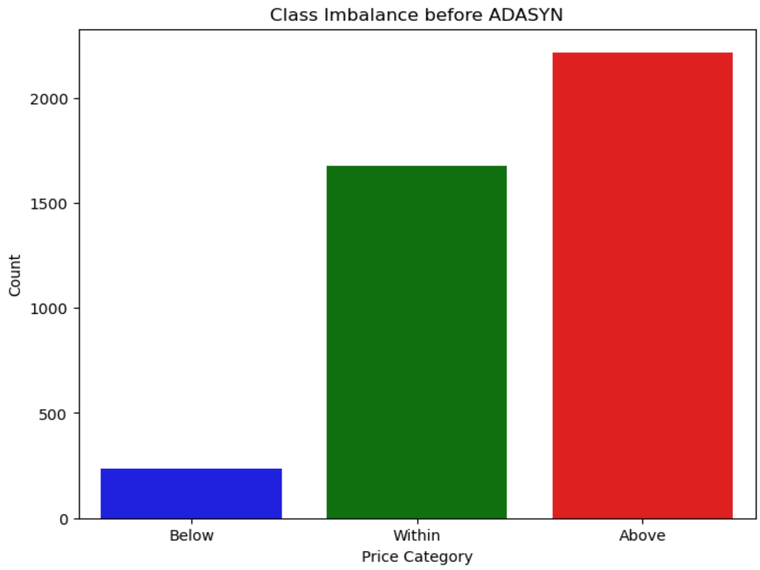

Figure 2. Class Imbalance after ADASYN

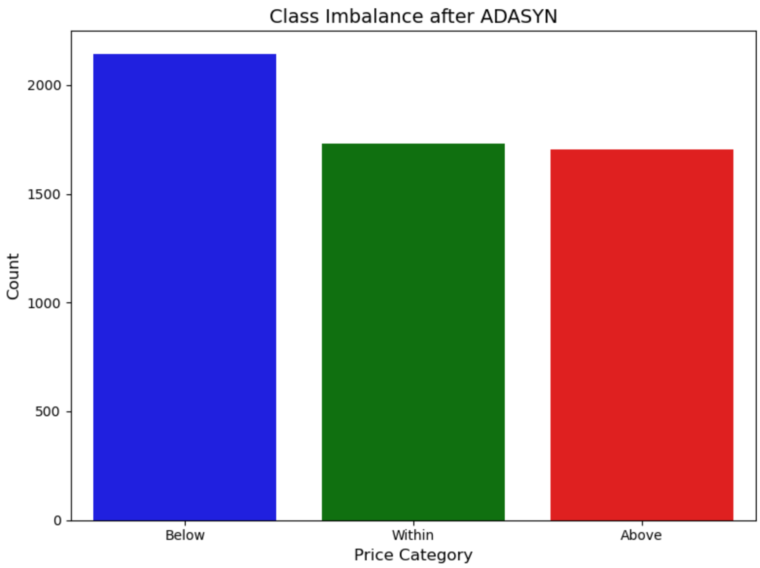

The final dataset comprised 89 features (later reduced with RFE), representing a comprehensive mix of numeric, categorical, temporal, and text-based attributes. These features provided the foundation for predictive modeling, ensuring that the dataset balanced interpretability, diversity, and reliability to meet the demands of a sophisticated hedonic pricing framework.

4. Results

4.1. Model Performance

The performance of the final model was evaluated using the results summarized in *Table 1* (Overall Model Performance Metrics) and *Table 2* (Class-Specific Model Performance Metrics). The overall accuracy of the model was 94.22%, meaning that the model correctly classified most instances. Key metrics such as weighted precision (94.32%), weighted recall (94.22%), and weighted F1 score (94.25%) further confirm the model's robust performance across all classes, indicating that the strategy implemented to address class imbalance effectively contributed to accurate predictions.

The breakdown of class-specific performance metrics highlights the model's strong and equitable performance across all categories. The "Below" category demonstrated exceptional performance, with an F1 score of 96.51%, reflecting the model's ability to identify instances accurately and consistently in this class. Similarly, the "Within" and "Above" categories achieved strong F1 scores of 94.81% and 90.82%, respectively. These results underscore the model's effectiveness in overseeing all target classes, maintaining high predictive accuracy despite class imbalance.

Table 1: Overall Model Performance Metrics

Metric	Value
Accuracy	94.22%
Weighted Precision	94.32%
Weighted Recall	94.22%
Weighted F1 Score	94.25%

Table 2: Class-Specific Model Performance Metrics

Metric	Below	Within	Above
Precision	97.26%	95.55%	89.37%
Recall	95.80%	94.10%	92.37%
F1 Score	96.51%	94.81%	90.82%

4.2. Feature Importance

The feature importance table (*Table 3*) highlights the key drivers of the model's predictions. Among the ten selected features, "Original Currency USD" is the most impactful, contributing 25.4% to the model's decisions. This reflects how artworks listed in USD, typically by the American market, affect buyer perceptions of estimate accuracy, potentially influencing whether they view the pre-sale estimates as less, more, or accurately aligned with the expected hammer price. Seasonal variables, such as "Season Spring" (12.7%) and "Season Summer," highlight the importance of timing in the auction context. Artistic techniques also play a prominent role, with "Techniques Gouache" (10.2%), "Techniques Acrylic," "Techniques Pencil," and "Techniques Pastel" collectively providing nuanced insights into material preferences. Less influence on the prediction is given by the artwork's surface, with only "Surfaces Board" (8.6%) included in the model. Additionally, "Price

Estimated" and "Full Description Encoded" (both 8.6%) display the integration of diverse data types, from continuous numerical estimates to text-encoded descriptions. This combination of categorical, continuous, and encoded variables demonstrates the comprehensive approach employed by the model to capture the multifaceted nature of the auction domain.

Table 3. Feature Importance Table

Features	**Contribution to Model's Predictions**
Original Currency USD	25.4%
Season Spring	12.7%
Techniques Gouache	10.2%
Price Estimated	8.6%
Full Description Encoded	8.6%
Surfaces Board	8.6%
Season Summer	8.4%
Techniques Acrylic	6.8%
Techniques Pencil	5.5%
Techniques Pastel	5.2%

4.3. Interpretation of Results

The delta probability plot (*Figure A1*) for binary features displays how the absence or presence of these variables impacts the likelihood of each class ("Below," "Within," "Above"). For instance, "Original Currency USD" dramatically increases the probability of the "Within" class when present while reducing the likelihood of "Above." Conversely, all the artworks' proper features (techniques and surfaces) significantly boost the probability of

“Above” and reduce the chances of “Within.” This form is seen again for the seasonal variables, but with less influence. These patterns illustrate suggesting how technical information about the artwork in its description can push the buyer to spend more than expected, and how American auctions tend to be more precise in these hammer price estimates.

The partial dependence plots (PDPs) and binned probabilities plot for the features “Price Estimated” and “Full Description Encoded” reveal complementary insights into their roles in influencing the model’s predictions. While PDPs illustrate the isolated effect of these features by averaging over all other predictors, binned probability plots reflect their combined interactions with the full feature set.

For “Price Estimated,” the PDPs (*Figure A2*) show that lower values strongly favor the “Below” class, with gradual decrease in probability as the values increase. Concurrently, the probability of the “Within” class rises steadily with higher values, while the “Above” class does not impact particularly the final prediction when considered in isolation. However, the binned probabilities plot (*Figure A3*) suggests a more nuanced relationship. For higher “Price Estimated” values, the probability of the “Above” class increases, reflecting interactions with other features that amplify its effect in these cases. This highlights the importance of “Price Estimated” in distinguishing between “Above” and “Within” predictions in the context of higher values, as the model leverages additional feature interactions to refine these predictions.

Similarly, for “Full Description Encoded,” the PDPs (*Figure A4*) reveal that lower encoded values significantly favor the “Below” class, with a steep drop in probability as the values increase. The “Within” class, in contrast, sees a sharp rise in probability with increasing values, while the “Above” class shows minimal sensitivity to this feature when considered in isolation. The binned probabilities plot (*Figure A5*), however, paints a different picture for the

“Above” class, indicating that higher values of “Full Description Encoded” are often associated with increased probabilities for “Above.” This discrepancy underscores the impact of feature interactions: while the isolated contribution of the feature may be minor for “Above,” its combined influence with other features drives these predictions.

4.4. Insights and Implications

These interpretations provide valuable insights into the dynamics of auction outcomes and the factors influencing them. The results suggest that including detailed technical information about the artwork (such as techniques, surfaces, and even seasonal influences) can significantly impact the hammer price, often pushing it above the estimated range. This aligns with the notion that well-described and technically supported artworks are more likely to attract attention and justify higher bids.

The analysis also reveals that estimates provided in USD tend to be more precise, with fewer deviations observed in the final outcomes. This may reflect the currency's role as a global benchmark, reducing variability and enhancing predictability for international buyers. Interestingly, higher price estimates correlate with increasingly unpredictable outcomes, as evidenced by the diverging probabilities for the “Within” and “Above” classes in both the partial dependence and binned plots. This suggests that as the estimated value grows, the auction results are less constrained by the estimate, reflecting the influence of collectors' willingness to compete for high-value pieces.

Overall, the findings highlight how a combination of precise pricing, detailed descriptions, and contextual factors like currency and seasonal timing contribute to shaping buyer behavior and auction outcomes. By leveraging these insights, auction houses and sellers can better tailor their strategies to optimize results.

5. Discussion

5.1. Analysis in the Context of Prior Findings

The findings of this analysis provide nuanced confirmations and extensions of the themes explored in the Introduction and Literature Review. The inherent difficulty in art valuation, attributed to the complex interaction of subjective and objective factors, was reflected in the model's reliance on diverse data types, ranging from text-encoded descriptions to numerical estimates. The imbalanced dataset, with a predominance of higher-than-expected sale prices, aligns with previous observations about the art market's tendency to favor higher-value sales, while unsold items often go unreported, creating a systemic bias in available data. This imbalance mirrors auction house practices like buy-in strategy, which prioritize the seller's reserve price over market inclusivity, thereby limiting the dataset's representation of the "Below" class.

The study also highlights how features associated with material attributes significantly impact auction outcomes, demonstrating that technical information can increase perceived value and push hammer prices above estimates. These findings align with prior assertions about art's dual appeal as both a financial investment and a source of "psychic" returns, where collectors derive satisfaction and emotional fulfillment from owning culturally significant pieces.

The influence of auction houses on valuation was also underscored by the model's ability to predict auction outcomes with high accuracy. By incorporating variables like auction season and currency denomination, the analysis reflected the strategic considerations highlighted in the Literature Review, where auction houses leverage timing and marketing tools to shape buyer behavior. This reinforces the notion that valuation is not purely an appraisal of intrinsic worth, but a performative act shaped by institutional and market dynamics.

5.2. Behavioral Finance Insights for the Art Market

The analysis sheds light on behavioral finance dynamics in the art market, particularly how buyer preferences and market strategies intersect to influence outcomes. The prominence of features related to descriptive and technical attributes suggests that buyers place substantial weight on credentialing practices, which serve to authenticate and elevate the perceived value of artworks. This underscores the social construction of value in the art market, where reputation and narrative often outweigh intrinsic qualities.

The probabilistic patterns observed for higher price categories suggest a degree of speculative behavior, where buyers are willing to exceed estimates for high-value items, likely driven by competitive auction settings. This behavior reflects broader economic incentives, where cultural capital and investment potential intertwine, motivating buyers to act beyond purely financial considerations. Similarly, the variability in outcomes for higher price estimates highlights the unpredictability of buyer psychology, where willingness to bid aggressively can disrupt expected valuations.

Overall, the findings reveal that valuation methods, while structured, are influenced by buyer behavior and perceptions, particularly in how descriptive and technical attributes shape decision-making. These insights highlight the importance of understanding buyer engagement and market dynamics to interpret auction outcomes more accurately.

6. Conclusion and Future Research Directions

This study has explored the intricate dynamics of art valuation, shedding light on the challenges, methodologies, and behavioral factors that underpin this unique market. Through an in-depth analysis of Sotheby's auction data and a robust predictive modeling framework, the study contributes to the discourse on art as both a cultural and financial asset, offering valuable insights into the factors influencing artwork prices.

The findings underscore the complexity of art valuation, affirming its dependence on both quantifiable features and subjective cultural constructs. Key elements, such as the medium, surface, and descriptive attributes of artworks, emerged as significant drivers of price outcomes, validating the multi-faceted nature of valuation discussed in earlier chapters. The incorporation of textual and categorical features highlighted the importance of expert narratives and buyer perceptions, demonstrating how encoded cultural significance translates into financial value. Additionally, addressing class imbalance using Adaptive Synthetic Sampling (ADASYN) illuminated the challenges inherent in modeling art prices. The overrepresentation of artworks sold within or above estimated ranges points to institutional practices, such as buy-in strategies, which filter out lower-value transactions from public view. This reinforces the influential role of auction houses as gatekeepers shaping valuation outcomes through strategic cataloging, pricing, and buyer engagement.

Behavioral tendencies, such as speculative fervor surrounding trending artists or emotional premiums associated with well-documented artworks, were also evident in the analysis. These findings align with behavioral finance theories, highlighting the interchange between cultural appreciation and market speculation. Credentialing practices, which enhance an artwork's perceived value through associations with renowned artists, mediums, or provenances, emerged as pivotal factors driving prices beyond pre-sale estimates.

Despite its contributions, this work acknowledges the limitations inherent in current valuation models and datasets. These gaps offer promising directions for future research that could address critical shortcomings and provide a deeper understanding of the art market's complexities.

One key area for future exploration is the Veblen effect, where rising prices and sales volumes increase the desirability and perceived value of artworks. This phenomenon, common in luxury markets, holds relevance in art, where exclusivity and social status influence demand. High-profile sales of iconic pieces often create a feedback loop, enhancing perceived value as buyers compete not only for the artwork itself but also for the prestige associated with its acquisition (Campos and Barbosa 2008). Quantifying this relationship between social status, prestige, and price escalation could provide further insights into cultural and financial drivers in the art market.

Another critical direction involves addressing the limitations of the data used in this study. While the analysis relied on Sotheby's auction data over the last five years, the exclusive focus on a single auction house limits generalizability. Future research should incorporate private sales and data from other auction houses to offer a more comprehensive understanding of market dynamics. This would ensure that findings reflect a broader spectrum of valuation practices and buyer behavior.

A further significant gap lies in the exclusion of unsold or bought-in artworks from existing valuation models. These works, which fail to meet their reserve prices, reflect important market dynamics but are often omitted from transaction datasets, leading to endogeneity biases. Incorporating these unsold pieces using econometric approaches like the Heckman correction model could provide a more accurate understanding of market behavior (De Ridder, Eriksen, and Scholtens 2024). By simultaneously modeling the likelihood of an artwork being sold and its price, this approach accounts for selection bias and captures the influence of market constraints, such as reserve pricing strategies.

Another promising area for research is the previously cited afternoon effect, where auction prices tend to decline for lots sold later in the session. This anomaly challenges the assumption of consistent pricing mechanisms and highlights the role of temporal dynamics such as bidder fatigue, liquidity constraints, and strategic bidding behaviors (Beggs and Graddy 1997). Analyzing these temporal patterns could yield insights into how auction sequencing and buyer behavior impact outcomes, enabling auction houses to optimize their strategies.

Finally, the integration of emerging technologies such as blockchain and non-fungible tokens (NFTs) presents another compelling avenue for exploration. Blockchain enhances transparency in provenance tracking and valuation, while NFTs are redefining the art market by enabling digital ownership and new forms of speculative behavior. Understanding how these technologies disrupt traditional valuation models and influence buyer confidence will be essential for stakeholders navigating this evolving landscape.

In conclusion, this study provides a nuanced understanding of the mechanisms driving artwork prices, bridging the cultural and financial realms of the art market. By identifying the roles of descriptive, technical, and contextual features, as well as highlighting buyer behaviors and auction house practices, the findings contribute to both academic discourse and practical applications. Auction houses can leverage these insights to refine their cataloging and marketing strategies, while policymakers may explore initiatives to enhance transparency and reduce market distortions caused by speculation.

Future research, addressing the gaps outlined above, holds the potential to advance valuation methodologies further. By integrating temporal dynamics, unsold works, and emerging technologies, new studies can enhance predictive accuracy and deepen our understanding of the art market. As art continues to evolve as a cultural and financial asset, its unique blend of subjectivity and quantifiable metrics remains fertile ground for further inquiry, offering valuable insights for collectors, investors, and researchers alike.

References

Agnello, Richard J. 2002. “Investment Returns and Risk for Art: Evidence from Auctions of American Paintings.” *Eastern Economic Journal* 28, no. 4: 443–63.

Ashenfelter, Orley, and Kathryn Graddy. “Auctions and the Price of Art.” *Journal of Economic Literature* 41, no. 3 (2003): 763–87.

Beggs, Alan, and Kathryn Graddy. “Declining Values and the Afternoon Effect: Evidence from Art Auctions.” *The RAND Journal of Economics* 28, no. 3 (1997): 544–65.

Bryan, Michael F. 1985. "Beauty and the Bulls: The Investment Characteristics of Paintings." *Economic Review*. Federal Reserve Bank of Cleveland, issue QI: 2–10.

Campos, Nauro F., and Renata Leite Barbosa. 2009. “Paintings and Numbers: An Econometric Investigation of Sales Rates, Prices, and Returns in Latin American Art Auctions.” *Oxford Economic Papers* 61 (1): 28–51.

Cassese, Sabino, and Lorenzo Casini. 2012. "Public Regulation of Global Indicators." In *Governance by Indicators: Global Power through Quantification and Rankings*, Oxford University Press.

Collier, C. 2005. “Artful misbehavior: An application of behavioral finance to the art market.” *Sotheby's Institute of Art, London*.

Coslor, Erica. 2016. “Transparency in an Opaque Market: Evaluative Frictions between ‘Thick’ Valuation and ‘Thin’ Price Data in the Art Market.” *Accounting, Organizations and Society* 50 (April): 13–26.

Crepaz, Lukas, Christian Huber, and Tobias Scheytt. 2016. “Governing Arts through Valuation: The Role of the State as Network Actor in the European Capital of Culture 2010.” *Critical Perspectives on Accounting* 37 (June): 35–50.

David, Géraldine, Kim Oosterlinck, and Ariane Szafarz. 2013. "Art Market Inefficiency." *Economics Letters* 121 (1): 23–25.

De Ridder, Adri, Steffen Eriksen, and Bert Scholtens. 2024. “The Art of Valuation: Using Visual Analysis to Price Classical Paintings by Swedish Masters.” Edited by Simon Porcher. *PLOS ONE* 19 (1).

Edwards, Sebastian. “Presidential Address: The Economics of Latin American Art: Creativity Patterns and Rates of Return.” *Economía* 4, no. 2 (2004): 1–35.

Frey, Bruno S., and Reiner Eichenberger. 1995. “On the Rate of Return in the Art Market: Survey and Evaluation.” *European Economic Review* 39 (3–4): 528–37.

Guo, Mingjun, Xuerong Li, and Yunjie Wei. 2024. “Bibliometric Analysis of the Art Market: From Art Price to Market Efficiency.” *Data Science and Management* 7 (4): 349–60.

Karpik, Lucien. 2010. “Valuing the unique: The economics of singularities.” Princeton University Press.

Lamont, Michèle. “Toward a Comparative Sociology of Valuation and Evaluation.” *Annual Review of Sociology* 38 (2012): 201–21.

Plante, Maude, Clinton Free, and Paul Andon. 2021. “Making Artworks Valuable: Categorisation and Modes of Valuation Work.” *Accounting, Organizations and Society* 91 (May): 101155.

Plattner, Stuart. 2000. "Profit Markets and Art Markets." In *Commodities and Globalization: Anthropological Perspectives*: 113–134.

Singer, Leslie, and Gary Lynch. “Public Choice in the Tertiary Art Market.” *Journal of Cultural Economics* 18, no. 3 (1994): 199–216.

Appendix

Figure A1. Impact of Binary Features (Delta Probabilities for Classes)

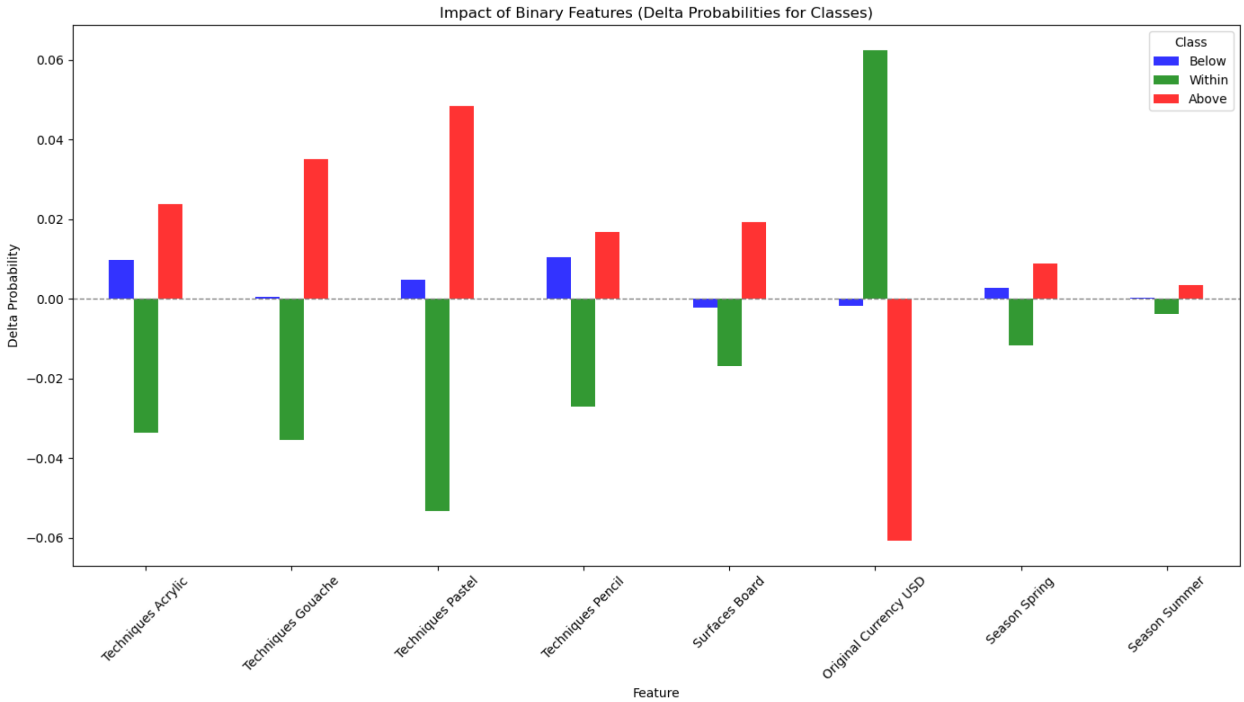

Figure A2. Partial Dependence Plots for 'Price Estimated'

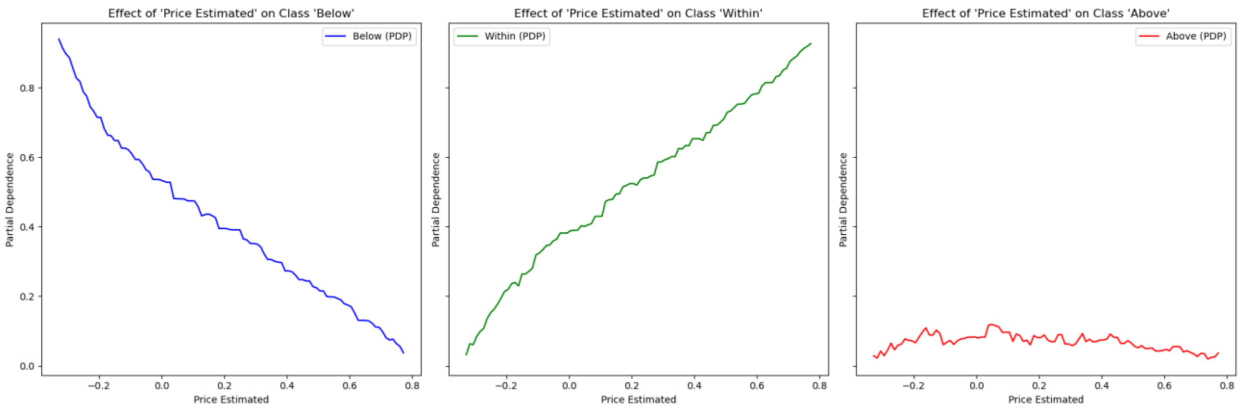

Figure A3. Class-Wise Average Predicted Probabilities by 'Price Estimated' (Binned Probabilities)

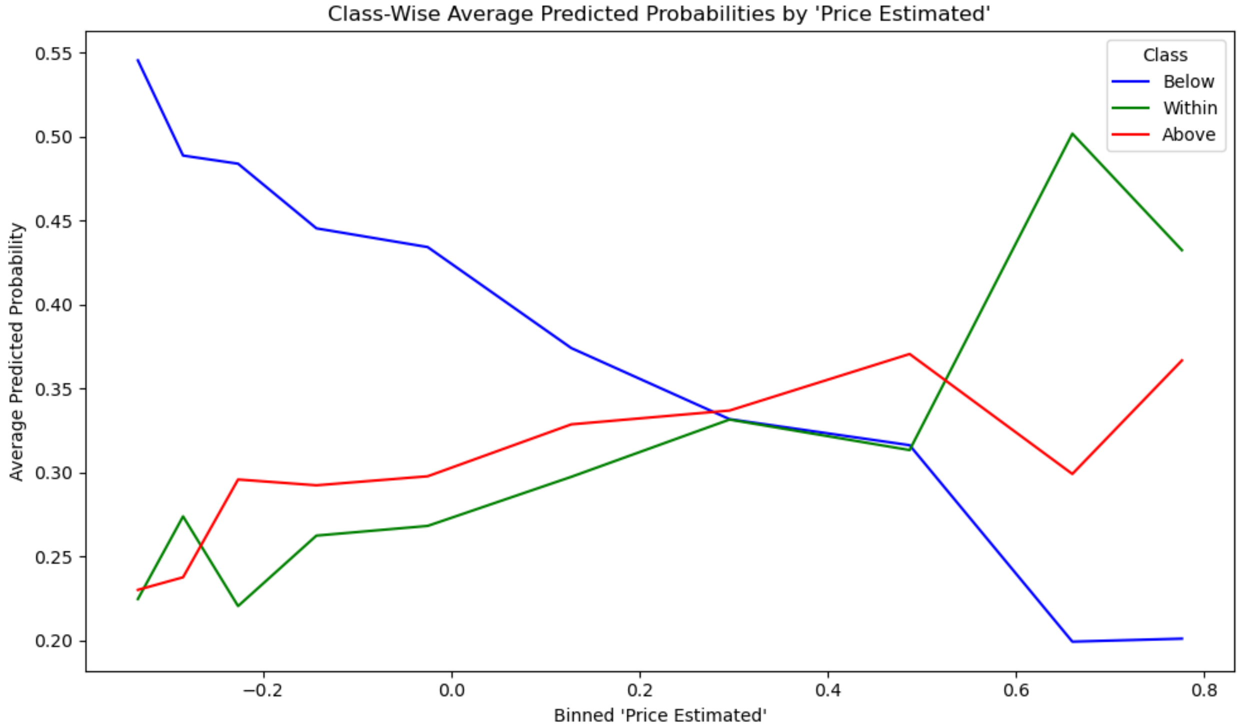

Figure A4. Partial Dependence Plots for 'Full Description Encoded'

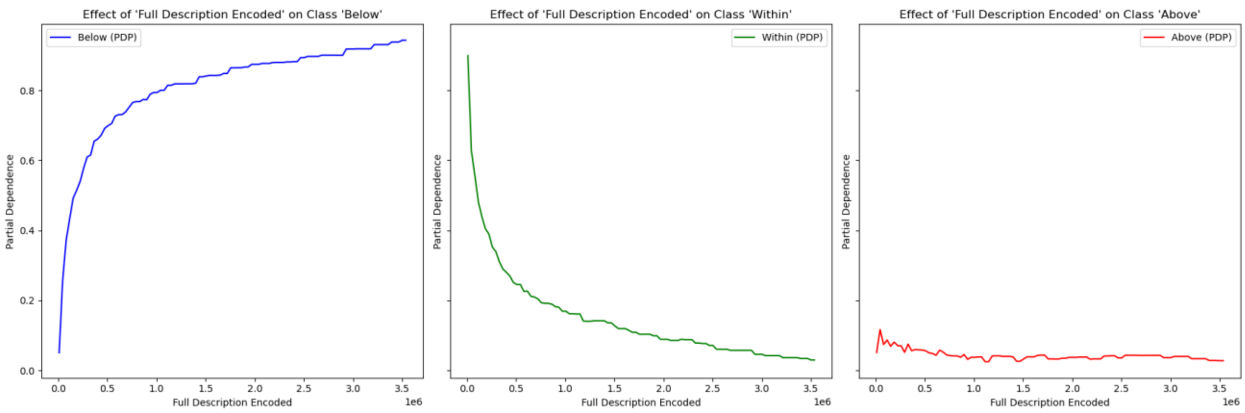

Figure A5. Class-Wise Average Predicted Probabilities by 'Full Description Encoded' (Binned Probabilities)