



SOABER: An Innovative Approach to Maximize Bidder Surplus in Simultaneous Online Art Auctions

Mayukh Dass

&

Wolfgang Jank

&

Galit Shmueli

<http://eprints.exchange.isb.edu/267/>

Working Paper

Indian School of Business

2010

SOABER: An Innovative Approach to Maximize Bidder Surplus in Simultaneous Online Art Auctions

Mayukh Dass¹, Wolfgang Jank², Galit Shmueli²

¹Rawls College of Business, Texas Tech University

²Robert H Smith School of Business, University of Maryland

Abstract

This paper presents a novel intelligent bidding system, called SOABER (Simultaneous Online Auction BidDER), which monitors simultaneous online auctions of high-value fine art items. It supports decision making by maximizing bidders' surplus and their chances of winning an auction. One key element of the system is a dynamic forecasting model, which incorporates information about the speed of an auction's price as well as about the level of competition within and across other auctions. Other elements are a wallet estimator, which gauges bidders' willingness to pay, and a bid strategizer that embeds the forecasting model into a fully automated decision system. We illustrate the performance of our intelligent bidding system on an authentic dataset of online art auctions for Indian contemporary art. We compare our system with several simpler ad-hoc approaches and find it to be more effective in terms of extracted surplus and the resulting winning percentage.

Keywords: Bidder Surplus Optimization, Dynamic Price Forecasting, Functional Data Analysis, Intelligent Bidding System, Simultaneous Online Art Auctions.

Introduction

During the last decade, the world art market has undergone a tremendous amount of transition with reports of better returns on investment compared to stock markets. According to *Artprice.com*, art prices have seen a growth of 80% during the last decade. The increasing interest in art in general has also helped popularize emerging art markets such as Indian contemporary art, Chinese art, and Russian art. Moreover, with the current tumult of the financial markets, art markets may see an even further boost of popularity as, historically, art has prospered in times of economic hardship.

While art markets have become very popular, the barriers of entry for buyers are higher compared to traditional stock markets due to the difficulty in evaluating art pieces accurately and objectively. Every piece of art is unique and often of extremely high value, which makes pricing difficult compared to classical investments. Auctions are therefore the typical way for selling art, and they are conducted regularly by auction houses such as Christies and Sotheby's. In recent years, online auctions have become very popular in the art marketplace, and much of their popularity can be attributed to the success of eBay. According to Forrester Research Reports, online auction sales will reach \$65 billion by 2010¹, a large share of the total US auction market of \$270 billion (NAA Industry Report 2007). Reasons for the popularity of online auctions are their worldwide reach and borderless and timeless nature. The auction mechanism also generates a fair price to both seller and buyer. Moreover, many people find entertainment value in the auction mechanism.

Fine art is typically sold in a *simultaneous online auction* (SOA) setting. These art auctions are held by specialized online auction houses dedicated to selling only one type of art genre (e.g., SaffronArt.com and BidandHammer.com sell only Indian contemporary art, Attinghouse.com sells only Chinese art and Southeast Asian art) with a price tag ranging from a few thousands to a few millions of US dollars. Although these auctions are similar to those on eBay, they differ both in terms of the auction design and the types of items that they sell. SOAs sell multiple objects simultaneously in a first-price ascending auction format. This means that auctions start and end at the same time for all items. Since

¹<http://www.forrester.com>

many items are highly complementary, bidders are usually interested in purchasing more than one item, which increases competition among participants. In contrast, items on eBay are sold in a variant of the second-price sealed-bid auction (Krishna 2003) and are held independently of each other. While eBay bidders also have the opportunity to bid on more than one item at a time, competition across these auctions is typically of smaller scale since eBay bidders' typically do not compete against each other on multiple auctions at the same time. Moreover, eBay masks bidder identities (especially for high-valued items), thereby making it difficult for bidders to identify their competitors. Another difference between eBay auctions and SOAs is the type of closing rule. While eBay auctions close at a fixed time point² (e.g., exactly 7 days after the auction starts), SOAs tend to have soft-closing rules where the time automatically extends after the last bid. A soft-close auction format not only encourages bidders to bid early (Roth and Ockenfels 2002), but also discourages *sniping*³ in the last moments. Finally, SOAs are organized by a single seller, i.e., the auction house, whereas eBay provides an auction platform for many different sellers.

In online auctions, and in particular in SOAs, bidders often face information overload. Competing auctions, competing bidders, and sometimes even competing auction formats (Anthony and Jennings 2003; Haruvy et al. 2008) make the decision process extremely hard and confusing. Bidders who want to maximize their surplus and winnings are likely to be overwhelmed by the large amount of streaming bidding information on the multiple simultaneous auctions. Therefore, there is a need for data-driven decision-support mechanisms that can assist bidders in deciding on which items to bid and how much to bid on each item. Such a system must capture all of the streaming data, convert it into decisions, and update the decisions as new information becomes available. In this paper, we introduce SOABER (SOA BiddER system), an intelligent bidding system based on dynamic forecasting that can facilitate SOA bidders to maximize their surplus and optimize their bidding strategies.

Recently, there has been a growing interest in the academic community in developing intelligent bidding systems and software agents to assist bidders in making rational and optimized bidding decisions,

² Except for Buy-it-Now auctions

³ Sniping is a strategic bidding activity where bids are submitted in the last moments of the auction to allow minimal time for other bidders to react to this bid. Such behavior is prominent in eBay auctions as the auction closes promptly at a specific time.

particularly in simultaneous auctions. Most of the current research uses various machine learning methods such as fuzzy neural network (He et al. 2003, 2004, 2006), supervised learning (Stone et al. 2003) or genetic algorithms (Anthony and Jennings 2003) to forecast the final auction price. Our work differs from this body of work both in terms of the problem setting and the bidding decision process. The majority of existing studies have used data from simultaneous auctions of flight tickets, hotel rooms, and entertainment events (using the TAC Ad auction game)⁴. Compared to fine art, these items possess less heterogeneity and are relatively low in value and thus involve lower monetary risk. Hence, the first difference is the context. Second, rather than relying on a single machine learning algorithm for forecasting the final price, our proposed system combines a forecaster, a wallet estimator, and a bid-strategizer to produce suggestions as to which items to bid on, and how much to bid on each. The forecaster uses ongoing auction characteristics such as price dynamics and bidder competition to predict the final price dynamically. The wallet estimator estimates the bidders' willingness to pay, and the bid strategizer uses information from the forecaster and the wallet estimator to make the final bidding decision. For each of these components we choose models and algorithms that are practical in terms of the available information, and at the same time capture as much of the relevant information as possible. These three system components are modular and can be replaced with alternatives, to allow flexibility in other types of SOAs. We use this modularity to compare the performance of our selected components with alternate configurations.

In addition, our proposed system contributes by being the first dynamic predictive system for the art market. Although none of the extant literature on art market computation and prediction involve bid history analysis, two valuation methods, namely repeat sales index (Baumol 1986; Goetzmann 1993; Mei and Moses 2002; Pesando 1993) and hedonic price index (Chanel, Gerard-Varet and Ginsburgh 1996) are used to estimate the valuation of art items. Both these methods analyze previous sales data of artists to estimate the current value of their art works. Currently, most auction houses use one or both of these

⁴ <http://tac.eecs.umich.edu/>

methods to generate pre-auction price estimates of the art works (Artnet 2010), but not for estimating or forecasting the final price of items.

The rest of the paper is organized as follows. The next section describes the mechanism of simultaneous online art auctions and in particular that on SaffronArt.com. We describe and explore our data in section three. In section four, we introduce and describe the intelligent bidding system SOABER and its three components. Section five discusses performance evaluation. We set up several competing methods and several bidding scenarios. In section six, we apply SOABER to the data from SaffronArt.com and compare its performance to the alternative methods under the various bidding scenarios, showing its superiority. We conclude in section seven with managerial implications and future research directions.

Simultaneous Online Art Auctions

During the last decade, the global art market has gone through a tremendous amount of transition. Reports of better return on investment on art items (Mei and Moses 2002) as compared to stock markets have encouraged the establishment of many new art hedge funds and mutual funds. They have also encouraged new bidders and art enthusiasts to attend auctions with an intention to buy art items that they can enjoy and that can generate profit over time. According to *Artprice.com*, art prices have seen a growth of 80% during the last decade (1997-2007). The increasing interest in the overall art market has also helped popularize emerging art markets such as Indian contemporary art, Chinese art, and Russian art to new collectors and investors.

Another prominent change in the art marketplace is the emergence of online auction houses such as SaffronArt.com and BidandHammer.com. These auction houses sell high quality, genuine art work comparable with the art sold by the traditional auction houses such as Christie's and Sotheby's. Among all the online auction houses, Saffronart.com has emerged as one of the leading auction houses for Indian contemporary art. From 2000-2009, Saffronart.com generated \$107m (market share = 29.56%), which is more than Sotheby's \$70m (market share = 19.47%), and just second to Christie's \$114m (market share = 31.69%) during that period. They sold more art items (n=3361) than Christie's (n=1837), Sotheby's (n=

1766) or any other auction houses selling Indian contemporary art. Saffronart.com online art auctions use a simultaneous auction format, which is popular when selling a wide range of high-priced items. Examples are FCC radio spectrum auctions (Milgrom 1998), art auctions (Dass, Seymour and Reddy 2010), U.S. Treasury bill auctions (Rothkopf, Pekec and Harstad 1998) and timber and car auctions (Kwasnica and Sherstyuk 2007).

For this study, we collected the bid histories of all the auctions held on SaffronArt.com in their December 6-8, 2005 event. In this event the auctions lasted three days and 199 art items were sold. SaffronArt.com holds such 3-day auction events from three to five times a year. In such an event, all auctions start at the same date and time, but their closing times are sequential in the following way: Auctions are divided into 8 groups, with each group consisting of 20-30 auctions. After the third day, all auctions belonging to the same group close simultaneously, but each group closes sequentially, typically within 30 minutes from the previous group. This procedure gives bidders an opportunity to counter any late bids if they are competing across multiple groups of auctions. Bidders interested in participating in SaffronArt.com auctions are required to go through an elaborate registration process. Unlike on eBay, where bidders merely need a valid email address and a credit card, bidders on SaffronArt.com go through a lengthy credit check and are required to provide names of references. Such a lengthy enrollment process increases participation of serious bidders. Like Christie's or Sotheby's, SaffronArt.com examines items that are offered to them by individual owners and then sells the items on the owners' behalf. Since the complete auction is held online, the auction house arranges live viewing events prior to the auction date, where bidders are invited to attend and inspect the art items in person. This increases the legitimacy of and trust in the auction process.

Data Used in this Study

Like other online auction houses, SaffronArt.com posts detailed bid histories of items for sale on their website while the auction is in progress. The bid history includes information on each submitted bid, its time, amount and the bidder's ID. Apart from the bidding activity information, each bid history also

includes information about the item: name of the artist, physical characteristics of the item (size and media), pre-auction estimates, the item's expected value based on analyses by auction house art experts, and provenance of the item. The auction house also provides information about previous (and comparable) items sold by the same artist. Since all items are of high value, the auction house tries to provide as much information as possible in order to help bidders make rational bidding decisions. The auction house also supplies additional information about the auction format, the bidding rules, and the closing schedules.

SaffronArt.com organizes from three to five auction events per year. In each of these events, they typically sell 100-200 contemporary Indian paintings and sculptures. Although the bid history is posted online *during* the auction, it is promptly removed as soon as the auction closes and there is no way of revisiting it afterwards (unlike eBay, where information is stored for at least 15 days after an auction closes). This policy complicates the data collection task. To overcome this obstacle, we developed a java-based web agent that visits the bid history during the very last moments of the auction and captures the entire information automatically. Since the auction is extended after a late bid, our web agent is able to capture the complete bid history.

We collected all consecutive bids of auction data for the entire set of 199 art items (lots) from 70 different artists. In this auction event, which took place between December 6-8, 2005, 256 bidders participated, posting 3060 bids. The average number of bids per item was 15 and the average number of bidders participating in an auction was 6. The average value realized for all 199 items was USD \$62,065, ranging between \$3,135 and \$1,486,100. Other descriptive statistics of our data are shown in Table 1.

Please Insert Table 1 about here

SOABER (SOA Bidder System)

We propose an intelligent bidding system with the goal of maximizing bidders' surplus and their probability of winning an auction. The bidding system consists of three basic modules (see Figure 1):

1. A forecasting module,
2. A wallet estimating module,
3. A bid strategizer module.

Please Insert Figure 1 about here

Our system works as follows: Consider an individual bidder who wants to make a bidding decision. This bidding decision involves two components. First, since several items are sold in simultaneous auctions, the bidder must decide which item to bid on. After deciding on the item, the second decision is how much to bid on that item. The first decision is more complex since it involves considering both the bidder's own willingness to pay (WTP) (i.e., his/her own "WTP wallet") as well as estimating that of all competing bidders' willingness to pay (i.e., the competitors' "WTP wallets"). Rational behavior dictates that a bidder will only bid in an auction if his/ her WTP wallet is larger than the current high-bid; a bidder will also avoid auctions that are expected to close at a higher price than his/ her own WTP wallet. We estimate an auction's expected final price using the forecasting module. Moreover, a bidder will take into account the competition's ability to outbid him/ her. If a bidder suspects that a competitor's WTP wallet is much larger than his own, his chances of winning are low and, as a consequence, that item does not present a good bidding opportunity. In other words, a bidder will choose to bid on items for which competitors' WTP wallets are not (much) larger than his own WTP wallet. The second component of our bidding system therefore estimates the competitors' WTP wallets.

As for the second decision of *how much* to bid, the bidding system decides on the bid-amount using the output from the forecasting module. The bid-strategizer combines price forecasting and WTP wallet size

estimation to arrive at an automated decision process, determining which item to bid on and how much to bid.

Forecasting Module

Our intelligent bidding system relies on real-time dynamic price forecasts. Wang et al. (2008) show that dynamics in online auctions matter and that they lead to improved real-time forecasting. We therefore adopt the dynamic forecasting ideas and extend them to the SOA context by integrating additional information on within- and across-auction competition and dynamics.

Price forecasting in online auctions, and particularly in SOAs, is challenging due to their dynamic environment. One aspect of this environment is the changing bid density, where the number of bids per time unit changes dramatically throughout the auction (Roth and Ockenfels 2002; Rothkopf, Pekec and Harstad 1998; Shmueli, Russo and Jank 2007). The resulting unequally-spaced time-series of bids deem traditional forecasting models (which assume evenly spaced measurements) inadequate. Furthermore, price dynamics across different auctions follow different price paths. In other words, the speed at which the price travels during the auction and the rate at which this speed changes varies across auctions. Therefore, traditional forecasting models, which do not account for such instantaneous change, fail to accurately predict auction prices. To incorporate the dynamic nature into a forecasting model, we take a functional data modeling approach.

Functional Data Analysis (FDA) is a statistical methodology that operates on functional observations such as the price curves in online auctions. While FDA has received a lot of enthusiasm within the statistics literature, it is only slowly entering the information systems literature (Bapna, Jank and Shmueli 2008; Jank and Shmueli 2006; Reddy and Dass 2006; Steward, Darcy and Daniel 2006). In this paper, we employ FDA to capture the dynamic components of an SOA and to build a real-time forecasting model for price. The underlying idea is to represent the price path during an auction as a continuous curve. Then, following functional principles, the price curves of individual auctions are estimated using smoothing techniques (Ramsay and Silverman 2005). From the price curves, *price*

dynamics are estimated via curve derivatives. The first derivative measures *price velocity*, or how fast the price changes during the auction. The second derivative measures *price acceleration*. See Jank and Shmueli (2006) for a more detailed discussion of price dynamics. The price curves and price dynamics are subsequently incorporated into the forecasting model, in addition to other auction information, to produce real-time price forecasts.

A Basic Forecaster

The smooth price curves and their dynamics serve as the basis for our dynamic forecaster. Following the ideas of Wang et al. (2008), we capture 3 conceptually different pieces of information: static and time-varying information, the price path until time of prediction, and the price dynamics information. The dynamic forecasting model of price at time t ($y(t)$) is given as:

$$y(t) = \alpha + \sum_{i=1}^Q \beta_i x_i(t) + \sum_{j=1}^J \gamma_j D^{(j)} y(t) + \sum_{l=1}^L \eta_l y(t-l) + \varepsilon(t) \quad (1)$$

where $x_1(t), \dots, x_Q(t)$ is a set of static (pre-auction) and time-varying predictors, $D^{(j)} y(t)$ denotes the j^{th} derivative of price at time t , and $y(t-l)$ is the l^{th} price lag. In our context, the static predictors, which do not change over the course of the auction, include the opening bid, item characteristics (size of the item and type of art work), and artist characteristics (artist type, average price per sq. inch of the artist's sold items in the previous year's auction); time-varying predictors, which do change as the auction progresses, include the number of bids. A brief description of all model-components is given in Table 2.

Using equation (1), the resulting h -step ahead forecast, given information until time T , is given by:

$$\tilde{y}(T+h|T) = \hat{\alpha} + \sum_{i=1}^Q \hat{\beta}_i x_i(T+h|T) + \sum_{j=1}^J \hat{\gamma}_j \tilde{D}^{(j)} y(T+h|T) + \sum_{l=1}^L \hat{\eta}_l \tilde{y}(T+h-1|T) \quad (2)$$

The above equation (1.2) faces two challenges. First, the price dynamic components $D^{(j)} y(t)$ are coincident indicators, and therefore must be forecasted prior to their use in equation (2). We solve this by forecasting the price dynamics using a polynomial-trended linear regression model with static and time-

varying predictors and autoregressive (AR) residuals. The second challenge is that the static predictors do not change during the auction, and therefore their estimated coefficients are confounded with the price function. We solve this by transforming the static variables into time varying predictors by considering each static variable's impact on the price evolution. This is done by fitting a functional regression model of price on each of the static predictors and then using their resulting time-varying estimated coefficient as a time-varying predictor in equation (1). For more details, see Wang et al. (2008).

Please Insert Table 2 about here

Extending the basic forecaster to SOAs

As pointed out earlier, one important distinction between SOAs and traditional online auctions is the high level of competition that SOAs induce. We therefore set out to include the extent of competition into our forecasting model. To that end, we first define suitable *measures* of bidder competition. Bidder competition can manifest itself in two different forms: competition within the same auction and competition across different auctions. We thus term these two competition types as *within-auction competition* and *across-auction competition*, respectively (Dass 2008).

We gauge within-auction competition via the number of sequential bids (head-to-head rivalry) between two bidders. Although there may be other forms of bidder competitions in the auctions such as triadic competitions (competition among 3 bidders) or competition among bidder subgroups, we only consider dyadic bidder competition, which is the fundamental level of bidder competition. For every auction, we first determine the number of unique pairs of bidders (j, k) participating in that auction. Then for each of these bidder pairs, we count the number of times (n_{jk}) the two bidders bid sequentially (e.g., $A \rightarrow B \rightarrow A$). We then measure within-auction competition as the maximum of this count (across all bidder pairs). Within-auction competition (wa) for auction i can thus be written as

$$wa_i = \max(n_{jk}) \text{ for } j = 1 \cdots B_i - 1, \text{ and } k = j + 1 \cdots B_i \quad (3)$$

where B_i denotes the number of bidders in auction i , and n_{jk} is the number of sequential bids between bidder j and bidder k .

In contrast to within-auction competition, across-auction competition measures the competitive reach of a bidder pair across several auctions. Like the previous measure, we first determine the number of unique bidder pairs. Then, for each bidder pair, we count the number of auctions in which the pair is competing simultaneously. For each auction, we subsequently compute the maximum count across all pairs. Therefore, across-auction competition (aa) for auction i is given by

$$aa_i = \max(c_{l_{jk}}) \text{ for } j = 1 \cdots B_i - 1, \text{ and } k = j + 1 \cdots B_i \quad (4)$$

where B_i denotes the number of bidders in auction i and $c_{l_{jk}}$ the number of common auctions bid by bidders j and k . This across-auction competition measure uses similar underlying concept as used by Gupta (2002), but extends it to represent the competitive intensity at dyadic bidder level.

We include our two competition measures into Equation 1 as follows:

$$y(t) = \alpha + \sum_{i=1}^Q \beta_i x_i(t) + \sum_{j=1}^J \gamma_j D^{(j)} y(t) + \sum_{i=1}^Q \kappa_i w a_i(t) + \sum_{i=1}^Q \lambda_i a a_i(t) + \sum_{l=1}^L \eta_l y(t-l) + \varepsilon(t) \quad (5)$$

In the next section we discuss how we implement the dynamic forecaster within the intelligent bidding system.

Implementing the Dynamic Forecaster on SOA art data

Our process starts with first estimating the functional objects related to an auction's price path. To that end, we use monotone smoothing splines (Ramsay and Silverman 2005; Simonoff 1996), which guarantee continuous and monotonically non-decreasing price representations. See Appendix A for further details on smoothing. The smoothed price curves and their first derivatives (i.e., price velocity) for our 199 auctions are shown in Figure 2. The average price and price velocity plots show that the price formation is fast at the beginning and near the end of the auction. Also note that the average price velocity

(i.e., the rate of change in price) nearly doubles towards the end of the auction compared to the beginning of the auction.

To estimate our forecaster's coefficients, we randomly partition the data into a training set (70% or 137 auctions) and a holdout set (30% or 59 auctions), where the training set is used to estimate the model, and the holdout set is later used to measure predictive accuracy. The resulting estimated coefficients for the forecaster are shown in Table 3. Further, we test the sensitivity of our forecaster to different training and holdout set proportions by gradually moving from a 10% training-90% holdout configuration to a 90% training-10% holdout allocation. Results (Figure 2) show that our forecaster's performance is rather unaffected by the proportions of training and holdout sets. In fact, the MAPE remains almost constant for configurations 50%-50%, 60%-40%, 70%-30%, 80%-20%. and 90%-10%.

Please Insert Figure 2 and 3 about here

Please Insert Table 3 about here

As the primary goal of this paper is predictive, our main emphasis is on the forecasting performance of the models, rather than on inference. From Table 3, we see that four *static* predictors are significant: opening bid (at the 5% level), previous auction history of the artist, size of the art item, and media of art (i.e. canvas or paper work) (at the 1% level). The effect of these static predictors on the final price has already been shown in prior research. In particular, the positive effect of the opening bid reflects the direct relation between an item's value and the choice of the starting price. This is in accordance with findings from prior studies (Bajari and Hortacsu 2003; Czujack, Flores and Ginsburgh 1996). The negative effect of an artist's previous year's values on this year's price could imply that bidders are looking for "bargains" (i.e., artists that had low values in the previous year) and are willing to bid rather aggressively for them. We also find a negative effect of the size of the art works on price, which is in accordance with the findings of Czujack et al. (1996). Prior study on the effect of the art medium on the price formation process (Reddy and Dass 2006) found no significant effect of canvas on price, but interestingly, we do find a significant negative effect of the predictor.

In addition to static predictors, our model also has dynamic components such as time-varying coefficients and static variables that are transformed into time varying predictors via functional regression weighting. The impact of each of these predictors can now be assessed at different times of the auction. Figure 4 displays the time-varying *coefficient-curves* (together with associated confidence bands). Examining the different coefficients, we see that size has a significant negative effect at the beginning of the auction, but becomes positive only at the end of the auction. This indicates that smaller art objects are more expensive than bigger ones, in accordance with the findings of Czujack et al. (1996), and that bidders take this into consideration early in the auction. The auction history is found to have an initial negative effect, but becomes positive as the auction progresses. This indicates that while bidders may be shopping for “bargains” early on, as the auction comes to a close they trust more artists who have previously done well. We also find the art medium (i.e., works on canvas) to have a negative effect early in the auction, but it becomes positive as the auction progresses. Interestingly, this is different from what an earlier study found with a different auction dataset (Reddy and Dass 2006), where no significant effect was found on the price formation. Finally, we find that the opening bid has an initial positive effect that later switches to a negative effect at the auction end. This finding is similar to that of Reddy and Dass (2006) and suggests that bidders initially draw information from the opening bid, but then gradually discount this information as more signals come in from competing bidders.

Note that the effects of within-auction competition on price are significantly positive in the earlier part of the auction, but later become insignificant. This result is in line with the findings from earlier studies on bidder competition (Ariely and Simonson 2003; Dass 2008; Heyman, Orhun and Ariely 2004). We also find a negative effect of across-auction competition in the first part of the auction, in line with the findings of Gupta (2002) and Dass (2008). Our estimated model for *price velocity* is a linear regression with quadratic trend ($K=2$) and three static predictors. Figure 5 illustrates the accuracy of the price velocity prediction for four randomly selected auctions (#12, #26, #39 and #54). For each of these auctions we see that the true and forecasted price curves are close for all but auction #39, where the velocity is under-predicted.

To evaluate predictive accuracy we use the mean absolute percentage error (MAPE) computed on the holdout set. MAPE is computed as the difference between the forecasted curve and the true curve. Figures 7a and 7b compare the predictive power of our model (denoted DF) with other competing forecasting models (DF-0, DF-I) which we will discuss in later sections. Recall that we observe the first 75% (or 54 hours) of the 3-day auction, and forecast the last 25% (or 18 hours). Figure 6 shows price forecasts for the same four auctions as in Figure 5. We see that while the velocity (Figure 5) is forecasted very accurately for auction #12, its price is not. More importantly, we can see that the forecaster accomplishes dynamic and real-time forecasts, and “customizes” its forecast for each individual auction, depending on that auction’s dynamic environment.

Please Insert Figures 4, 5, 6, 7a and 7b about here

Robustness of the Dynamic Forecaster

The intelligent bidding system requires ongoing forecasts during the entire SOA duration. In the above analysis we forecast the last $T=18$ hours. We now evaluate the robustness of our results to the choice of T . Intuitively, the forecast accuracy should decrease as we forecast further into the future. To this end, we compare several values of T : forecasting only the last $T=7$ hours prior to the end of the auction (10% time remaining), the last 11 hours (15%), 14 hours (20%), 22 hours (30%) and 29 hours (40%). For each of these time periods, we proceed exactly as in Figure 6. The results are shown in Figure 8, where we also compare the robustness of our model to the competing models. We find that the predictive accuracy is nearly unaffected by T . We conclude that our forecasting model achieves very high prediction accuracy. In the next section, we explain how we can use this predictive power for making bidding decisions.

Please Insert Figure 8 about here

Wallet Estimating Module

To build an intelligent bidding system, it is important not only to determine *how much* one should bid, but also whether one bidder will have better odds of winning an auction compared to another (i.e., *which* auction to bid on). To that end, it is essential to understand one's competition. In particular, it is essential to understand the competitors' financial willingness to pay. We refer to the amount that a competitor is capable of spending as the competitors' *WTP wallet size* (or simply, *wallet*). See also Perlich and Rosset (2008).

There are two ways for gathering information about competitors' wallets: 1) participating bidders can either voluntarily provide this information at the beginning of the auction or 2) we can estimate this information from the bid history. Although the first approach is likely more accurate, it is not always feasible since individuals may view their WTP wallet as competitive advantage and may be unwilling to reveal this information. Our intelligent bidding system therefore assumes that a bidder knows his *own* willingness to pay, but that he has to estimate *competitors' WTP* from the available bid history.

We estimate competitors' WTP wallets in the following way. Using the training set, we first determine, for each bidder, the number of auctions that he participates in simultaneously. Then, for each of these simultaneous auctions, we determine the bidder's maximum bid. The sum of maximum bids across all simultaneous auctions measures a bidder's *total WTP wallet* (since he would be held responsible by the auction house for exactly that amount in case of winning all auctions where he places a bid); in similar vein, the average of maximum bids measures the bidder's *average WTP wallet*. We apply this procedure to all bidders in the training set, which yields wallet estimates of all competitors *individually*. However, in an auction, a bidder competes against all competitors simultaneously. As a result, the bidder will only win that auction if his WTP is larger than the *maximum* of all competitors' average WTP wallets; we refer to the resulting value as the *maximum competitors' auction WTP wallet*, or auction wallet, for short. The distribution of the auction wallet is shown in Figure 9.

The 'auction wallet' is the amount a bidder must match in order to have a realistic chance of winning that auction. However, the auction wallet depends on the value of the item (i.e., the maximum

competitor's willingness to pay will be larger in auctions selling more valuable items). In order to adjust the auction wallet relative to an item's value, we regress it on the pre-auction estimate. An item's pre-auction estimate is an estimate of the item's value, provided by several art experts and published by the auction house before the start of the auction. It indicates the quality of an item and bidders typically use it as a reference point to select the auctions they want to participate in. We use this regression model to predict the auction wallet for the holdout sample. Given this prediction, a bidder will only place a bid if his own WTP is not (significantly) smaller than the estimated auction wallet of all competitors.

Bid Strategizer Module

The third component in our intelligent bidding system is responsible for placing the price forecaster and the wallet estimator into a decision framework. The system first creates a priority list by ranking all auctions from lowest to highest anticipated price, and it then sequentially moves down the list. If the forecasted price of an auction is less than the bidder's own WTP *and* if the bidder's own WTP is larger than the auction wallet of all competitors, then the system will place a bid⁵. Otherwise the system will not bid in that auction and it will move on to the next auction in the list. If the system places a bid, it will bid the forecasted price. See Figure 1 for a flowchart of the process. We call this process an "intelligent strategy".

Please Insert Figure 9 about here

Bidding System Evaluation

The goal of the intelligent bidding system is to optimize bidders' use of their own WTP wallets and to increase their chances of winning an auction. We therefore evaluate our system using two key performance indicators: bidder surplus and percentage of items won. The latter indicator is directly available from the bidding system, which computes the number of auctions each bidder wins. We define a

⁵ We also explored a strategy where the bidder's WTP wallet is not highest than the maximum of the competitors' WTP wallets, but rather higher than the 75th and 90th percentiles. See Appendix B for details.

“win” if the forecasted price (i.e., the amount the system bids on behalf of the bidder) is higher than the actual realized price of the auction. Bidder surplus is computed as the difference between the forecasted price and the bidder’s average WTP wallet size; i.e., his average financial commitment per item. We examine the performance of the intelligent bidding system under four different forecasters (DF-00, DF-0, DF-I, DF), and under four types of bidding strategies (intelligent, early bidding, late bidding, and risk-averse). We describe each of these forecasters and bidding strategies below.

We assume that our bidding system makes a recommendation for a randomly selected bidder (and only for one bidder at a time) bidding on a randomly selected auction from the holdout sample. Since different bidder WTP wallets (large vs. small) can have a different impact on the performance of our bidding system, we draw a bidder’s WTP wallet randomly from the bidder WTP wallet distribution as shown in Figure 11b. This bidder now bids according to the different bidding strategies.

Alternative Forecasters

Our forecaster is one that incorporates both static and dynamic information about price, competition and other auction components. We call this forecaster DF (Dynamic Forecaster). To evaluate the effectiveness of the dynamic forecaster on the performance of the intelligent bidding system, we compare it to three simpler forecasters: The first, DF-I, is identical to our dynamic forecaster in terms of the information that it uses, except that it excludes the price dynamics component. The second forecaster, DF-0, is a regression-based forecaster that takes into account only information that is known at the start of the auction (Lucking-Reiley et al. 2007) but adds to that the price at the time of forecasting. The third forecaster, DF-00 assumes that data in the past are correlated with data in the future, but it does not make explicit use of any auction characteristics. In this case, we compute the average (final) price of all the auctions in the training set and forecast every auction from the holdout set using that number. Since our bidding system is modular, we can simply exchange, say, DF by DF-I, DF-0, or DF-00 and measure the corresponding results.

Alternative Bidding Strategies

To evaluate the usefulness of our automated and intelligent bidding system, we consider and simulate three alternative bidding strategies that are often observed in online auctions: early bidding, late bidding, and risk-averse bidding. These three strategies distinguish themselves from our approach in that they are not automatic and imitate human bidders. Early bidders are bidders who place bids at early stages of the auction. This is similar to *Early Evaluators* as classified by Bapna et al. (2004). Late bidders are the bidders who bid near the end of the auction (termed *Opportunists* by Bapna and his colleagues).

Simulating Early-bidding (NAÏVE-I): To simulate early bidding, the bidder is allowed to bid at the beginning of the auction with 90% of his average WTP wallet value (his/ her average financial commitment per item). The bidder wins the auction if the bid amount is more than the final realized price of the auction and his/ her WTP wallet value is greater than the auction wallet for the item.

Simulating Late-bidding (NAÏVE –II): To simulate late bidding, the bidder is allowed to bid after the second-last bid with an amount that is 20% more than the second last bid but less than the bidder's average WTP wallet. The bidder wins the item if the amount is more than the realized price in the auction and his/ her WTP wallet value is greater than the auction wallet for the item.

Simulating Risk-averse bidding (NAÏVE – III): Before an art auction starts, the auction house carefully prepares the event catalog which includes each item's pre-auction estimate. These estimates are prepared by art experts and indicate the current market value of an item. If a bidder is risk-averse, s/he will only bid up to that pre-auction estimate. Therefore, we simulate a risk-averse bidding strategy by allowing the bidder to bid on the item if his/ her average WTP wallet is greater than the pre-auction estimate and greater than the auction wallet for the item.

Results

We implement the three alternative forecasters and three alternative bidding strategies on the holdout sample, using a forecasting window from $T=0.76$ (approx. 18 hours before the auction end) to $T=0.99$ (approx. 43 minutes before the end). We repeat the experiment 50 times. We find that SOABER (using the DF forecaster) outperforms all other combinations both in terms of bidder surplus and percentage of items won (Figure 10).

Please Insert Figure 10 about here

Figure 9 suggests that bidder surplus tends to be higher for our intelligent bidding system (using the DF forecaster) compared to the DF-1, DF-0 and DF-00 forecasters and to the three naïve bidding strategies. Interestingly, the winning percentage of our system is very close to that of the late bidding strategy (naïve –II)⁶. This is not too surprising because late bidding is an aggressive bidding strategy aimed at maximizing wins (but not necessarily surplus) and also found to be a popular and successful strategy in other online auction formats such as in E-Bay (Bajari and Hortacsu 2003). To investigate the robustness of our decision criteria to varying bidding thresholds (i.e., change the % of competitors' WTP wallets), we modify our bidding system to allow bidders to bid if their wallet is 90%, 75% and 50% of the auction wallet. Results from this additional experiment are shown and discussed in Appendix B.

To verify the robustness of our approach, we analyzed data collected from a different auction held in May 10-12, 2006. In this auction event, 146 items from 41 artists were sold to 195 participating bidders, posting 1866 bids. The average number of bids per item was 13 and the average number of bidders participating in an auction was 6. The average value realized for all 146 items was USD \$85,787, ranging between \$4,496 and \$687,500. We found that our model DF is still the most effective forecaster

⁶ We also performed a pair-wise T-test between IBS with DF and IBS with DF-00, DF-0, DF-I, NAÏVE-I, NAÏVE-II, NAÏVE-III and found IBS with DF-I to generate significantly (C.I. = 95%) more bidder surplus at each of the forecasting windows ($t=0.76$ to $t=0.99$) than other methods. Winning percentage of DF was found not to be different than DF-I and NAÏVE-II.

in comparison to the alternative models. We also found that our approach outperformed other heuristics both in terms of bidder surplus and winning percentage. Results can be obtained from the authors upon request.

Conclusion and Future Directions

During the last two decades, the art market has become one of the most dynamic markets in the world, posting 25.4% growth annually (ArtPrice 2006). In 2007, global public auctions (which constitute only 20% of the overall art market) generated \$9.2B, which is up 43.8% compared to 2006 fine art market revenue (ArtPrice 2007). The remaining 80% of the market is dominated by private dealing and sales through art exhibitions, which are expected to result in a total market capacity of \$50B. Indian contemporary art, the context of our study, is one of the most dynamic segments of the art market with a growth of 480% during the last 10 years. As a dynamic art movement, it is now in fourth position, just behind English Pop Art (ArtPrice 2007). Indian artists have also enjoyed the sharpest inflation during the last 10 years. For example, a \$100 investment in the works of the prominent Indian artist Francis Newton Souza would have generated an average of \$7,227 in 2007. With the growing use of online auctions, the increasing art prices and the real time availability of streaming bidding data, there is an urgent need to provide art investors, collectors, and auction houses with better decision-support tools.

In this paper we propose a novel intelligent bidding system based on dynamic forecasting that supports decision making in simultaneous online art auctions. The system monitors and evaluates streaming bidding information on multiple simultaneous auctions, and suggests which items to bid and how much to bid. The system maximizes bidder surplus and the chances of winning an auction relative to alternative bidding strategies. We also point out that SOABER illustrates a flexible dynamic bidding decision framework that can be applied beyond the boundaries of art markets. For instance, such a system can be applied to eBay auctions where all types of items are auctioned. It can also be used in other simultaneous auction settings such as jewelry auctions, FCC Spectrum auctions, timber and car auctions provided bid histories are readily available. Although the system is developed for simultaneous auctions,

it can be modified for sequential auctions of high value items such as Intellectual Property Auctions. We hope that this research will encourage more advanced decision tools in other dynamic financial markets.

Appendix A: Estimating Price Curves from Observed Bid Histories

Data Pre-Processing: The first step is to recover a smooth price curve from the observed bid history. This recovery stage often needs to be preceded by data preprocessing. Let $y_i^{(j)}$ denote the bid placed at time t_{ij} . To better capture bidding activity at the beginning and near the end of the auction, we transform the bids into their log score. Next, we linearly interpolate the raw data and sample it at a common set of time points $t_i, 0 \leq t_i \leq 3$, where $i = 1, \dots, n$ in order to account for the irregular spacing of the bid arrival. Thus, each auction can be represented by a vector of equal length

$$y^{(j)} = (y_1^{(j)}, \dots, y_n^{(j)}) \quad (0.1)$$

which forms the basis for the smooth price curves.

Recovering the Underlying Price Function: To recover the underlying price curves, we use penalized monotone curves (Ramsay and Silverman 2005; Simonoff 1996), which provide both small local variation and overall smoothness. They also readily yield higher-ordered derivatives of the target price curve as desired in our case. We first start with selecting an appropriate basis function for the price dynamics. Our choice of the b-spline basis function is based on its common used in cases when the data are not periodic. Next, for every auction, we express a price function $w(t)$ as a linear combination of a basis function $\phi_k(t)$. Therefore,

$$w(t) = \sum_{k=1}^K c_k \phi_k(t) \quad (0.2)$$

where c_k is a constant and k ranges from 1 to K basis functions. Then, we fit the data by minimizing the error sum of squares by

$$SSE = \sum_{j=1}^n [y_j - f^{(j)}(t_j)]^2 \quad (0.3)$$

where y_j is the price of item j observed in time t_j . j is 1...100 in our case and $f(t)$ is the price function that fits the observed values. A roughness penalty function is imposed to measure the degree of departure from the straight line

$$PEN_m = \int [D^m f(t)]^2 dt \quad (0.4)$$

where $D^m f$, $m = 1, 2, 3, \dots$, is the m^{th} derivative of the function f . The goal is to find a function $f^{(j)}$ that minimizes the penalized residual sum of squares

$$PENSS_{\lambda,m}^{(j)} = \sum^n \left(y_i^{(j)} - f^{(j)}(t_i) \right)^2 + \lambda \times PEN_m^{(j)} \quad (0.5)$$

where the smoothing parameter λ provides the trade-off between fit $[(y_i^{(j)} - f^{(j)}(t_i))^2]$ and variability of the function (roughness) as measured by PEN_m .⁷ We use the monospline module developed by (Ramsay 2003) for minimizing $PENSS_{\lambda,m}^{(j)}$.

Appendix B: Robustness Test of the Decision Criteria

Decision Criteria	Bidder Surplus			Winning Percentage		
	Mean	Median	Standard Dev.	Mean	Median	Standard Dev.
Bidder's WTP wallet is more than Competitor's WTP wallet	\$184,243	\$178,124	\$37,806	5.92%	6.00%	2.57%
Bidder's WTP wallet is more than 90% of Competitor's WTP wallet	\$169,740	\$169,110	\$39,406	7.08%	8.00%	2.88%
Bidder's WTP wallet is more than 75% of Competitor's WTP wallet	\$149,857	\$151,528	\$40,921	9.00%	10.00%	2.99%
Bidder's WTP wallet is more than 50% of Competitor's WTP wallet	\$117,573	\$116,077	\$31,417	14.00%	14.00%	3.41%

⁷ Sensitivity tests are performed with different values of p (4, 5, 6 were used) and λ (14 different values between 0.001 to 100 were used). We find the model fit to be insensitive to different values of p and λ . However, the RMSE for the model is the lowest with $p=4$ and $\lambda = 0.1$. Thus, we use these smoothing parameters in recovering the price curves.

References

- Anthony, P. and N.R. Jennings (2003), "Developing a Bidding agent for multiple heterogeneous auctions," *ACM Transactions on Internet Technology*, 3(3), pp. 185–217.
- Ariely, D. and I. Simonson (2003), "Buying, Bidding, Playing, or Competing? Value Assessment and Decision Dynamics in Online Auctions," *Journal of Consumer Psychology*, 13 (1&2), 113-23.
- ArtPrice (2006), "2006 Art market trends," <http://img1.artprice.com/pdf/trends2006.pdf>.
- ArtPrice (2007), "2007 Art market trends," <http://img1.artprice.com/pdf/trends2007.pdf>.
- ArtPrice (2007), "Contemporary Indian art." *ArtPrice Report*.
- Artnet (2010), "Market Analyst: White Paper," *White Paper*
- Bajari, P. and A. Hortacsu (2003), "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 34 (2), 329-55.
- Bapna, R., P. Goes, A. Gupta, and Y.W. Jin (2004), "User Heterogeneity and Its Impact on Electronic Auction Market Design: An Empirical Exploration," *MIS Quarterly*, 28, 21-43.
- _____, W. Jank, and G. Shmueli (2008), "Price Formation and Its Dynamics in Online Auctions," *Decision Support Systems*, 44 (3), 641-56.
- Baumol, W. (1986), "Unnatural Value: or Art Investment as a Floating Crap Game," *American Economic Review*, May, 76, 10-14.
- Chanel, O., Louis-Andre G. V. and V. Ginsburgh (1996), "The Relevance of Hedonic Price Indices: The Case of Paintings," *Journal of Cultural Economics*, 20, 1-24.
- Czujack, C., R. Flores Jr., and V. Ginsburgh (1996), *On Long Run Co-Movements Between Paintings and Prints* (Amsterdam).
- Dass, M. (2008). "Two Essays on Simultaneous Online Auctions," PhD. Dissertation, University of Georgia.
- Dass, M., L. Seymour, and S.K. Reddy (2010), "An Investigation of Value Updating Bidders in Simultaneous Art Auctions," *Journal of Probability and Statistics*, forthcoming
- Gupta, S. (2002), "Competition and Collusion in a Government Procurement Auction Market," *Atlantic Economic Journal*, 30 (1), 13-25.
- Ghani, R. and H. Simmons (2004), "Predicting the End-Price of Online Auctions," *International Workshop on Data Mining and Adaptive Modeling Methods for Economics and Management*, (Pisa, Italy).

- Goetzmann, W.N. (1993), "Accounting for Taste: Art and the Financial Markets over Three Centuries," *American Economic Review*, 1993, 83, 1370-6.
- Haruvy, E., P. Popkowski Leszczyc, O. Carare, J. Cox, E. Greenleaf, W. Jank, S. Jap, Y-H Park, and E. Rothkopf (2008). "Competition Between Auctions," *Marketing Letters*, 3-4, 431-448.
- He, M., H.F. Leung, and N.R. Jennings (2003), "A Fuzzy Logic Based Bidding Strategy for Autonomous Agents in Continuous Double Auctions," *IEEE Trans. Knowledge Data Engineering*, 15 (6), 1345-1363.
- _____ and N.R. Jennings (2004), "Designing a Successful Trading Agent: A Fuzzy Set Approach," *IEEE Trans. Fuzzy Syst.* 12 (3), 389-410.
- _____, N.R. Jennings and A. Praugel-Bennett (2006), "A Heuristic Bidding Strategy for Buying Multiple Goods in Multiple English Auctions," *ACM Transactions on Internet Technology*, 6 (4), 465-496.
- Heyman, J.E., Y. Orhun, and D. Ariely (2004), "Auction Fever: The Effect of Opponents and Quasi Endowment on Product Valuations," *Journal of Interactive Marketing*, 18 (4), 8-21.
- Jank, W. and G. Shmueli (2006), "Functional Data Analysis in Electronic Commerce Research," *Statistical Science*, 21 (2), 155-66.
- _____, and S. Zhang (2008), "An Automated and Data-Driven Bidding Strategy for Online Auctions," Working paper, University of Maryland.
- Krishna, V. (2002), *Auction Theory*, San Diego: Academic Press .
- Kwasnica, A.M. and K. Sherstyuk (2007), "Collusion and Equilibrium Selection in Auctions," *The Economic Journal*, 117, 120-45.
- Lucking-Reiley, D., D. Bryan, N. Prasad, and D. Reeves (2007), "Pennies from eBay: the Determinants of Price in Online Auctions," *Journal of Industrial Economics*, 55 (2), 223-33.
- Mei, J. and M. Moses (2002), "Art as an Investment and the Underperformance of Masterpieces," *American Economic Review*, 92 (5), 1656-68.
- Milgrom, P. (1998), "Game Theory and the Spectrum Auctions," *European Economic Review*, 42 (3-5), 71-78.
- NAA Industry Report (2007), National Auctioneers Association Report, Online Resource, <http://www.auctioneers.org> .
- Perlich, C. and S. Rosset (2008), "Quantile Modeling for Wallet Estimation," *Statistical Methods in eCommerce Research*, Jank & Shmueli, Ed.: Wiley.
- Pesando, J.E. (1993), "Art as an Investment: The Market for Modern Prints" *American Economic Review*, December, 83, pp. 1075-1089.
- Ramsay, J.O. (2003), Matlab, R, and S-PLUS Functions for Functional Data Analysis, Online Resource, <ftp://ego.psych.mcgill.ca/pub/ramsay/FDAfuns>.

- _____ and B.W. Silverman (2005), *Functional Data Analysis*, 2nd ed., New York: Springer Verlag.
- Reddy, S.K. and M. Dass (2006), “Modeling Online Art Auction Dynamics using Functional Data Analysis,” *Statistical Science*, 21(2), 179-193.
- Roth, A.E. and A. Ockenfels (2002), Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” *American Economic Review*, 92 (4), 1093-103.
- Rothkopf, M.H., A. Pekec and R. M. Harstad (1998), “Computationally Manageable Combinational Auctions,” *Management Science*, 44 (8), 1131-47.
- Russo, R.P., G. Shmueli and N.D. Shyamalkumar (2008), “Models of Bidder Activity consistent with Self-similar Bid Arrivals,” *Statistical Methods in eCommerce Research*, Jank & Shmueli, Ed.: Wiley.
- Shmueli, G. and W. Jank (2008), “Modeling Dynamics in Online Auctions: A Modern Statistical Approach, in: Kauffman, R., and Tallon, P. Economics,” *Information Systems and Electronic Commerce: Empirical Research*, M.E. Sharpe, Armonk, NY, ISBN: 978-0-7656-1532-9.
- _____, R.P. Russo and W. Jank (2007), “The BARISTA: A Model for Bid Arrivals in Online Auctions,” *Annals of Applied Statistics*, 1(2), 412-441.
- Simonoff, J.S. (1996), *Smoothing Methods in Statistics*, 1st ed., New York: Springer-Verlag.
- Stewart, K.J., D.P. Darcy and S.L. Daniel (2006), “Opportunities and Challenges Applying Functional Data Analysis to the Study of Open Source Software Evolution,” *Statistical Science*, 21 (2), 167-178.
- Stone, P., R.E. Schapire, M.L. Littman, J.A. Csirik and D. McAllester (2003), “Decision-Theoretic Bidding based on Learned Density Models in Simultaneous, Interacting Auctions,” *Journal of Artificial Intelligence Research*, 19, 209-242.
- Wang, S., W. Jank and G. Shmueli (2008), “Explaining and Forecasting Online Auction Prices and their Dynamics using Functional Data Analysis,” *Journal of Business and Economic Statistics*, 26 (3), 144-160.

Tables:

	Mean (SD)	Median	Min.	Max.
No. of Unique Bidders/ Lot	6.35 (2.47)	6	2	14
No. of Unique Lots Bid / Bidder	4.93 (7.95)	3	1	65

No. of Bids/lot	15.47 (7.46)	15	2	48
Opening Bid in \$	\$19,343 (\$36,663)	\$6,400	\$650	\$300,000
Bid in \$/lot	\$34,824 (10,957)	\$11,576	\$650	\$1,486,100
Pre-Auction Low Estimates of the Lots	\$24,128 (45,747)	\$8,000	\$795	\$375,000
Pre-Auction High Estimates of the Lots	\$31,065 (60,351)	\$10,230	\$1,025	\$475,000
Realized Value of the Lots in USD(\$)	\$62,065 (133,198)	\$22,000	\$3,135	\$1,486,100
Realized Sq. Inch Price of the Lots in USD(\$)/ Sq. Inch	\$108.77 (225.49)	\$45.12	\$1.40	\$1,865.42

Table 1: Summary Data Description

Category	Covariates	Description
Static Predictors		
<i>1</i>	Opening Bid	Opening bid is the first bid in the auction.
	Size of the Item	It is the dimension of the artwork in square area.
	Type of Artwork	Artworks can be categorized into works on paper and works in canvas. We used an indicator variable in our model to indicate whether the item is a canvas work or not.
	Artist Reputation	The artists are categorized into established artists and emerging artists.
	Previous Auction History	The price/sq. inch of the artworks of the artists in the previous years
Current Price		
<i>2</i>	Current and previous price (price path)	Price lags at time t , $t-1$...

Price Dynamics		
3	Price Velocity	First derivative of price at time t
Competition		
4	Within-auction competition	This indicates the current level of within-auction competition in the auction.
	Across-auction competition	This indicates the current level of across-auction competition in the auction.

Table 2: Predictors Used in the Model

Types of Predictors	Covariates	Coefficients (Std. Error)
Static Predictors	Opening Bid	0.002* (0.0012)
	Previous Auction History of the artist	-0.0027** (0.0002)
	Size	-0.0012** (0.0023)
	Canvas	-0.0004** (0.0008)

** Significant at 0.01 level

* Significant at 0.05 level

Table 3: Parameter Estimates of the Models in the Training Set

Figures

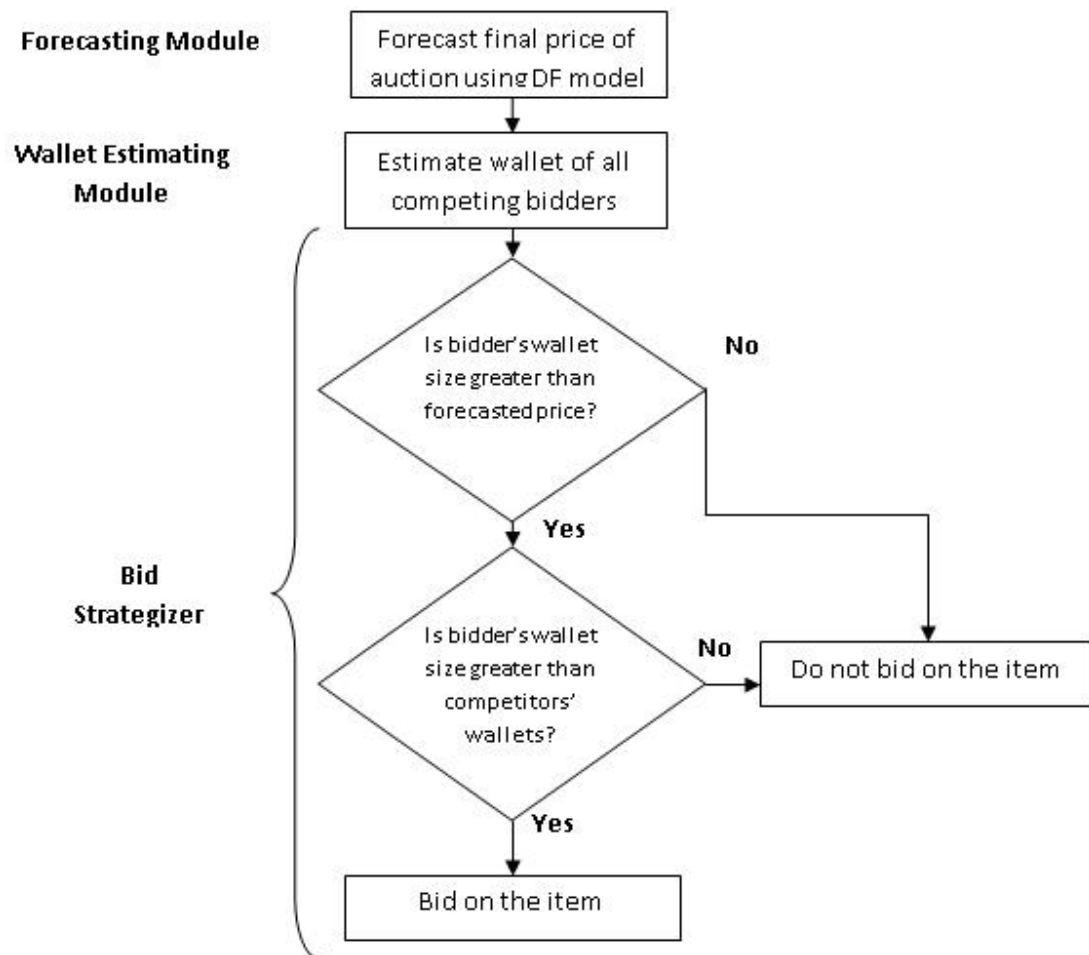


Figure 1: SOABER Framework

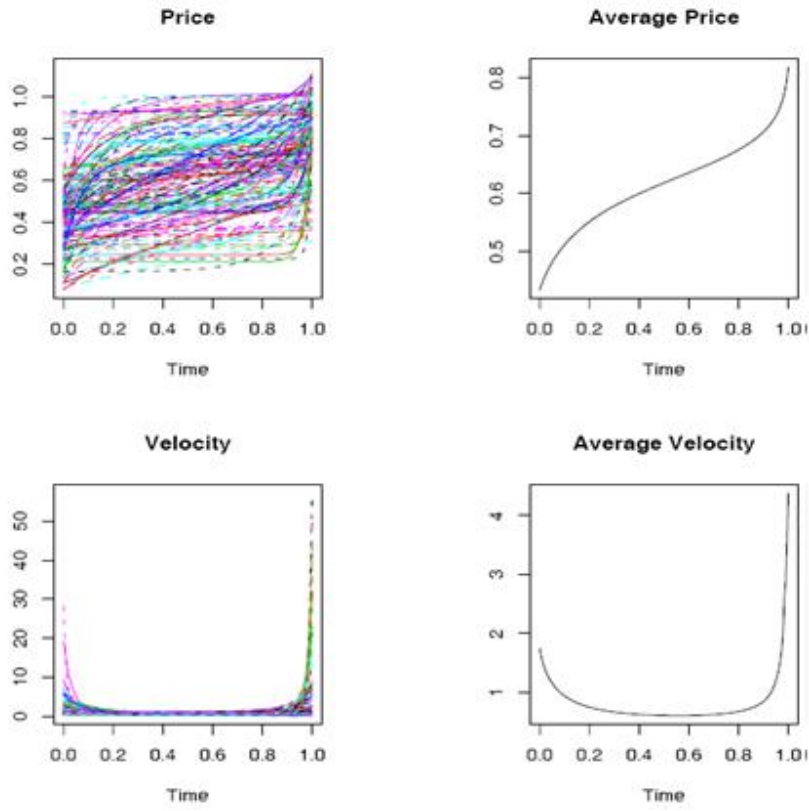


Figure 2: Price Dynamics for 199 Lots Sold in the Online Art Auction

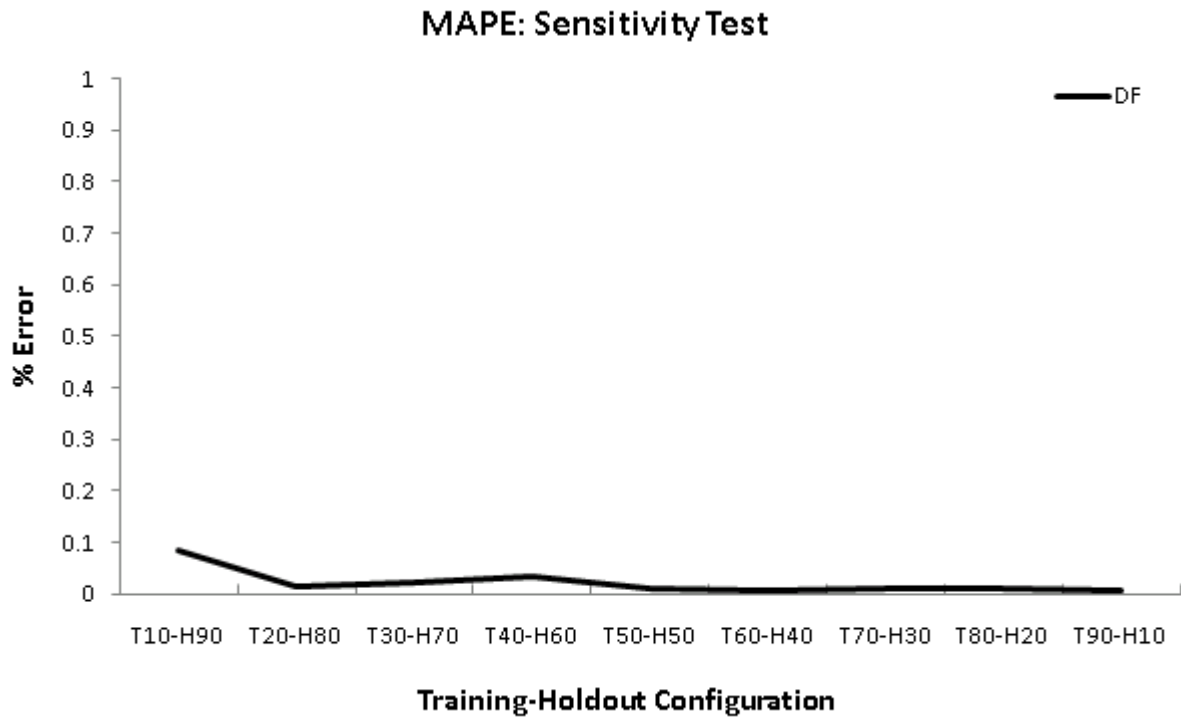


Figure 3: Sensitivity test of the Forecaster

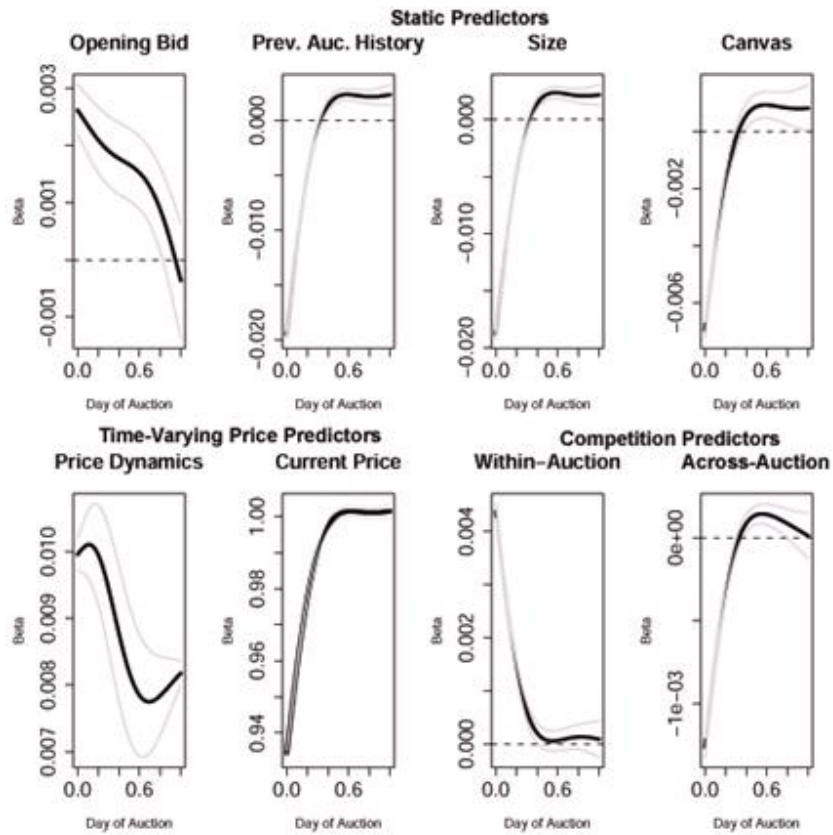


Figure 4: Parameter Estimates of Time-varying Covariates in the Model with Training Set

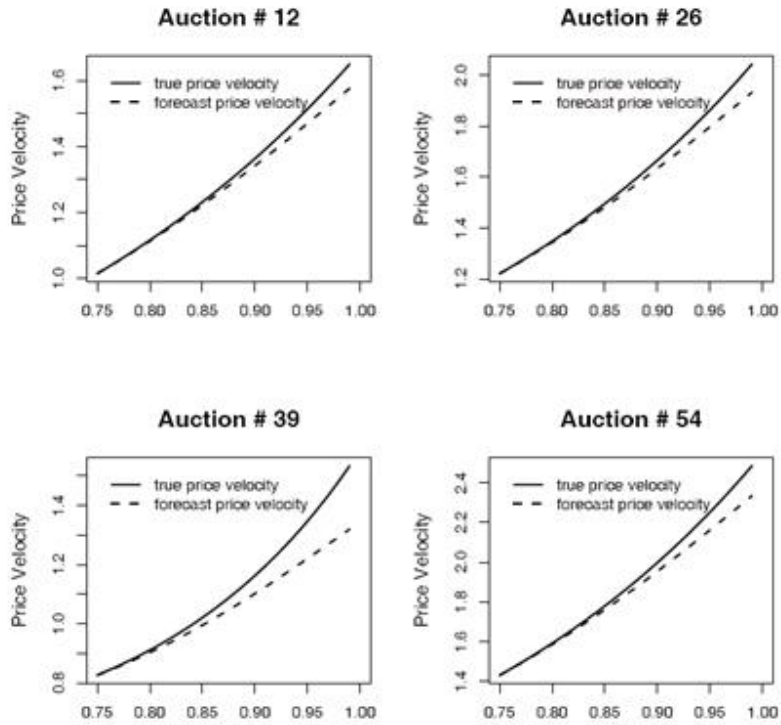


Figure 5: Performance of Forecasting Price Dynamics of the Last 18 Hours for Four Sample Lots using DF

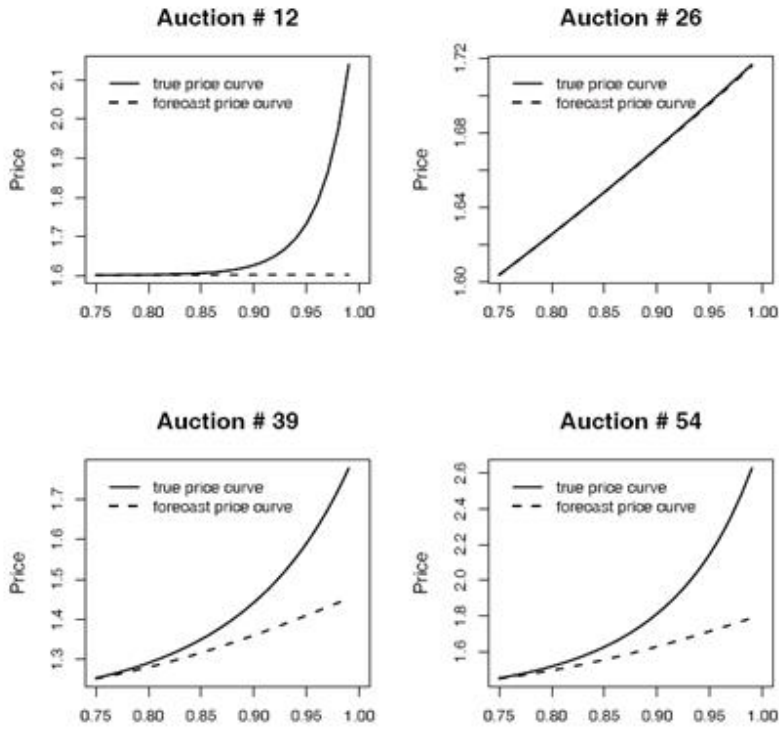


Figure 6: Dynamic Forecasting of the Last 18 Hours for Four Sample Auctions using DF




Figure 7a: MAPE of DF Model




Figure 7b: MAPE of Competing Models

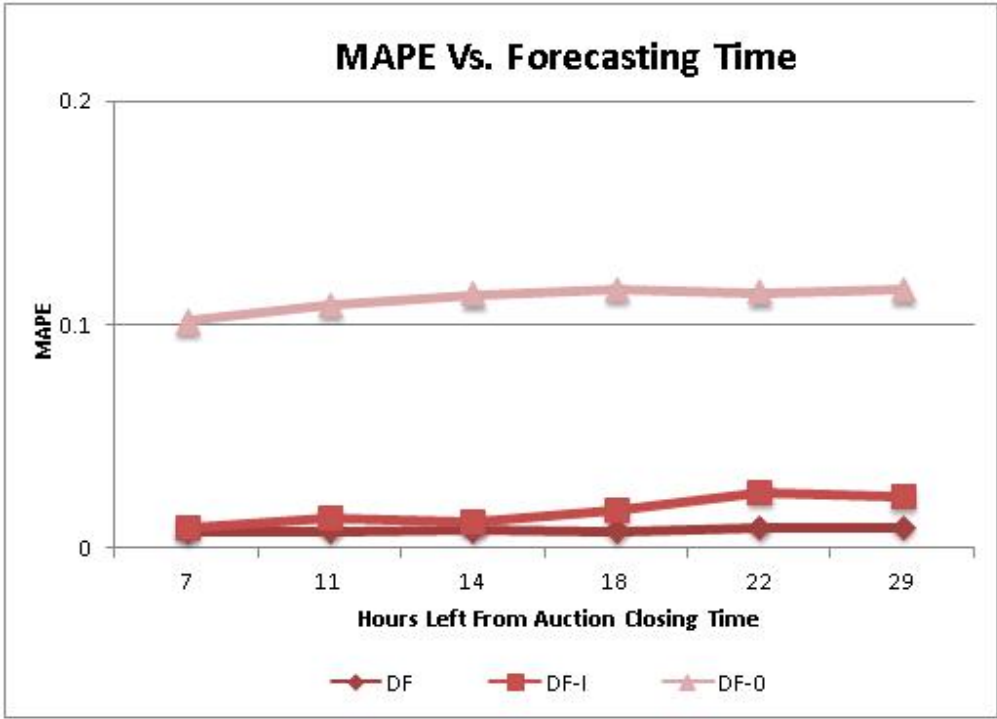


Figure 8: Comparison of Final Price Forecasting Models at different Forecasting Time

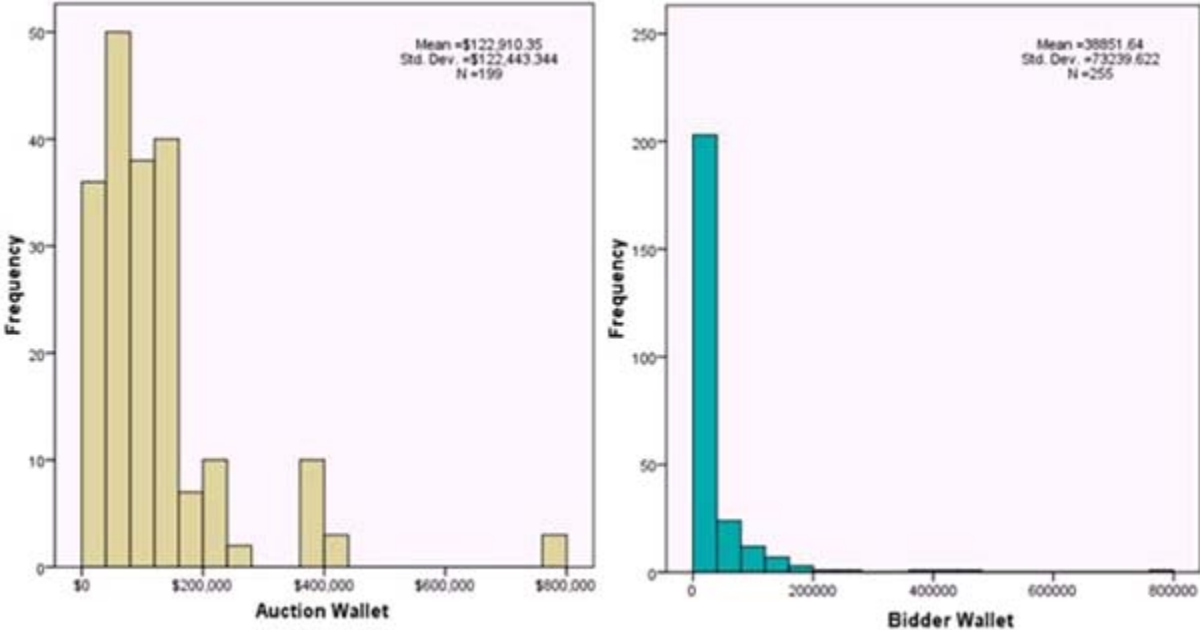


Figure 9: Auction Wallet and Bidder Wallet Distributions

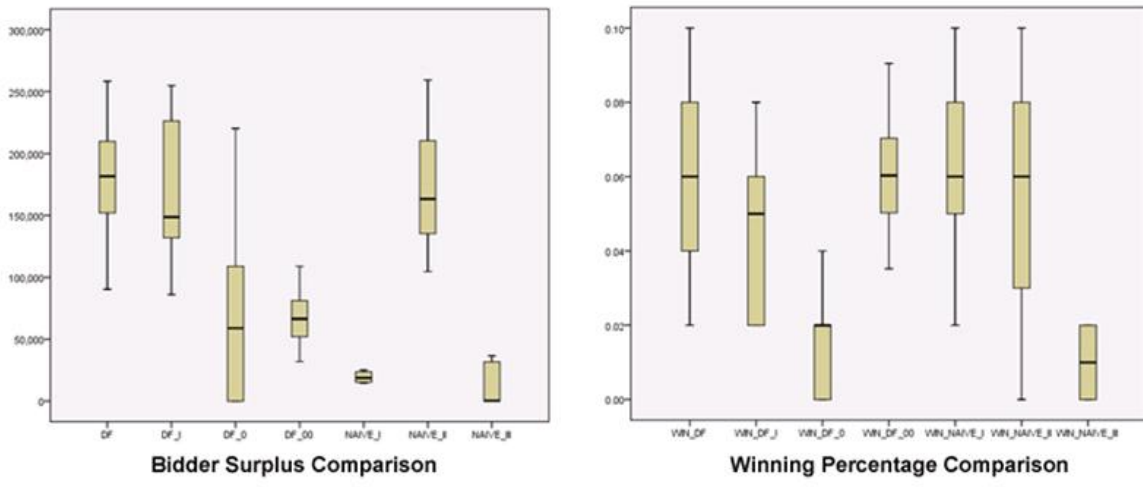


Figure 10: Bidder Surplus and Winning Percentage Comparison