

Machines and Masterpieces: Predicting Prices in the Art Auction Market

Mathieu Aubry, Roman Kräussl, Gustavo Manso, and Christophe Spaenjers*

This version: 5 March 2019

Abstract

We study the accuracy and usefulness of automated (i.e., machine-generated) valuations for illiquid and heterogeneous assets. We assemble a database of 1.1 million paintings that were auctioned between 2008 and 2015. We use a popular machine-learning technique—neural networks—to develop a price prediction algorithm based on both non-visual and visual artwork characteristics. Our out-of-sample valuations predict auction prices dramatically better than valuations based on a standard hedonic pricing model. Moreover, they help explaining price levels and sale probabilities even after conditioning on auctioneers’ pre-sale estimates. Machine learning is particularly helpful for assets that are associated with higher levels of ex-ante price uncertainty. Finally, we show that it can help overcome experts’ systematic biases in expectations formation.

Keywords: asset valuation, auctions, experts, big data, machine learning, computer vision, art.

JEL Codes: C50, D44, G12, Z11.

*Aubry (mathieu.aubry@imagine.enpc.fr) is at École des Ponts ParisTech, Kräussl (roman.kraussl@uni.lu) is at Luxembourg School of Finance, Manso (manso@haas.berkeley.edu) is at Haas School of Business, University of California at Berkeley, and Spaenjers (spaenjers@hec.fr) is at HEC Paris. The authors thank Will Goetzmann, Katy Graddy, Stefano Lovo, Christophe Pérignon, Luc Renneboog, Kelly Shue, and conference participants at the 2018 ULB Art Market Workshop and the 2019 HEC Paris Data Day for helpful comments. This research is supported by the following grants of the French National Research Agency (ANR): EnHerit (ANR-17-CE23-0008) and “Investissements d’Avenir” (LabEx Ecodec/ANR-11-LABX-0047).

1 Introduction

Many asset markets see a rising importance of automated (i.e., machine-generated) valuations. This is particularly true for durable assets, such as property or collectibles, where trading is infrequent and assets are heterogeneous. In such markets, valuations are tightly—but in a complex manner—linked to a wide range of utility-generating—but not necessarily easy-to-quantify—characteristics. Human eyes and experience have therefore historically been deemed indispensable in valuing houses or paintings. However, advances in computer vision and machine learning may change this, which could have a major impact on how these markets function. Newly-developed automated valuation mechanisms can make markets more liquid, by leading to increased information availability and novel disintermediated trading opportunities for asset owners. In the real estate market, so-called “i-buyers” are already making offers to potential property sellers in an instantaneous and completely online fashion (The Economist (2018)). But having more accurate price predictions could also be productivity-enhancing and risk-reducing for long-existing businesses such as real estate agents and auction houses, banks lending against durable assets, and insurance companies.

However, despite the fastly-increasing importance of automated valuations, not much is known about their accuracy and about the conditions under which they may be helpful. In this paper, we therefore study machine-generated valuations in a setting where human expertise has up until now been considered an essential factor, namely the art auction market. Art prices are particularly difficult to predict because every object is (nearly-)unique and potential bidders’ tastes—and thus willingness-to-pay—may exhibit substantial variation. The art auction market is also an appealing setting to examine the potential of automated valuations because we have a “human” benchmark, namely the pre-sale estimates that are published by auction houses for

each lot and reflect their evaluation of the object.

We use data on 1.1 million painting auctions at more than 350 different auction houses worldwide between 2008 and 2015 from a proprietary database of art sales. For each lot, the database contains information related to the artist, the artwork, and the auction—and also an image of the artwork. We use the data for the period 2008–2014 (our “training set”) to generate different types of statistical hammer price predictions for auctions in the first half of the year 2015 (our “test set”). These predictions can be compared to—and used together with—the auctioneers’ pre-sale estimates for the same works.

Our first—and main—set of predictions are generated through popular machine-learning techniques based on neural networks. Neural networks can be considered as a method to define very large parametric models, in which the parameters are learned from the observations in an iterative and stochastic manner. Our most basic prediction, which we label \hat{p}_{ML}^{txt} , relies on independent variables derived from the textual and numerical data in the database (e.g., artist, year of creation, materials, size, etc.). Next, we use so-called “convolutional neural networks”, which are often used in image-recognition tasks, to develop an algorithm that also considers the image of each work. This gives us a second prediction, which we denote by \hat{p}_{ML}^{img} . Crucially, neither prediction is exploiting any auction-related data (e.g., auction house, location), which are endogenous to auctioneers’ evaluation of the object.

The second type of prediction that we use is also “automated”, but relies on a much less sophisticated (and more traditional) method, namely hedonic regressions. Following Rosen (1974) and real estate scholars, academics studying the art market have linked prices to artwork characteristics, typically employing linear regression models (e.g., Anderson (1974), Renneboog and Spaenjers (2013)). We estimate a standard hedonic model on the training set, and use the

regression coefficients to generate price predictions \hat{p}_{HR} for all artworks in the test data set.

We then use this unique data set to study the distributional characteristics of our automated valuations. We also examine their relative performance in predicting prices, and whether they have explanatory power even after conditioning on auctioneers' pre-sale estimates. Finally, we analyze in depth why and when machine learning is particularly useful for valuing assets. We can summarize our empirical findings as follows.

First, we show that the long right tail in the distribution of art auction prices—with certain works by certain artists fetching very high prices—is mirrored in the distribution of machine-learning valuations but not in that of hedonic predictions. The lack of non-linearities and interaction effects in a standard hedonic model imply that valuations will be much closer to each other than actual prices.

Second, when regressing hammer prices against predictions, we find that our machine-learning predictions perform dramatically better (R^2 of 0.730 and 0.748 for \hat{p}_{ML}^{txt} and \hat{p}_{ML}^{img} , respectively) than predictions based on the hedonic model ($R^2 = 0.046$). Interestingly, the incremental predictive power of images appears to be relatively limited. Even our best machine-learning algorithm does not do as well as auctioneers ($R^2 = 0.877$), but a comparison of the explanatory power shows that it explains only about 15% less of the variation in price outcomes. (Note that auction house experts have access to more information about the artwork's quality and history. Also, we cannot rule out the possibility that auction house estimates *affect* bids.)

Third, we study whether machine-learning predictions have any predictive power after controlling for pre-sale estimates. In other words, we are testing the informational efficiency of auctioneers' estimates. We find that, unlike hedonic predictions, machine-learning predictions help explaining hammer prices conditional on auctioneers' evaluations. Adding our most sophis-

ticated machine-learning valuation to a model with only the auction house estimate increases the R^2 by 1.2%, and is 14.7% more likely to lead to a more accurate than to a less accurate price prediction. We also document substantial predictability of “buy-ins” (i.e., auctions where the highest bid remains below the secret reserve—typically set at a level just below the low estimate—and thus goes unsold). When the auction house estimate is low compared to the machine-learning valuation, the buy-in probability is about 25%, while this probability increases to 45% when the pre-sale estimate is high compared to the automated valuation. Machine-learning valuations thus do not only help predict prices conditional on selling, but also—because of the relation between auction house estimates and reserve prices—the probability of selling in the first place.

Fourth, we hypothesize that machine learning should be more beneficial in settings where prediction is more difficult. Certain artists are associated with a wider heterogeneity in the characteristics of their output or in potential buyers’ tastes. The prices of works by such artists will then consistently exhibit more dispersion. They will also be more difficult to predict accurately—especially by human experts—as disentangling the different drivers of cross-sectional and temporal variation in prices becomes more challenging. We show empirically that there is indeed persistence in the artist-level standard deviation of realized price-to-estimate ratios—a measure that we call “P/E volatility”. Moreover, we find that our machine-learning valuations have a substantially higher incremental explanatory power—over and above auctioneers’ pre-sale estimates—for those artists with the highest levels of relative price volatility. (Machine learning also is more helpful for works by artists with lower average price levels and higher transaction volume.)

Fifth, and finally, we show that machine-learning valuations can help overcome human biases in expectations formation. We present evidence that auctioneers have a tendency to ignore negative information—leading to systematic patterns in their prediction errors. More specifically,

pre-sale estimates are on average too high for works by artists that are associated with relatively low prices or returns in recent years, suggesting that auctioneers are reluctant to adjust their valuations downwards. In these subsets of assets, machine-learning valuations are particularly likely to be able to improve prediction accuracy.

1.1 Related Literature

Our paper contributes to different strands of literature. First, there exists a growing body of work that applies machine-learning techniques to “predict” asset prices and expected returns. A number of recent papers use machine learning to study the cross-section of equity returns (see Gu et al. (2018) and the references therein). Closer to our empirical setting, Lee and Sasaki (2018) find strong predictive power of online home value estimates from Zillow.com in explaining house transaction prices, even when controlling for house and neighborhood characteristics. However, it is unknown what information enters the estimates; moreover, the authors do not have access to human experts’ valuations for the same properties. Also related to our work, because similar in data and methods, is a recent paper by Glaeser et al. (2018) that uses computer vision techniques to link houses’ (or neighboring houses’) appearance to home values.

Second, some recent papers have studied the relative strengths and weaknesses of “men” vs. “machines” in financial-economic decision-making, e.g., in investment management (Abis (2017)) or new venture financing (Catalini et al. (2018)).

Third, our paper also relates to the literature on art prices and auctions. Hedonic regressions are a popular method to analyze the price drivers of artworks and to control for quality differences between works, but the existing literature has not examined their potential for out-of-sample valuation. A number of studies have analyzed whether auction house pre-sale estimates are

unbiased and informationally efficient (Bauwens and Ginsburgh (2000), Ashenfelter and Graddy (2003), Mei and Moses (2005), McAndrew et al. (2012)); these papers, which often use relatively small samples, come to conflicting conclusions.

2 Art Auction Data

2.1 Art Auctions

Art auctions are typically organized as “English” (i.e., ascending-bid, open-outcry) auctions. Prior to the auction, auction house experts publicly share a low and a high estimate for each lot.¹ Auction houses base these estimates on the quality, condition, rarity, etc. of the lot and, crucially, on recent auction prices for similar objects. Each consignor sets a reserve price (in agreement with the auction house), which is the lowest price she is willing to accept. Auction houses do not disclose this reserve, but it cannot exceed the low estimate. If the highest bid at the auction meets or exceeds the reserve price, the object will be sold at this price—the “hammer price”.² If the highest bid remains below the reserve price, the item is said to be “bought in”; it does not sell and instead returns to the consignor.

2.2 Data

The analysis in this paper relies on proprietary data coming from the Blouin Art Sales Index (BASI), which tracks auction sales at hundreds of auction houses worldwide, including the two most important ones, namely Christie’s and Sotheby’s. The data have been used before by

¹These estimates can be found online and in printed sales catalogues. Exceptionally, these estimates are only available “on request”, typically for very valuable items.

²The auction house will charge a “buyer’s premium” on top of the hammer price. Moreover, the consignor has to pay a “seller’s commission”. We do not consider transaction costs here.

Korteweg et al. (2016).

We use data on paintings offered at auctions over the period between start-2008 and mid-2015. In total, our data set contains information on about 1.1 million lots at 372 different auction houses—some of which have different locations—of works by about 125,000 different artists. About two thirds of these auction lots have been sold, while the remaining one third were bought in because the highest bid remained below the consignor’s reserve price. (The year 2008 is the first one for which the underlying database has complete coverage of buy-ins.) For each lot, the database contains information related to the artist (artist name, birth and death year, nationality), the artwork (title, size, year of creation, markings (e.g., signed, dated), some details on materials), and the auction (auction house, auction date, pre-sale low and high estimate, a buy-in indicator, hammer price (if sold)). All estimate and price data are converted to U.S. dollars using the spot rate at the time of the sale. Crucially, we also have access to a high-quality image of each painting through BASI.

In the analysis below, we will use the 985,188 observed lots over the period 2008–2014 as our “training set”, on which we will develop our (hedonic and machine-learning) price prediction algorithms. Our “test set” comprises the 104,404 painting auctions recorded for the first six months of 2015.³ Table 1 shows some descriptive statistics for both subsamples of our data set. About two thirds of all consigned lots sell successfully. In our training data, the average (median) hammer price is \$61,492 (\$3,526), with a long right tail of (extremely) expensive paintings. The difference between median and mean prices (and estimates) points to substantial skewness in the

³The results are not very different when using data for the complete calendar year. However, limiting the analysis to the first six months after the end of the training data allows for a fairer comparison between experts and machines. The pre-sale estimates for the lots that we consider have been determined around the end of 2014. If we also included lots for the second half of 2015, then we would compare automated valuations based on information until the end of 2014 to pre-sale estimates that incorporate information about transactions and market conditions until at least the summer of 2015.

value distribution. In what follows, we will work with natural logs of prices.

Table 1: **Descriptive statistics**

This table reports descriptive statistics separately for the data set that we use for training our automated valuation methods (painting auctions over the period January 2008 to December 2014) and for the data set that we use for testing the accuracy of the different price predictions (painting auctions over the period January to June 2015).

	2008–2014 Training set	Jan–June 2015 Test set
<i>N</i> observations	985,188	104,404
<i>N</i> distinct artists	116,694	37,094
<i>N</i> distinct auction houses	369	267
% with pre-sale estimates	97.8%	96.9%
Median low estimate (\$)	3,000	2,000
Mean low estimate (\$)	37,769	39,194
Median high estimate (\$)	4,186	3,000
Mean high estimate (\$)	53,673	55,450
% sold	65.2%	66.3%
Median hammer price (\$)	3,526	2,228
Mean hammer price (\$)	61,492	67,697
Median price-to-low-estimate ratio	1.025	1.000
Mean price-to-low-estimate ratio	1.446	1.367
Median price-to-high-estimate ratio	0.775	0.750
Mean price-to-high-estimate ratio	1.030	0.969

2.3 Auctioneers’ Pre-Sale Estimates

As explained before, the auction house typically shares a “low” and a “high” estimate prior to an auction. Table 1 includes some descriptive statistics for both estimates and for price-to-estimate ratios. The median sale is associated with a price that is virtually equal to the low pre-sale estimate and about three quarters of the high estimate; the mean price-to-estimate ratios are higher. Below, we will use the logged average of the low and the high estimate, which we

can label as \hat{p}_{AH} , as a predictor of the price outcome.⁴

Auction houses will typically argue that their estimates serve as “as an approximate guide to current market value and should not be interpreted as a representation or prediction of actual selling prices”. Yet, at the same time, the estimates are said to be representing auction house experts’ opinion “about the range in which the lot might sell at auction” (quotes taken from Sotheby’s website).

One concern could be that auction houses may strategically choose to be relatively aggressive or conservative in their estimates, in an attempt to affect bidder participation or bids conditional on participation. Yet, auction theory says that “honesty is the best policy” (Milgrom and Weber (1982)). Moreover, even if estimates are shaded upwards or downwards, higher estimates should still imply higher expected prices, which is what will really matter for our empirical analysis.

Auction houses’ pre-sale estimates will be based on some of the easily-observable tangible characteristics of the artwork presented and used below in the development of our price prediction algorithms, such as the identity of the artist, materials, size, etc. However, auctioneers will also take into account artwork condition, provenance, and other quality factors not observable to the econometrician working with art sales databases such as ours. Therefore, we might reasonably expect auction house estimates to be better predictors of prices than valuations—automated or not—based on a smaller information set.

There exists another reason for why auctioneers might be expected to “beat” our algorithms. Namely, any random noise element in auctioneers’ estimates will spill over into bids if potential buyers anchor on those estimates. By contrast, our own automated valuations were of course

⁴The size of the spread between the low and the high estimate does not show much variation once controlling for auction house and low estimate, and is thus unlikely to contain any relevant information about the auctioneer’s confidence in her own estimate. For example, in our training data, 175 out of the 177 lots with a low estimate of \$100,000 offered at Christie’s in the U.S. have a high estimate of \$150,000.

not available to potential buyers at the time of the auctions that we study.

3 Automated Valuations

We will now use the training set—works auctioned over the years 2008–2014—to develop algorithms that predict the price of any artwork based on its characteristics. Both when using machine learning (cf. subsection 3.1) and when running standard linear hedonic regressions (cf. subsection 3.2), we train the algorithm on hammer prices for successful sales only.⁵ All prices are log-transformed. Moreover, as we will drop artworks with mid estimates below \$1,000 in our empirical tests in the next section, we winsorize all prices at this level prior to our training.⁶ After we develop our models, we can apply them to the test set—auctions in the first half of the year 2015—to generate out-of-sample price predictions that can be compared to auction house experts’ pre-sale estimates.

3.1 Machine-Learning Predictions

3.1.1 Methodology

The machine-learning technique that we employ in this paper is neural networks. Neural networks can be seen as a way to define very large parametric models. Their parameters—also called “weights”—are typically learned from observations in an iterative and stochastic manner. The backbone of the architecture we use is a “multi-layer perceptron”, which alternates between linear operations and non-linearities.

⁵Crucially, results are very similar when using an imputed price equal to 75% of the low estimate for bought-in works. (Ashenfelter and Graddy (2011) estimate the average reserve to be approximately equal to 71% of the low estimate. McAndrew et al. (2012) find an average reserve of 73.5% of the low estimate.

⁶We also winsorize a handful of prices at \$50 million. Any price variation above this level is arguably largely idiosyncratic. In any case, we do not have any works with an estimate exceeding this level in our test data.

To generate representations from the artwork images, we use convolutional neural networks (CNNs). CNNs are often used in image-recognition tasks. They are neural networks designed to have the capacity to learn very complex functions of images' pixel values, while taking advantage of the spatial structure of an image in which nearby pixels are correlated. CNNs are able to predict very reliably semantic and texture information from an image. Given sufficient training data, they can predict the genre, creation date, or creator of an artwork, as well as human aesthetic judgments (Karayev et al. (2013), Tan et al. (2016), Strezoski and Worring (2017)).

Because neural networks can have hundreds of thousands—or even millions—of parameters, they can often perfectly explain any data they are trained with. Therefore, it is crucial to apply them first to validation data during training (to optimize meta-parameters such as their architecture and training procedure), and then to a test set that they have not seen during training and that has not been considered for validation (to test their out-of-sample performance).

3.1.2 Variables and Predictions

Using the methods described above, we come up with two different machine-learning predictions. Our most basic prediction, which we label \hat{p}_{ML}^{txt} , relies only on the following independent variables derived from the textual and numerical information in the database:

- **Artist and artist nationality.** Table 1 showed that we have more than 100,000 distinct artists in the training data. Together, these artists represent 168 different nationalities.
- **Artwork creation year.** We have precise information on the creation year for about half of all observations. A large majority of the works for which we have this information date from the twentieth century.

- **Artwork size.** Width and height are included in the database for nearly all observations. We winsorize these size variables at the 0.5% lowest and highest values in each year. After winsorizing, the median (mean) width in the training data is 55 (65) centimeters, while the median (mean) height is 52 (63) centimeters.
- **Artwork markings.** We create three dummy variables that equal one if the artwork is (1) signed; (2) dated; or (3) inscribed by the artist. These categories are not mutually exclusive. In the training data, at least one of these indicator variables equals one for 82.0% of all works.
- **Artwork title.** We create eight indicator variables for the following groups of terms that are used frequently in artwork titles:⁷ (1) untitled, sans titre, senza titolo, ohne titel, sin titulo, o.t.; (2) composition, abstract, composizione, komposition; (3) landscape, paysage, paesaggio, seascape, marine, paysage; (4) still life, flowers, nature morte, bouquet de fleurs, nature morta, vase de fleurs; (5) figure, figura, character; (6) nude; (7) portrait, mother and child; (8) self-portrait, self portrait. These categories are not mutually exclusive. In the training data, at least one of these indicator variables equals one for 22.7% of all works.
- **Artwork materials.** We create 18 indicator variables for the following terms that appear frequently in the description of the materials and support: (1) oil; (2) watercolor; (3) acrylic; (4) ink; (5) gouache; (6) bronze; (7) mixed media; (8) pastel; (9) lithograph; (10) poster; (11) etching; (12) pencil; (13) canvas; (14) board; (15) panel; (16) paper; (17) masonite; (18) wood. These categories are not mutually exclusive. In the training data, only 2.8% of all lots fall outside of all of these categories. For more than 75% of all lots,

⁷To come up with this classification, we consider the 50 most frequent titles in our sample, and manually create groups of related words.

exactly two dummies equal one (as would be the case, for example, if the description reads “oil on canvas”).

The algorithm also has access to the year of sale. When making out-of-sample predictions for the test set (i.e., early 2015), it will do so as if these observations are from the final year of the training set (i.e., 2014). In principle it can put more weight on more recent observations.

Next, we develop an algorithm that considers the image of each painting in the way described before. This gives us the prediction denoted by \hat{p}_{ML}^{img} . This prediction thus relies on non-visual characteristics, visual characteristics, and any interactions between the two.

Crucially, neither \hat{p}_{ML}^{txt} nor \hat{p}_{ML}^{img} uses any auction-related variables, such as the identity and location of the auction house, or the month in which the auction takes place. Where and when a lot is being offered is endogenous to auctioneers’ evaluation of the market value of the object, so a prediction algorithm that uses such information cannot be considered fully “automated”. Moreover, a method that necessitates such information would not be useful to most market participants, as it would not allow to value an item as long as it is not coming up for auction.

3.2 Hedonic Predictions

A natural benchmark, in particular for \hat{p}_{ML}^{txt} , is the prediction generated by a standard linear hedonic regression model applied to the artwork’s characteristics. More specifically, we can estimate the following model using ordinary least squares on the observations in the training set:

$$p_{i,t} = \alpha + X_i' \beta + \gamma_t + \varepsilon_i, \tag{1}$$

where $p_{i,t}$ is the log-transformed hammer price of painting i , X_i is a vector of hedonic variables, and γ_t are year fixed effects. We here use the following hedonic variables: artist fixed effects,

artwork height and width (and their squares), and the artwork marking, title, and material dummies introduced before. Unlike hedonic models estimated to measure quality-controlled changes in average price levels over time (e.g., Renneboog and Spaenjers (2013)), we do not include controls for auction-related variables. Instead, we limit ourselves to truly exogenous artwork characteristics, for reasons outlined before.

Table 2 shows the hedonic regression coefficients. (We do not show standard errors or significance levels, but nearly all coefficients are highly statistically significant.) The results are in line with findings in previous literature. For example, substantially higher prices are paid for works that are bigger, signed or dated, self-portraits, and oils. The (in-sample) R^2 is 0.789, mainly thanks to the many artist fixed effects; estimating the same model without artist dummies yields an R^2 of 0.177.

We can then use the estimated coefficients reported in Table 2 to generate out-of-sample price predictions \hat{p}_{HR} for all lots without missing values on any of the variables included in the hedonic regression model. In line with what we did before, we make predictions as if the out-of-sample observations are from the year 2014 by using the coefficient on the fixed effect for that year.

4 Results

We will in this section relate price outcomes in the test data (i.e., in the first half of 2015) to our different predictions. Unless otherwise noted, we focus on hammer prices. We also impose a number of data filters. We exclude a low number of works by artists who did not have any auctions over the training period; our hedonic model would not even generate a price prediction for such artists. We also drop objects with mid estimates below \$1,000, for which variation in

Table 2: **Hedonic regression coefficients**

This table reports estimated ordinary least squares coefficients for the hedonic regression model shown in Eq. (1). The dependent variable is the logged hammer price. Height and width are measured in meters. The model is estimated over all transactions in our training data set, which covers the period 2008–2014.

Year fixed effects	Yes	Materials: oil	0.535
Artist fixed effects	Yes	Materials: watercolor	0.039
Height	1.131	Materials: acrylic	0.293
Height squared	-0.258	Materials: ink	-0.203
Width	1.247	Materials: gouache	0.230
Width squared	-0.248	Materials: bronze	0.586
Markings: signed	0.211	Materials: mixed media	0.129
Markings: dated	0.140	Materials: pastel	0.091
Markings: inscribed	0.050	Materials: lithograph	-2.173
Title: untitled	-0.173	Materials: poster	-0.638
Title: composition	-0.144	Materials: etching	-1.639
Title: landscape	-0.132	Materials: pencil	-0.310
Title: still life	-0.050	Materials: canvas	0.232
Title: figure	-0.068	Materials: board	0.084
Title: nude	-0.099	Materials: panel	0.236
Title: portrait	-0.212	Materials: paper	-0.167
Title: self-portrait	0.565	Materials: masonite	0.114
		Materials: wood	0.167
N			625,449
R^2			0.789

(logged) prices is substantial but not very meaningful economically.⁸ Finally, we drop sales where the sale price is below 10% of the low estimate or above ten times the high estimate. Some of these outliers may be cases where either the price or the estimate is incorrectly recorded in the database, or some (to us) unobservable event happened between estimation and auction (e.g., a re-attribution).

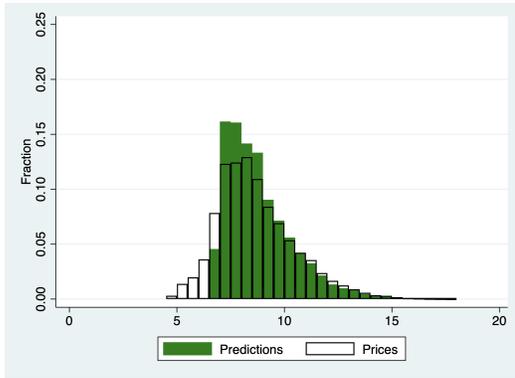
4.1 Distributions of Predictions

We start our analysis by plotting the distributions for all of the previously-introduced price predictions, and by comparing them to the distribution of realized hammer prices. The results are shown in Figure 1. Strikingly, the long right tail of prices is captured adequately by our machine-learning valuations, but not at all by the hedonic ones. (The left tail of prices constitutes of items that had auction house estimates in the lower \$1,000s, but sold for less than \$1,000.) Hedonic regressions are not well-suited to capture the whole distribution of prices because of the low dimensionality of the parameter space. Take, for example, the works of Pablo Picasso. All Picasso paintings will have very similar hedonic valuations, largely driven by the estimated coefficient on the artist fixed effect. Differences between hedonic Picasso predictions will be due to the average price differences—aggregated across all artists—between works with different sizes, markings, materials, etc. By contrast, our machine-learning predictions can take into account that large oil paintings by Picasso carry an above-average premium, especially if they are—or look like they are—from, for example, 1907.

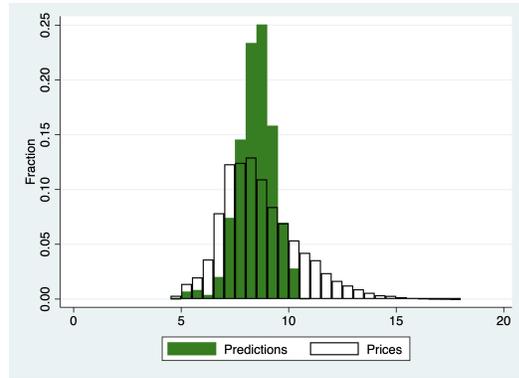
⁸Such objects are unlikely to have any resale value. Virtually none of these lots were offered by main auction houses Christie’s or Sotheby’s.

Figure 1: **Distributions of predictions and prices**

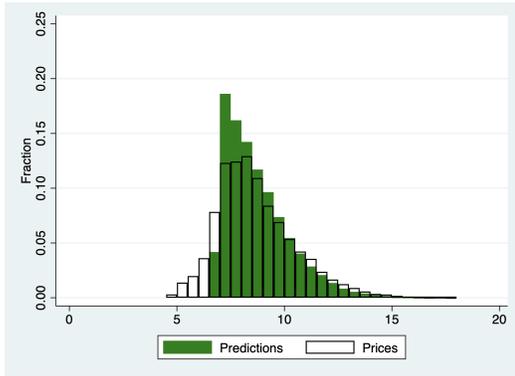
This figure shows the distributions of logged hammer prices and different price predictions over all transactions in our test data set, which covers the period January–June 2015. The predictions are auction house estimates (\hat{p}_{AH}) in subfigure (a), hedonic predictions (\hat{p}_{HR}) in subfigure (b), and machine-learning valuations without and with relying on image information (\hat{p}_{ML}^{txt} and \hat{p}_{ML}^{img}) in subfigures (c) and (d).



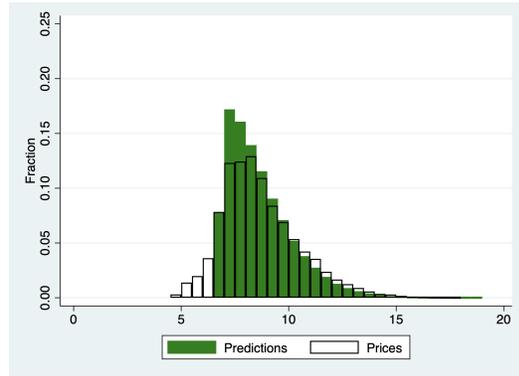
(a) Prediction used: \hat{p}_{AH}



(b) Prediction used: \hat{p}_{HR}



(c) Prediction used: \hat{p}_{ML}^{txt}



(d) Prediction used: \hat{p}_{ML}^{img}

4.2 Comparison of Predictive Power

To analyze how hammer prices line up with the different predictions, we estimate the following regression model using ordinary least squares in the test data:

$$p_i = \alpha + \beta \hat{p}_i + \varepsilon_i, \quad (2)$$

where p_i is the log hammer price of artwork i and \hat{p}_i is the prediction for the same artwork. Columns 3–4 of Table 3 show the results for the different machine-learning predictions, which can be compared to the auction house pre-sale estimates in column 1 and the hedonic predictions in column 2.

Table 3: **Predictions and prices**

This table reports estimated ordinary least squares coefficients for the regression model shown in Eq. (2). The dependent variable is the logged hammer price. The models are estimated using the transactions in our test data set, which covers the period January–June 2015. Standard errors, which are two-way clustered at the artist and auction month level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
<i>Prediction:</i>	\hat{p}_{AH}	\hat{p}_{HR}	\hat{p}_{ML}^{txt}	\hat{p}_{ML}^{img}
Prediction	1.072 *** (0.010)	0.460 *** (0.090)	1.012 *** (0.011)	1.011 *** (0.011)
Constant	-0.758 *** (0.105)	4.796 *** (0.777)	-0.178 (0.104)	-0.135 (0.105)
N	45,734	44,721	45,734	45,734
R^2	0.877	0.046	0.730	0.748

What do we learn from these results? First, auction house estimates explain about 88% of the variation in hammer prices. Second, a standard hedonic model performs extremely poorly in this out-of-sample setting, despite the large in-sample R^2 documented before. Third, even when only using non-visual variables, our machine-learning predictions do dramatically better than

the hedonic predictions. They do not explain as much of the variation in price results as pre-sale estimates, but the R^2 s in columns 3 and 4 of Table 3 are still 83% and 85%, respectively, of that in column 1. Fourth, the incremental explanatory power of images is relatively limited. The R^2 in column 4 is only 2.5% higher than that in column 3. Fifth, average realized prices line up with machine-learning predictions near a 45-degree line: the constants are close to zero, and the slope coefficients are close to one. (By contrast, auction house estimates seem to slightly overestimate market values for cheaper works and underestimate them for more expensive works.)

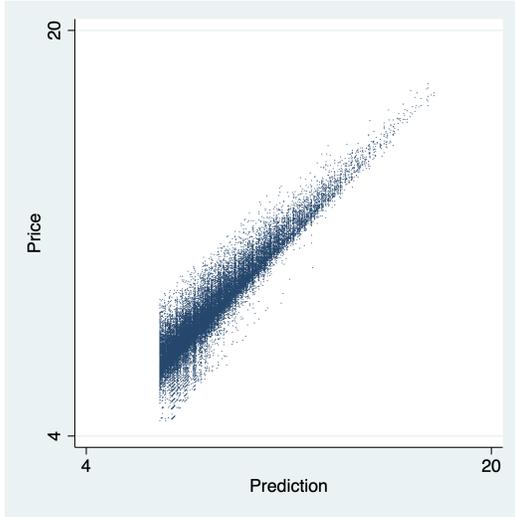
These can also be visualized through scatter plots. The different panels in Figure 2 show the predictions \hat{p}_{AH} , \hat{p}_{HR} , \hat{p}_{ML}^{txt} , and \hat{p}_{ML}^{img} on the horizontal axis and hammer prices on the vertical axis. The plots based on the machine-learning predictions have a shape similar to that based on the auction house estimates, but exhibit more noise. The plot showing the relation between hedonic predictions and hammer prices looks very different. As indicated before, hedonic valuations tend to be clustered together much more.

4.3 Test of Efficiency of Auctioneers' Pre-Sale Estimates

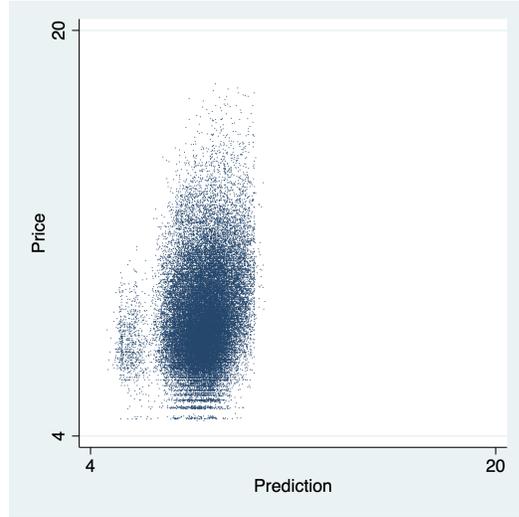
We now turn to studying whether hedonic and machine-learning predictions of hammer prices have some additional explanatory power after controlling for pre-sale estimates. Another way of seeing this is as a test of the informational efficiency of auctioneers' estimates. Column 1 of Table 4 repeats the results reported in the first column of Table 3. The next columns then add the other predictions, which were first orthogonalized with respect to \hat{p}_{AH} , to the regression model. The bottom rows in columns 2–4 show, first, the relative change in R^2 compared to the model in column 1, and, second, the likelihood that the predicted valuation coming out of the regression model is closer to the observed price.

Figure 2: **Predictions and prices**

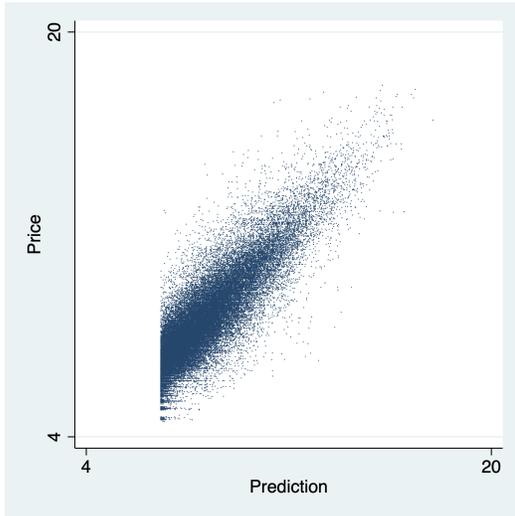
This figure plots logged hammer prices against different price predictions over all transactions in our test data set, which covers the period January–June 2015. The predictions are auction house estimates (\hat{p}_{AH}) in subfigure (a), hedonic predictions (\hat{p}_{HR}) in subfigure (b), and machine-learning valuations without and with relying on image information (\hat{p}_{ML}^{txt} and \hat{p}_{ML}^{img}) in subfigures (c) and (d).



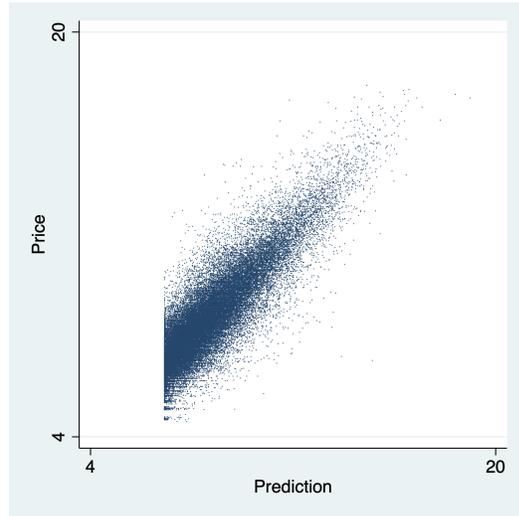
(a) Prediction used: \hat{p}_{AH}



(b) Prediction used: \hat{p}_{HR}



(c) Prediction used: \hat{p}_{ML}^{txt}



(d) Prediction used: \hat{p}_{ML}^{img}

Table 4: **Efficiency of pre-sale estimates**

This table reports estimated ordinary least squares coefficients for regression models where the dependent variable is the logged hammer price. Column 1 only has \hat{p}_{AH} as an explanatory variable, and thus repeats the first column of Table 3. Columns 2–4 then add the other predictions, which were first orthogonalized with respect to \hat{p}_{AH} . The models are estimated using the transactions in our test data set, which covers the period January–June 2015. Standard errors, which are two-way clustered at the artist and auction month level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
<i>Automated valuation (orthog.):</i>		\hat{p}_{HR}	\hat{p}_{ML}^{txt}	\hat{p}_{ML}^{img}
\hat{p}_{AH}	1.072 *** (0.010)	1.073 *** (0.010)	1.038 *** (0.008)	1.033 *** (0.008)
Automated valuation		-0.016 * (0.007)	0.216 *** (0.031)	0.239 *** (0.029)
Constant	-0.758 *** (0.105)	-0.774 *** (0.106)	-0.438 *** (0.075)	-0.402 *** (0.074)
N	45,734	44,721	45,734	45,734
R^2	0.877	0.877	0.886	0.888
% increase in R^2 relative to (1)		-0.1%	+1.0%	+1.2%
% of predictions more accurate than (1)		50.1%	52.7%	53.4%

The results in Table 4 show that, in contrast to hedonic valuations, machine-learning valuations can help in predicting price outcomes conditional on pre-sale estimates. A relative increase in R^2 of 1.2% may not sound dramatic, but we need to consider the full extent of variation in artwork price levels that exists. From the results in the last column we can actually compute that adding a machine-learning valuation is $53.4\%/46.6\% - 1 = 14.7\%$ more likely to lead to a more accurate than to a less accurate price prediction. We thus believe the improvement in predictive power to be economically significant.

We have so far considered the relation between our predictions and hammer prices, which are only observable if an item sells successfully. However, our finding that we can improve on the pre-sale estimate to predict price outcomes suggests that there might also be some predictability of whether a lot will be bought in. More specifically, if the estimate is set relatively high for

a certain work, then the reserve—decided jointly upon by auctioneer and consignor, but never above the low estimate—is also likely to be relatively high. We can thus expect to see more buy-ins if our automated valuations are low compared to the pre-sale estimates. We test this hypothesis in Table 5, which shows the results for probit regressions—estimated over all lots in the test data set—where the dependent variable is a dummy that equals one if the item was bought in. In line with our expectations, we find that machine-learning artwork valuations help predicting buy-ins.

Table 5: **Predictability of buy-ins**

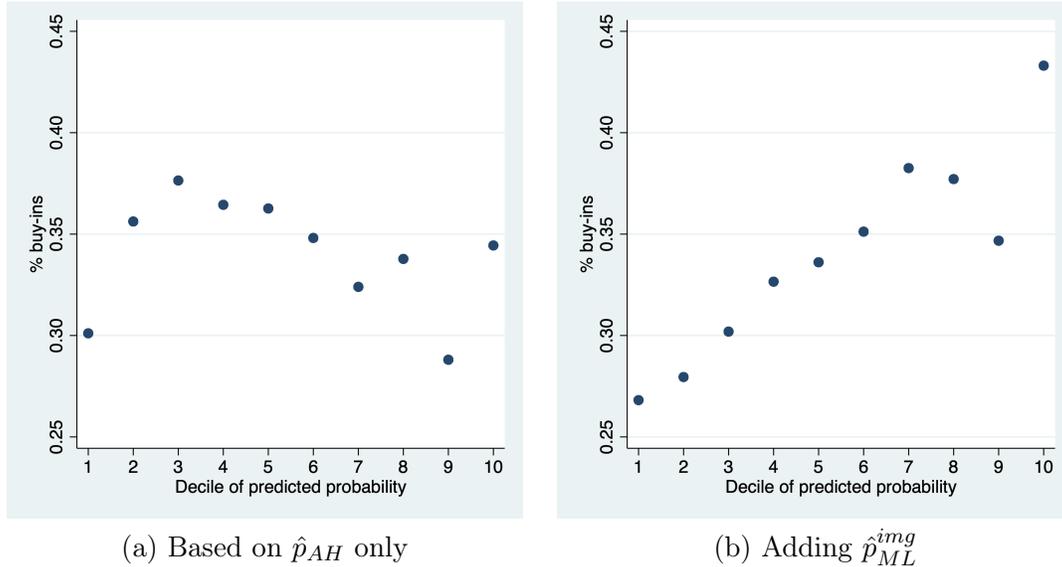
This table reports estimated probit coefficients for regression models where the dependent variable is a dummy variable that equals one if a lot is “bought in” (i.e., if the highest bid remains below the reserve price). Column 1 only has \hat{p}_{AH} as an explanatory variable. Columns 2–4 then add the other predictions, which were first orthogonalized with respect to \hat{p}_{AH} . Each model also includes a constant, which is not shown here. The models are estimated using the transactions in our test data set, which covers the period January–June 2015. Standard errors, which are clustered at the artist level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
<i>Automated valuation (orthog.):</i>		\hat{p}_{HR}	\hat{p}_{ML}^{txt}	\hat{p}_{ML}^{img}
\hat{p}_{AH}	-0.004 (0.005)	-0.004 (0.005)	0.009 ** (0.005)	0.018 *** (0.005)
Automated valuation		0.008 (0.008)	-0.097 *** (0.009)	-0.161 *** (0.009)
N	69,326	67,867	69,326	69,326
Pseudo R^2	0.000	0.000	0.003	0.007

To evaluate the economic significance of our results, we plot in Figure 3 the realized out-of-sample buy-in frequency as a function of deciles of the predicted buy-in likelihoods that follow from the probit models in columns 1 and 4 of Table 5. We can see that we find substantial predictability of buy-ins when adding our machine-learning valuations to the model. Subfigure (b) shows that the buy-in probability is about 25% when \hat{p}_{ML}^{img} is relatively high (compared to the auction house estimate), while this frequency approaches 45% when \hat{p}_{ML}^{img} is relatively low.

Figure 3: **Predictability of buy-ins**

This figure shows the frequency of buy-ins over all auction lots in our test data set, which covers the period January–June 2015, for each decile of predicted buy-in probability based on two different models from Table 5. In subfigure (a), the buy-in probabilities are predicted using the model in column 1 of Table 5, which only has \hat{p}_{AH} as an explanatory variable. In subfigure (b), the predicted probabilities come from column 4, which adds \hat{p}_{ML}^{img} to the model.



4.4 Variation in Prediction Difficulty

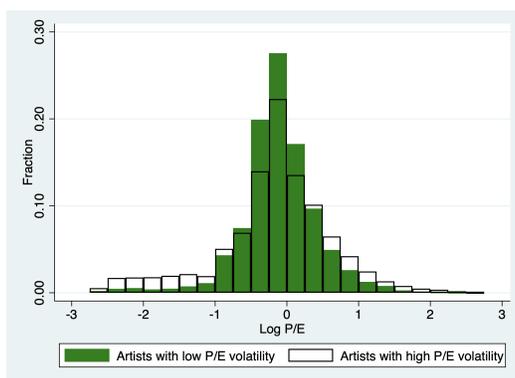
It is likely that some artists have a more heterogeneous output than others, and therefore exhibit more price dispersion. Also, some artists may be associated with more heterogeneity in (potential) buyers’ tastes and preferences, and therefore even their quality-controlled prices fluctuate (Lovo and Spaenjers (2018)). Prices of both types of artists will be more difficult to predict accurately—both by humans and machines, but especially by the former—as disentangling the different drivers of cross-sectional and temporal variation in prices becomes more challenging.

We here first check whether there indeed exists artist-level persistence in the extent to which auction prices deviate from pre-sale estimates. We compute for each artist the standard deviation of the logged price-to-estimate ratios (i.e., $p - \hat{p}_{AH}$) in the training data. We then rank all lots in the test data on this measure, and create four quartiles, where the first (resp. fourth) quartile

has the lots by the artists with the lowest (resp. highest) “P/E volatility”. Figure 4 compares the distributions of logged price-to-estimate ratios for the extreme quartiles. We clearly see a wider distribution for high-P/E-volatility artists. So artists that have historically (in the training data) exhibited larger price deviations from estimates indeed continue to show a higher relative price variation (in the test data).

Figure 4: **Persistence in prediction difficulty**

This figure shows different distributions of logged price-to-estimate ratios (i.e., $p - \hat{p}_{AH}$) in the test data set, which covers the period January–June 2015. We classify all lots in different quartiles based on the artist-level standard deviation of logged price-to-estimate ratios in the training data set, which covers the period 2008–2014. We then show the distribution for the first quartile (“Artists with low P/E volatility”) and the fourth quartile (“Artists with high P/E volatility”).



We now study whether machine-learning predictions are more helpful for the artworks that are harder to value (or, to be more precise, that are by artists who have historically been associated with less accurate auction house estimates). To do so, Table 6 repeats the analysis from Table 4 on split samples based for the first and the fourth quartile of P/E volatility. The regression coefficient on the (orthogonalized) \hat{p}_{ML}^{img} is indeed much higher for the latter set of artists. Our predictions also have a much higher incremental explanatory power—over and above auctioneers’ pre-sale estimates—in column 4 than in column 2.

We can more generally examine what are the drivers of the likelihood that machine learning helps generating a more accurate valuation by running a probit regression in which the dependent

Table 6: **Variation in efficiency of pre-sale estimates**

This table reports estimated ordinary least squares coefficients for regression models identical to those shown in columns 1 and 4 of Table 4, but now estimated separately for lots in the lowest and highest quartile of artist-level “P/E volatility” (measured as the artist-level standard deviation of the logged price-to-estimate ratios (i.e., $p - \hat{p}_{AH}$) in the training data set, which covers the period 2008–2014. The models are estimated using the transactions in our test data set, which covers the period January–June 2015. Standard errors, which are two-way clustered at the artist and auction month level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
	Low P/E volatility		High P/E volatility	
\hat{p}_{AH}	1.041 *** (0.007)	1.015 *** (0.006)	1.118 *** (0.021)	1.061 *** (0.015)
\hat{p}_{ML}^{img} (orthog.)		0.173 *** (0.027)		0.347 *** (0.039)
Constant	-0.469 *** (0.073)	-0.228 *** (0.058)	-1.209 *** (0.215)	-0.682 *** (0.143)
N	10,294	10,294	10,832	10,832
R^2	0.891	0.896	0.841	0.864
% increase in R^2		+0.6%		+2.7%
% of predictions more accurate than benchmark		53.0%		54.4%

variable is a dummy that equals one if \hat{p}_{ML}^{img} is closer to p than \hat{p}_{AH} . We show the results for two specifications in Table 7. In column 1, the only explanatory variable is the P/E volatility measure created and used before. As expected, we see a strong positive correlation. In subsequent columns, we include two other covariates, namely the average estimate and number of lots for the artist (as measured in the training data). Adding these controls does not alter the coefficient on the P/E volatility variable much, but provides some additional insights. First, auction houses do relatively better (or, in other words, machines do relatively worse) for more expensive artists. This is intuitive, and different factors could play a role: auctioneers may have more market knowledge for such artists; auctioneers may do more effort to produce accurate estimates for such artists; hard-to-quantify factors like condition and provenance may be more important; etc. Second, auctioneers do well for relatively *illiquid* artists, at least when controlling for their

price level. This may at first sound counterintuitive, but note that the machine is not fed data on which artists can be considered “substitutes”. Such art-historical knowledge may be relevant when valuing a work by an artist with few recent sales, as it will allow to consider prices in recent auctions by similar artists. Also, when not a lot of recent price points are available, auctioneers can still rely on soft information that they have about the current demand for a certain artist, while the machine is basically left in the dark.

Table 7: **Drivers of relative accuracy of machine-learning predictions**

This table reports estimated probit coefficients for regression models where the dependent variable is a dummy variable that equals one if \hat{p}_{ML}^{img} is closer to p than \hat{p}_{AH} . The three independent variables are measured at the artist level using data from the training data set, which covers the period 2008–2014. Each model also includes a constant, which is not shown here. The models are estimated using the transactions in our test data set, which covers the period January–June 2015. Standard errors, which are clustered at the artist level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
P/E volatility	0.175 *** (0.047)			0.140 *** (0.042)
Log average estimate		-0.165 *** (0.007)		-0.193 *** (0.009)
Log number of lots			-0.004 (0.007)	0.050 *** (0.009)
N	43,061	45,690	45,734	43,061
Pseudo R^2	0.001	0.020	0.000	0.023

4.5 Overcoming Biases in Pre-Sale Estimates

We have established that auctioneers’ estimates are not always efficient, in the sense that they can be improved upon as a prediction of the hammer price. We now want to explore one particular channel that may explain why this is the case. We know from prior work that humans often only slowly update expectations or beliefs. Moreover, participants in real asset markets

tend to be reluctant to adjust appraisals or selling prices downwards.⁹ We can therefore expect that recent changes in price levels—and especially recent downward price trends—may not always be reflected in auction house estimates, which would make alternative (automated) valuations more helpful.

We test our working hypothesis as follows. Like before, we classify all lots in the test data in quartiles, using two different metrics. First, we rank lots based on artist-level average price-to-estimate ratios in the training data. Second, we can identify more than 10,000 resales of identical items in the training data, and compute artist-level average annualized returns. We then compare the distributions of price-to-estimate ratios for lots by artists with relatively low recent price surprises or returns to those by artists with relatively high recent price surprises and returns. The results are shown in Figure 5. As we were expecting, we see that low recent prices or returns (in the training data) are associated with low price-to-estimate ratios (in the test data).

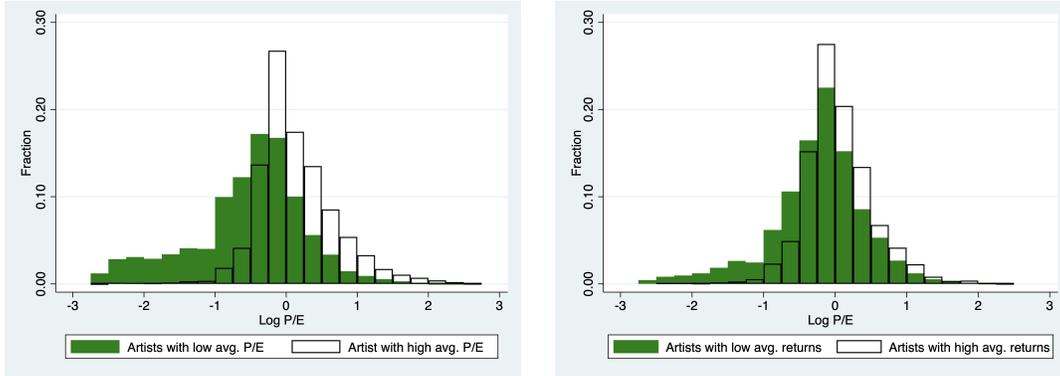
We then compute a number of statistics for the different quartiles that we constructed. First, we measure the proportion of lots for which the machine-learning valuation exceeds the auction house estimate. Second, we report the proportion of transactions for which a prediction model like the one estimated before gives a more accurate prediction once including the machine-learning valuation (cf. columns 1 and 4 of Table 4). The results are reported in Table 8. We see that the machine-learning valuations are much less likely to be above the pre-sale estimates and much more likely to improve accuracy for lots by artists with relatively low recent prices and returns.¹⁰

⁹This is well-known for the housing market, but is also true for collectibles markets (Velthuis (2007), Dimson and Spaenjers (2011)).

¹⁰We also repeat the analysis based on quartiles of artist-level price-to-estimate ratios *at the same auction house*, and the results are even stronger. This points to the importance of personal or institutional persistence in beliefs about market value.

Figure 5: **Systematic patterns in price deviations from pre-sale estimates**

This figure shows different distributions of logged price-to-estimate ratios (i.e., $p - \hat{p}_{AH}$) in the test data set, which covers the period January–June 2015. To create subfigure (a), we classify all lots in quartiles based on the artist-level average logged price-to-estimate ratio (i.e., $p - \hat{p}_{AH}$) in the training data set, which covers the period 2008–2014. We then compare the distribution for the first quartile (“Artists with low avg. P/E”) to that of the fourth quartile (“Artists with high avg. P/E”). In subfigure (b), we do a similar exercise but creating quartiles based on the artist-level average log return on observed resales in the training data.



(a) Sample split on recent P/E ratios

(b) Sample split on recent returns

Table 8: **Systematic patterns in relative accuracy of machine-learning predictions**

This table reports a number of statistics for different subsamples of the test data set, which covers the period January–June 2015. In Panel A, we classify all lots in quartiles based on the artist-level average logged price-to-estimate ratio (i.e., $p - \hat{p}_{AH}$) in the training data set, which covers the period 2008–2014. In Panel B, we do a similar exercise but creating quartiles based on the artist-level average logged return on observed resales in the training data.

Panel A: Lots ranked by average P/E of artist

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
% where $\hat{p}_{ML}^{img} > \hat{p}_{AH}$	26.0%	39.1%	44.2%	51.3%
% of predictions more accurate than benchmark	68.0%	47.8%	49.5%	50.9%

Panel B: Lots ranked by average returns on resales of artist

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
% where $\hat{p}_{ML}^{img} > \hat{p}_{AH}$	38.5%	41.7%	45.0%	46.3%
% of predictions more accurate than benchmark	59.7%	52.0%	50.5%	52.6%

5 Conclusion

We study the accuracy and usefulness of automated (i.e., machine-generated) valuations for illiquid and heterogeneous assets. We assemble a database of 1.1 million paintings that were auctioned between 2008 and 2015. We use a popular machine-learning technique—neural networks—to develop a price prediction algorithm based on both non-visual and visual artwork characteristics. Our out-of-sample valuations predict auction prices dramatically better than valuations based on a standard hedonic pricing model. Moreover, they help explaining price levels and sale probabilities even after conditioning on auctioneers’ pre-sale estimates. Machine learning is particularly helpful for assets that are associated with higher levels of ex-ante price uncertainty. Finally, we show that it can help overcome experts’ systematic biases in expectations formation.

References

- Abis, Simona (2017), “Man vs. machine: Quantitative and discretionary equity management.” Working paper. URL <https://www8.gsb.columbia.edu/researcharchive/getpub/25656/p>.
- Anderson, Robert C. (1974), “Paintings as an investment.” *Economic Inquiry*, 12, 13–26.
- Ashenfelter, Orley and Kathryn Graddy (2003), “Auctions and the price of art.” *Journal of Economic Literature*, 41, 763–787.
- Ashenfelter, Orley and Kathryn Graddy (2011), “Sale rates and price movements in art auctions.” *American Economic Review*, 101, 212–216, URL <http://www.aeaweb.org/articles?id=10.1257/aer.101.3.212>.
- Bauwens, Luc and Victor Ginsburgh (2000), “Art experts and auctions: Are pre-sale estimates unbiased and fully informative?” *Louvain Economic Review*, 66, 131–144.
- Catalini, Christian, Chris Foster, and Ramana Nanda (2018), “Machine intelligence vs. human judgement in new venture finance.” Mimeo.
- Dimson, Elroy and Christophe Spaenjers (2011), “Ex post: The investment performance of collectible stamps.” *Journal of Financial Economics*, 100, 443–458.
- Glaeser, Edward L., Michael S. Kincaid, and Naik Nikhil (2018), “Computer vision and real estate: Do looks matter and do incentives determine looks.” NBER Working Paper 25174. URL <http://www.nber.org/papers/w25174>.
- Gu, Shihao, Bryan Kelly, and Dacheng Xiu (2018), “Empirical asset pricing via machine learning.” Working paper. URL <https://ssrn.com/abstract=3159577>.
- Karayev, Sergey, Matthew Trentacoste, Helen Han, Aseem Agarwala, Trevor Darrell, Aaron Hertzmann, and Holger Winnemoeller (2013), “Recognizing image style.” *CoRR*, abs/1311.3715, URL <http://arxiv.org/abs/1311.3715>.
- Korteweg, Arthur, Roman Kräussl, and Patrick Verwijmeren (2016), “Does it pay to invest in art? A selection-corrected returns perspective.” *The Review of Financial Studies*, 29, 1007–1038.
- Lee, Yong Suk and Yuya Sasaki (2018), “Information technology in the property market.” *Information Economics and Policy*, 44, 1–7.
- Lovo, Stefano and Christophe Spaenjers (2018), “A model of trading in the art market.” *American Economic Review*, 108, 744–774.
- McAndrew, Clare, James L. Smith, and Rex Thompson (2012), “The impact of reserve prices on the perceived bias of expert appraisals of fine art.” *Journal of Applied Econometrics*, 27, 235–252.

- Mei, Jianping and Michael Moses (2005), “Vested interest and biased price estimates: Evidence from an auction market.” *The Journal of Finance*, 60, 2409–2435.
- Milgrom, Paul R. and Robert J. Weber (1982), “A theory of auctions and competitive bidding.” *Econometrica*, 50, 1089–1122.
- Renneboog, Luc and Christophe Spaenjers (2013), “Buying beauty: On prices and returns in the art market.” *Management Science*, 59, 36–53.
- Rosen, Sherwin (1974), “Hedonic prices and implicit markets: Product differentiation in pure competition.” *Journal of Political Economy*, 82, 34–55.
- Strezoski, Gjorgji and Marcel Worring (2017), “Omniart: Multi-task deep learning for artistic data analysis.” *CoRR*, abs/1708.00684, URL <http://arxiv.org/abs/1708.00684>.
- Tan, Wei Ren, Chee Seng Chan, Hernán E. Aguirre, and Kiyoshi Tanaka (2016), “Ceci n’est pas une pipe: A deep convolutional network for fine-art paintings classification.” In *2016 IEEE International Conference on Image Processing (ICIP)*, 3703–3707.
- The Economist (2018), “Tech firms disrupt the property market.” 13 September 2018.
- Velthuis, Olav (2007), *Talking Prices: Symbolic Meanings of Prices on the Market for Contemporary Art*. Princeton University Press.