E-COMMERCE

Investigating Concurrency in Online Auctions Through Visualization

Valerie HYDE, Wolfgang JANK, and Galit SHMUELI

This article presents graphical techniques for visualizing concurrency in online auctions. These include *rug plots*, which allow for a compact view of many simultaneous auctions while preserving the structure of individual auctions. We also use box plots, moving statistics plots, and autocorrelation plots, supplemented by statistical tests. Together, these are used to study synchronous events and to surmise trends in the data, as well as to raise new research questions. We illustrate our methods on data from *eBay.com*.

KEY WORDS: Electronic commerce; Functional data; Rug plot; Simultaneous events; Smoothing; Visualization.

1. INTRODUCTION

Concurrency of events is omnipresent in today's commerce, both within the electronic commerce world and even more so if we consider the online and offline domains together. For example, a consumer purchasing a luxury item such as a digital camera can gather information about price from a variety of different sources: by reading a store's sales circular, by checking the price at online retailers, by using comparison-shopping Web sites such as *bizrate.com* that search across many online retailers, by reading newspaper/online sales postings for used cameras by individuals (e.g., *craigslist.com*), by checking online auction sites such as *eBay.com*, and more. The availability of items from multi-channel sources gives consumers the power to compare prices and other related information (such as shipping-time and -cost, trustworthiness of the seller, travel time, and physical inspection). This also means that sources compete with one another and therefore, more than likely, also influence each other. There has been a growing literature investigating how these different channels affect each other, their sales, their prices, and so on (e.g., Etzion, Pinker, and Seiderman 2004; Gallien 2002; Vakrat and Seidmann 1999).

This article focuses on concurrency within a single channel, namely, online auctions on *eBay.com*. Online auctions are different from traditional brick-and-mortar auctions in that they occur simultaneously or within close temporal proximity. Online auctions tend to be longer, ranging over several days rather than minutes. Further, online auctions are not limited by the geographic barriers of traditional auctions. Sellers and bidders can be located in different parts of the country, or even the world, and still conduct business. The wide and growing popularity of online auctions makes the problem of concurrency an interesting and important topic for research.

A great majority of the literature that analyzes online auction data assumes *independence* across auctions. This assumption is typically made for reasons of simplicity and convenience, while in reality auctions for the same item, competing items, or even related (substitute) items will influence each other especially if they take place within a close time frame. On *eBay.com*, an identical product is often sold in numerous simultaneous auctions. Although each auction contains a replicate of the product, the resulting sales, prices, and even the number and level of bids during the ongoing auctions are clearly not independent of each other. Empirical research has mostly been based on auctions for the same item [e.g., mint condition Indian-head pennies (Lucking-Reiley, Bryan, Prasad, and Reeves 2005), coins (Bajari and Hortacsu 2003), personal digital assistants (Jank and Shmueli 2005; Shmueli and Jank 2005; Ghani and Simmons 2004), and rare coin auctions (Kauffman and Wood 2005)], that have closed during a certain time period (typically within a few of months). In such situations it is likely that there is dependence between the auctions because buyers have the option to select which of the competing items to bid on, and sellers have the option to decide when to post their item for sale using information on similar previously sold items.

Auction theory is mainly concerned with a single auction, and there has not been much theoretical research on concurrent auctions, even in the offline context. One article that does consider simultaneous auctions is Guerre, Perrigne, and Vuong (2000), who examined the underlying distribution of a bidder's private value in sealed-bid first-price offline auctions.

Prior empirical research concerned with the interplay of auctions has focused mainly on *sequential* classical (offline) auctions. That is, auctions that occur one-after-the-other, but not simultaneously. Sequential auction research is concerned with

Valerie Hyde is a Doctoral Student, Applied Math and Scientific Computation Program, University of Maryland, College Park, MD 20742. Wolfgang Jank is Assistant Professor, Department of Decision and Information Technologies, Robert H. Smith School of Business, University of Maryland, College Park, MD 20742. Galit Shmueli is Assistant Professor, Department of Decision and Information Technologies, Robert H. Smith School of Business, University of Maryland, College Park, MD 20742. Author names are listed in alphabetical order. Please address all correspondence to Galit Shmueli (E-mail: *gshmueli@rhsmith.umd.edu*).

 DOI: 10.1198/000313006X124163

price and quality of selection as one auction succeeds another auction.

Allen and Swisher (2000) found that auctions occurring later in an auction sequence tend to fetch a higher price than those in the beginning probably due to the limited supply at the end of a sequential auction. Deltas (1999) examined sequential ordering across numerous auctions for cattle, where it is common that higher value lots are sold at the beginning of an auction. Deltas hence found that prices decline throughout an auction, but also found a different rate of decline for auctions of different size: prices decline faster in auctions where only a small number of lots are sold.

While the insight gleaned from these analyses is useful in the traditional auction setting, it is scarcely applicable to the online context. Online auctions present new and challenging questions because of the prevalence of simultaneous events. Even within this particular universe, concurrency is prevalent, and it is challenging to evaluate its magnitude and effects. There have been only few attempts at quantifying concurrency in online auctions. Snir (in press) shows that the expected selling price in S sequential auctions is equivalent to the Sth order statistic. Although that work is primarily concerned with the final price, we are also interested in the effect of concurrency on the *process* of bidding. Zeithammer (2005) examined sequential auctions on eBay, and finds that bidders deflate their bids based on the future expected surplus. In particular, when bidders expect more auctions in the future, bids decrease. While the work does not examine price trends over time, it examines how information about the near future affects the bidding equilibrium. Specifically, it investigates whether it is better to bid high today and lose potential surplus tomorrow or bid cautiously today in hopes of surplus tomorrow. Finally, the recent article by Anwar, McMillan, and Zheng (2006) found that bidders on eBay typically place bids on multiple competitive auctions.

We are interested in the effect of concurrency not only on the final price of an auction but also on the relationship between the current bid-levels and high-bids in simultaneous ongoing auctions. Moreover, very few visualization methods have been developed for the special data challenges arriving in online auctions. Exceptions are the work of Jank, Shmueli, Plaisant, and Shneiderman (2006), Shmueli and Jank (2005), and Aris et al. (2005). Although we concentrate on online auctions, our proposed methods can be generalized to other applications where one is interested in studying the effect of concurrency on the final result of an event and also on the *event-evolution* itself. Our work adds to this line of research by specifically focusing on the aspect of concurrency and its impact in online auctions.

The goals of this article are two-fold: to study concurrent events via statistical exploratory methods (mainly visualization), and second, to develop new research questions based on these visualizations. The article is organized as follows: Section 2 discusses the characteristics of online auction data available on *eBay.com* and the set of data that we use throughout this work. Section 3 proposes several new visualizations and adaptations of classic graphical displays for capturing concurrency. Section 4 concludes the article.

2. DATA FROM ONLINE AUCTIONS

The number of different online auction sites is growing steadily. Despite different formats and rules, there is a common data-structure that can be found across most sites. This structure comprises of a time series that describes the bids placed over time (the bid history), and an associated set of features that describe the auction-setting, such as the seller rating, the auction duration, and the item category. We refer to these features as the auction attributes. Figure 1 shows an abridged snapshot of a closed auction from *eBay.com*, providing the auction attributes (top) and the bid history (bottom). We see that this was a five-day auction for a Palm M515 personal digital assistant (a handheld organizer). The seller, *oh-snap!*, has a feedback rating of 15 with 100% positive feedback. The closing price is \$144.50, and there are a total of 25 bids. This combination of information available to the bidder and seller suggests that concurrency can affect the bid histories (a bidder who places a bid on one auction can influence the bid history of other auctions) and/or the attributes (a seller uses information from previous or simultaneous auctions to decide on the shipping cost, the starting bid, or the reserve price). Figure 2 shows a snapshot of an eBay Web page that displays all open auctions closing within the next two days. Notice that the soonest closing auction (in 6 hours and 55 minutes) has only received three bids, and the highest bid is \$41. This is surprising since two other auctions closing the following day already received more bids and also a higher current price. One reason for this could be the lack of a picture on the display page which attracts more attention; another reason could be the less detailed product description.

2.1 Palm Pilot Data

Our data contain information on 236 closed auctions for a new Palm Pilot M515. The data were collected between March 11, 2003, and April 20, 2003, roughly a year after the Palm M515 was released to the market. EBay auctions last 3, 5, 7, or 10 days, depending on the length set by the seller. The most popular duration is seven days, as can be seen in Table 1. All auctions in our data resulted in a sale.

Even though the Palm Pilot has a known market value (\$250 at the time of the analysis), auctions do not always close near this value. Low prices can result, for instance, if the box is already open or if the seller has a questionable reputation. Auctions can close high if something special is offered with the product such as an accessory or free shipping. Prices also vary because bidders often get caught up in the excitement of bidding ("auction fever") and pay more than would be expected. The average selling price for all Palm Pilots in our data is \$234.50 with a median of \$234.50 and standard deviation of \$20.83. The least expensive Palm Pilot sold for \$172.50 and the most expensive auction closed at \$290. Table 1 provides descriptive statistics for all auctions.

To understand the structure of bid-history data, it is necessary to understand the auction rules and bidding mechanism. On eBay, the majority of auctions are second-price auctions which means that the winner is the bidder who placed the highest bid, and he or she pays the second highest price. Furthermore, eBay uses a so-called "proxy bidding" system where bidders place the highest value that they are willing

eBay.com Bid History for
Palm Pilot M515 hand held PC PDA HP Dell Sony wireless (Item # 5787366894)
Listed in category: Consumer Electronics > PDAs/Handheld PCs > Handheld Units

Winning bid: US $144.50
Ended: Jul-10-05 19:04:26 PDT
Start time: Jul-05-05 19:04:26 PDT
History:25 bids (US $14.95 starting bid)

Winning bidder: eduardopt (34)

Seller: oh-snap! (15)
Feedback Score: 15
Positive Feedback: 100%
Member since Dec-18-02 in United States
Item location: Columbus, Georgia

User ID	Bid Amount	Date of bid
eduardopt (34)	US $144.50	Jul-10-05 19:04:17 PDT
mongopoo (6)	US $142.00	Jul-09-05 12:46:58 PDT
prunoiu1970 (8)	US $130.00	Jul-09-05 11:37:09 PDT
eduardopt (34)	US $126.00	Jul-08-05 14:53:57 PDT
prunoiu1970 (8)	US $120.00	Jul-09-05 11:36:59 PDT
prunoiu1970 (8)	US $110.00	Jul-08-05 03:47:51 PDT
boxersmelinda2005 (5)	US $100.00	Jul-07-05 19:59:44 PDT
prunoiu1970 (8)	US $100.00	Jul-08-05 03:47:38 PDT
eduardopt (34)	US $96.00	Jul-07-05 18:45:01 PDT
boxersmelinda2005 (5)	US $90.00	Jul-07-05 19:41:19 PDT
boxersmelinda2005 (5)	US $85.00	Jul-07-05 19:40:32 PDT
boxersmelinda2005 (5)	US $80.00	Jul-07-05 19:38:58 PDT
boxersmelinda2005 (5)	US $76.00	Jul-07-05 19:38:42 PDT
boxersmelinda2005 (5)	US $72.00	Jul-07-05 19:38:17 PDT
boxersmelinda2005 (5)	US $70.00	Jul-07-05 19:37:46 PDT
boxersmelinda2005 (5)	US $60.00	Jul-07-05 16:28:48 PDT
shivasutra (2)	US $50.01	Jul-06-05 17:54:59 PDT
boxersmelinda2005 (5)	US $50.00	Jul-07-05 16:28:24 PDT
boxersmelinda2005 (5)	US $48.00	Jul-07-05 16:27:53 PDT
boxersmelinda2005 (5)	US $45.00	Jul-07-05 16:26:50 PDT
tmlcfmat (83)	US $25.00	Jul-06-05 12:48:35 PDT
shivasutra (2)	US $21.99	Jul-06-05 17:54:30 PDT
shivasutra (2)	US $18.88	Jul-06-05 17:54:15 PDT
shivasutra (2)	US $16.88	Jul-06-05 17:53:58 PDT
shivasutra (2)	US $14.95	Jul-05-05 20:31:03 PDT

Figure 1. Time series and attributes for a Palm Pilot M515 auction. Notice that bids are arranged in descending order by bid-amount. This order, however, does not reflect the arrival of the bids. Rather, it reflects the current auction high bid. These materials have been reproduced with the permission of eBay, Inc. Copyright 2006 eBay, Inc. All rights reserved.

to pay, and then eBay bids on their behalf by increasing the current price by only an increment. [For further details see *http://pages.ebay.com/help/buy/proxy-bidding.html*.]

2.2 Auction Price Evolution as a Continuous Curve

For studying concurrency of prices in competing auctions, we consider as the observation the complete price evolution that takes place during each auction. The first step is therefore to represent/estimate this evolution from the discrete bid data. Bids in online auctions are placed at varying time points. Typically, there is some bidding activity at the auction start, followed by a period of very little activity, cumulating in a surge of bidding at the very end of the auction (Shmueli et al. 2004). This last-moment bidding is often referred to as "sniping" (Bajari and Hortacsu 2003; Roth and Ockenfels 2002). The resulting bid histories are therefore time series that are unevenly spaced with sometimes very sparse and other times very dense areas. Although we could simply "connect the dots" to obtain the price of the auction at any given time, this would overfit the data (i.e., model the noise), thereby not providing a good representation of the bidding behavior. From a visualization point of view, a series of discrete unevenly spaced bids loses the temporal nature (i.e., a scatterplot of bid value vs. bid time) unlike a single smooth curve. Furthermore, such plots do not scale to multiple auctions.

An alternative is to represent the price as a continuous smooth curve. This type of curve representation is prevalent in functional data analysis (Ramsey and Silverman 2005), an approach where the observations of interest can be any type of continuous rep-

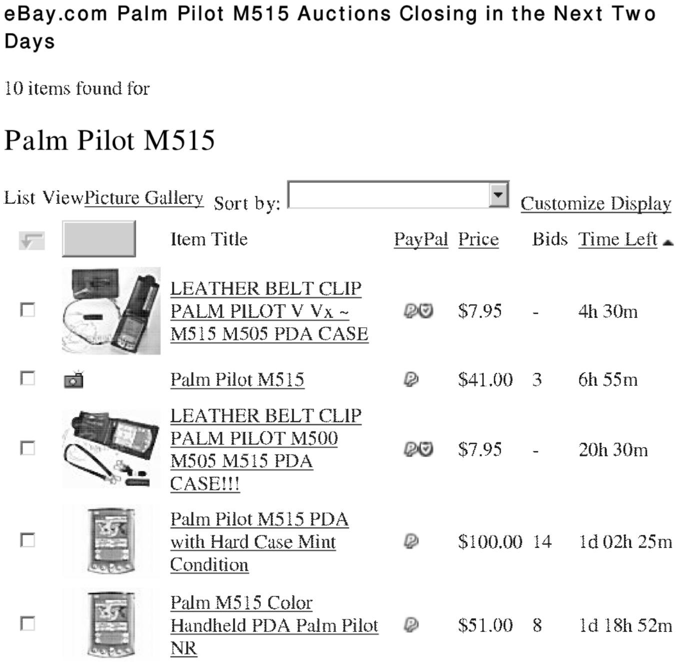

eBay.com Palm Pilot M515 Auctions Closing in the Next Two Days

10 items found for

Palm Pilot M515

List View Picture Gallery Sort by: Customize Display

	Item Title	PayPal	Price	Bids	Time Left
	LEATHER BELT CLIP PALM PILOT V Vx ~ M515 M505 PDA CASE		$7.95	-	4h 30m
	Palm Pilot M515		$41.00	3	6h 55m
	LEATHER BELT CLIP PALM PILOT M500 M505 M515 PDA CASE!!!		$7.95	-	20h 30m
	Palm Pilot M515 PDA with Hard Case Mint Condition		$100.00	14	1d 02h 25m
	Palm M515 Color Handheld PDA Palm Pilot NR		$51.00	8	1d 18h 52m

Figure 2. List of Palm Pilot M515 auctions closing within the next two days. Note that some of the auctions are for accessories, not the actual Palm M515 product. These materials have been reproduced with the permission of eBay, Inc. Copyright 2006 eBay, Inc. All rights reserved.

resentation, not merely a vector. Obtaining a continuous curve from discrete data is typically achieved through smoothing techniques. We use monotone smoothing splines for their flexibility and suitability in the auction context: since the bidding process is necessarily nondecreasing, the monotone smoothing splines guarantee monotone nondecreasing price curves. Our choice of knots is governed by bidding frequencies, such that more knots are present during the “active period” at the very end of the auction and much fewer during the “quiet period” in the middle of the auction [for further details on the choice of knots and smoothing parameter see Jank and Shmueli (2005)]. Figure 3 shows the results for three sample auctions. Even though the bid patterns in each of the auctions is very different, the same family of monotone splines (using the same knots, smoothing parameter, and polynomial order as in Figure 3) leads to reasonable approximating curves in each case, and avoids adding an extra source of variability across curves. The advantage of the curve representation is that it captures the complete price evolution in a more compact and easier to visualize way than the raw bid data. Furthermore, an appealing feature of smooth

Table 1. Descriptive Statistic for the Palm Pilot Auctions. The bottom row gives statistics for all auctions, the other rows break statistics up by auction length.

Data	Count	Minimum	Median	Mean (std)	Maximum
3 Day	43	177.50	242.50	239.20 (24.23)	290.00
5 Day	45	183.50	235.20	233.70 (22.95)	280.00
7 Day	144	186.50	233.50	233.60 (17.87)	283.50
10 Day	4	172.50	223.80	222.50 (47.30)	270.00
All	236	172.50	234.50	234.50 (20.83)	290.00

curves is that we can gauge their derivatives in order to learn about price dynamics: the price-velocity (first derivative) and price-acceleration (second derivative). We find the dynamics to be an important manifestation of the concurrency effect.

3. VISUALIZING CONCURRENT AUCTIONS

3.1 Rug Plots

Our first goal is to visualize a set of auctions by describing their entire price evolution for the purpose of exploring similarities and patterns in price evolutions of concurrent (or partially concurrent) auctions. One step in this direction is the *Calendar of Auctions* plot (Shmueli and Jank 2005) for displaying a sample of auctions over the data-collection period. Although it preserves chronological information, it does not preserve the bid history information. We propose a graphical method that expands upon the Calendar of Auctions plot, which displays the entire price-evolution curves or dynamics curves over calendar time. Because of its rug-like appearance we name these displays *rug plots*. Figure 4 displays the price curves and Figure 5 displays the price-velocity curves. Since the end of an auction is of special interest, we emphasize the endpoint of each curve by a darkened dot. Viewing rug plots at multiple resolutions allows seeing an overview but also local details. For instance, the right panels of Figures 4 and 5 display a subset of all curves, zoomed in between March 29, 2004, and April 8, 2004.

The price-velocity curve, which is the first derivative of the price-evolution curve, is always nonnegative (because the price curve is monotone). A zero velocity occurs when the price does not increase, representing a period of bidding inactivity. A small positive velocity means that the price slowly increases at that instant. Similarly, high velocity corresponds to rapid price increases.

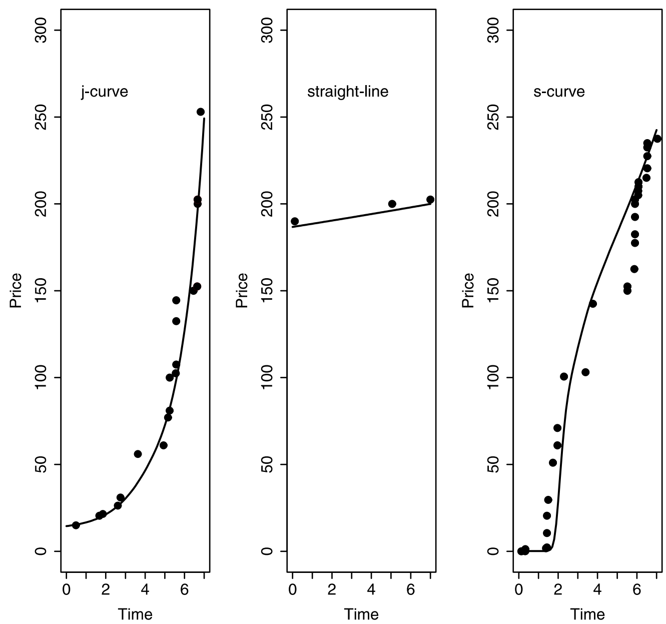

Figure 3. Bids and fitted price curve for three different auctions. Knots are placed at day (0, 1, 2, 3, 4, 5, 6, 6.25, 6.50, 6.75, 6.8125, 6.8750, 6.9375, 7). The smoothing parameter is λ = 0.01, and the polynomial order is 5.

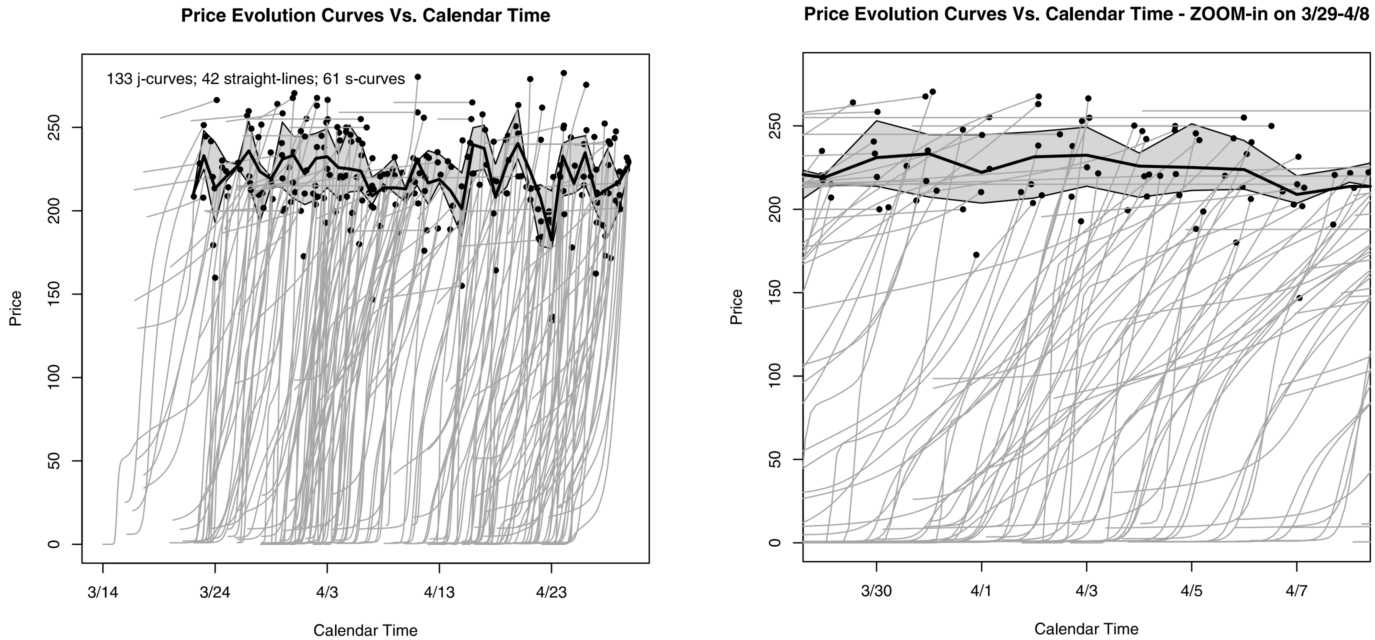

Figure 4. Price evolution over calendar time (left) and zoomed in to March 29-April 8 region (right). Thick curve and gray band are the daily median and IQR closing price.

The next step is to see what types of concurrency can be gleaned from rug plots. The rug plot display maintains the temporal information of each auction (or event, in general) in order to compare global trends to more local ones. Concurrency can affect the shape of a price curve or the shape of the dynamics curve if we assume that different curve shapes capture different bidding patterns (e.g., gradual vs. bursty bidding or early vs. late bidding). Taking an exploratory approach, we examine the price and velocity curves in the rug plot for certain prominent types and look for patterns of temporal proximity between types of curves.

3.1.1 Curve Shapes

The empirical online auction literature describes several bidding patterns. One of these is late bidding ("sniping"), which results in very high bidding frequency at the end of an auction. We therefore expect auctions with sniping to have price-velocity curves that spike towards the auction end. Bidding frequency

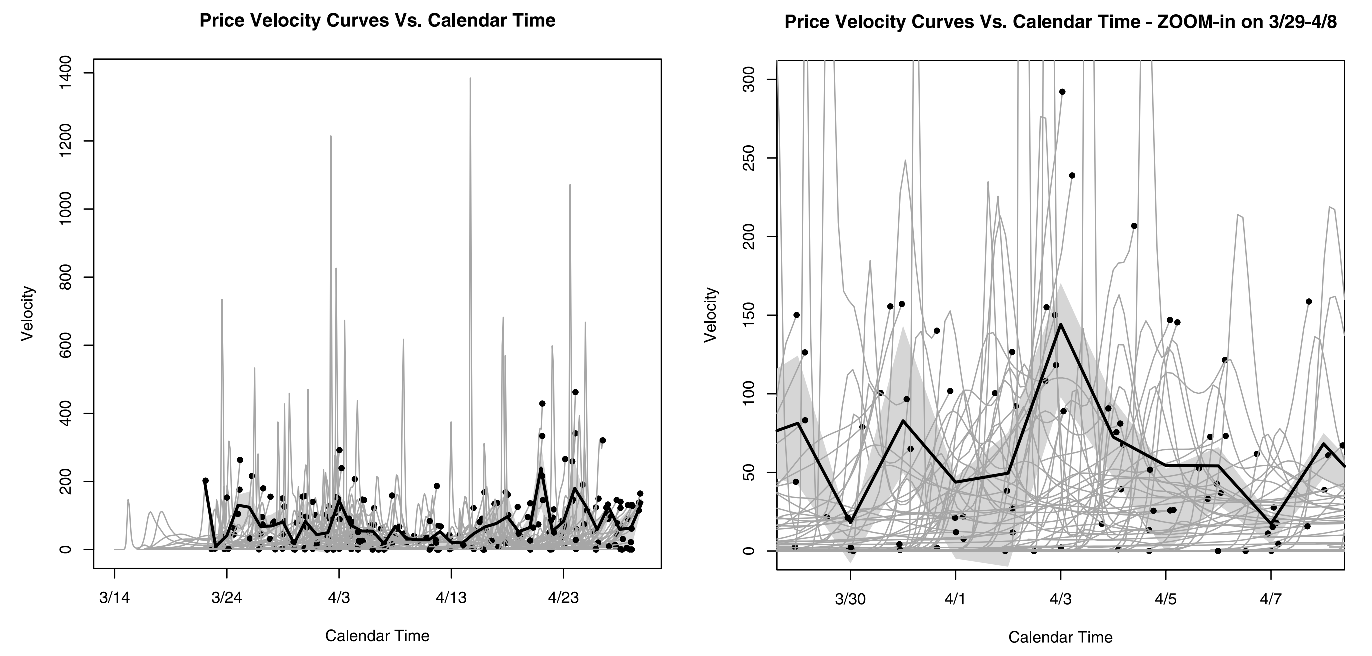

Figure 5. Price velocity over calendar time (left) and zoomed in to March 29–April 8 region (right). Thick curve and gray band are the daily median and IQR velocity.

should also manifest itself in the resulting curve: gradual bidding would lead to more gradual curves (and slow dynamics), whereas jump bidding would lead to bursty velocity curves. Relating different types of curves with bidding behavior can be useful in interpreting the rug plot results.

Examining the price curves in Figure 4 reveals three main shapes of price curves corresponding to the three shapes seen in Figure 3. The most popular (133 auctions) is a concave-up "j-shaped" price-curve which represents auctions with gradual price increases until mid-auction and a price jump towards the auction end, perhaps indicative of sniping. The corresponding velocity curve slowly increases and culminates in a peak. Further inspection reveals that most of the "j-shaped" auctions started at the lowest opening bid of $0.01.

A second type of price-curve that we identify is the straight line where the angle depends on the ratio of the opening and closing prices. The corresponding velocity curve is constant. These curves represent auctions with little to no bidding in the middle of the auction. Since all of our auctions transacted, a price-flat curve represents auctions with a high opening bid (set by the seller), and therefore, very little bidding. In our data, the 42 straight-line curve auctions vary widely in their opening and closing bids, but they all opened above $150.

The third typical shape in our data (61 auctions) is a stretched-out "S-shaped" curve which reflects auctions where the price increases slowly, then jumps up during mid-auction, and slowly increases to the close. The corresponding velocity curve spikes at the start of the price-jump, and then decreases when the price slows down again, resulting in a single "hump" shape. This can be indicative of jump bidding that occurs during mid-auction, where a single bidder raises the price drastically.

3.1.2 Temporal Groupings of Curve Types

When searching for auctions for a particular item on eBay, users are presented with information on all open auctions for that item as well as for auctions that closed in the last 15 days. A typical user would examine each of the open auctions (their current price, number of bids, and closing time) and perhaps some closed auctions to learn about closing prices. All these can influence how this bidder will place a bid. In addition, sellers can also use this information to schedule their auction, to set their opening bid, etc. We therefore expect to see temporal patterns in bidding (and posting) behavior in concurrent auctions. These patterns can manifest themselves as similarities in bidding behavior, for example, when bidders mimic a sniping behavior. Or they can show dissimilarities, if for instance, low-ending auctions lead new sellers to change strategy.

Our rug plots (Figures 4–5) show temporal effects such as clustering of shapes in certain periods or the lack of such clustering in others. One global observation is that the straight-line price curves appear to be scattered throughout the collection period. Many of these auctions have a high opening bid (sometimes higher than the interquartile range), suggesting an avenue of research to explain this phenomenon.

In general, during periods with many auctions it is useful to zoom in on the x-axis to better separate curves and make their shape more visible. For this reason we concentrate on three periods: the calendar start, the zoomed-in period of March 29–April 8 (more visible in the right panels of Figures 4–5), and the calendar end. During the beginning of the data period (before March 24), we see mostly "S-shaped" price curves, and a few straight-line price curves. The main distinction between the two appears to be the opening bid, with the "S-shaped" curves starting at lower prices. This is more visible in the velocity rug plot where the first group of auctions has a positive velocity "bump" at mid-auction (corresponding to the "S-shaped" price curves), followed by a period with nearly flat velocity curves. This raises questions about the bidding behavior during this period: What is going on during this flat price velocity period? Why was there no jump-bidding during this period?

The second period (March 29–April 8) contains many concurrent auctions (see the right panels of Figures 4–5). First, we see a global peak in velocity around April 3, as indicated by the median daily acceleration, indicating fast price increases during this period. Second, during this entire period we see two dominant types of auctions going on: a few "j-shaped" price curves (more easily identified in the zoomed-in velocity rug plot, as curves that end high), followed by a large group of "S-shaped" price curves (i.e., single-humped velocity curves). The price curves of both types appear to intersect mainly around $100–150 perhaps suggesting that bidders who see later closing auctions reaching a price in this range will tend to jump-bid to this amount in soon-to-close auctions, reasoning that the price will most likely reach (at least) this amount.

Finally, during the last period of the calendar, there appear to be more and more auctions with high opening bids. One possible explanation is that the Palm M515 has been auctioned long enough by this time that an opening bid of $50 or $100 seemed reasonable to bidders and sellers. The resulting price curves are mostly "j-shaped" curves, and many of them tend to close lower than the quartile daily price.

3.1.3 Other Groupings of Curves

The rug plot allows one to explore groupings not only of curve types but also of other relevant factors such as the number of bidders, the seller rating, and the auction duration. This can be done by color-coding each curve according to the level or category of the variable of interest. Such an extension allows investigating hypothesized concurrency effects such as: How does a high-rated seller's auction affect the bidding progress in a concurrent low-rated seller's auction? Or, do sellers choose the auction duration according to the durations of other open auctions for the similar item?

To show how this can reveal interesting effects, consider Figure 6, which displays the price curve rug plot from Figure 4 coded by auction duration. Besides the popularity of seven-day auctions (light gray), we see that some auction-durations seem to group temporally while others do not exhibit such a pattern. There are two periods when three-day (solid dark gray) auctions were very popular and a number of periods where five-day (dashed dark gray) auctions are prevalent. In contrast, 10-day auctions (black) are relatively rare and do not seem to group together.

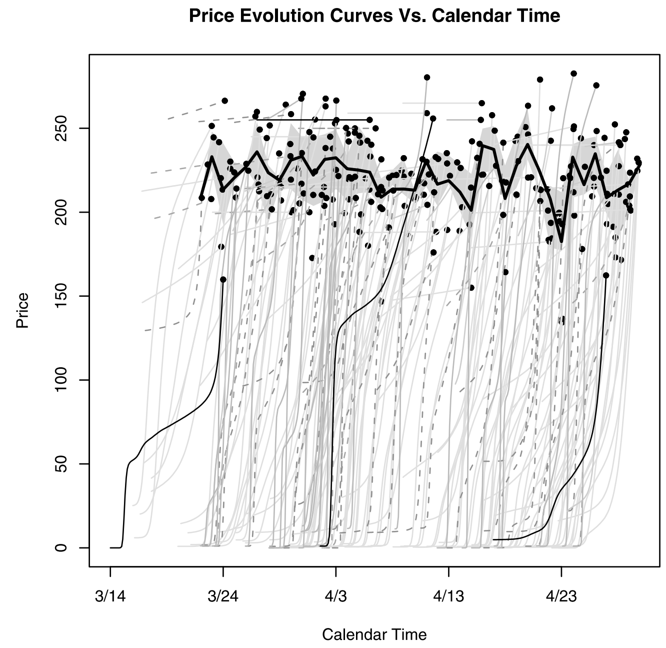

Figure 6. Price evolution by auction duration (light gray = 7-day, dashed dark gray = 5-day, solid dark gray = 3-day, black = 10-day.)

3.2 Time Grouped Box Plots

To study the relationship between closing prices of nearby auctions as well as trends over time, we grouped the auction closing price into categories based on the calendar date and ending time of the auction. Using multiscale boxplots for visualizing closing prices in a series of online auctions was proposed by Shmueli and Jank (2005). In particular, they proposed an interactive computation of boxplots according to temporal scales of interest ("STAT-zoom"). To incorporate the volume of auctions within each temporal window, the widths of the boxplots are proportional to the sample size, or alternatively the boxplots are coupled with histograms.

To study the relationship between selling prices of concurrent or nearly concurrent auctions we applied the same method. The levels of grouping we examined are weekly, daily, and half-daily. Figure 7 displays time-grouped boxplots. The width of the box is proportional to the square root of the number of auctions (McGill, Tukey, and Larsen 1978) to show the volume of auctions closing in a particular period. We see that the median of weekly closing prices is relatively constant over time whereas median prices are highly variable on the daily and half-daily scale. This is comparable to the stock market where prices are highly volatile within a day but less so when an entire week is considered. The within-day price fluctuations could be attributed to different time zones that bidders live in or to day/night differences in bidding activity. Although the width of the boxplot is visible in the weekly displays it is not easily differentiable at the daily and half-daily scales; therefore, we enhance the boxplots with histograms of auction closings. Figure 8 displays boxplots that are coupled histograms that contain the number of observations in each group. For the daily displays (left panel of Figure 8), the number of auctions closing on each day is between 2 and 10. In some cases where there are many auctions, the median closing price is small as would be expected by economic theory since as supply increases, price decreases. However, this is not

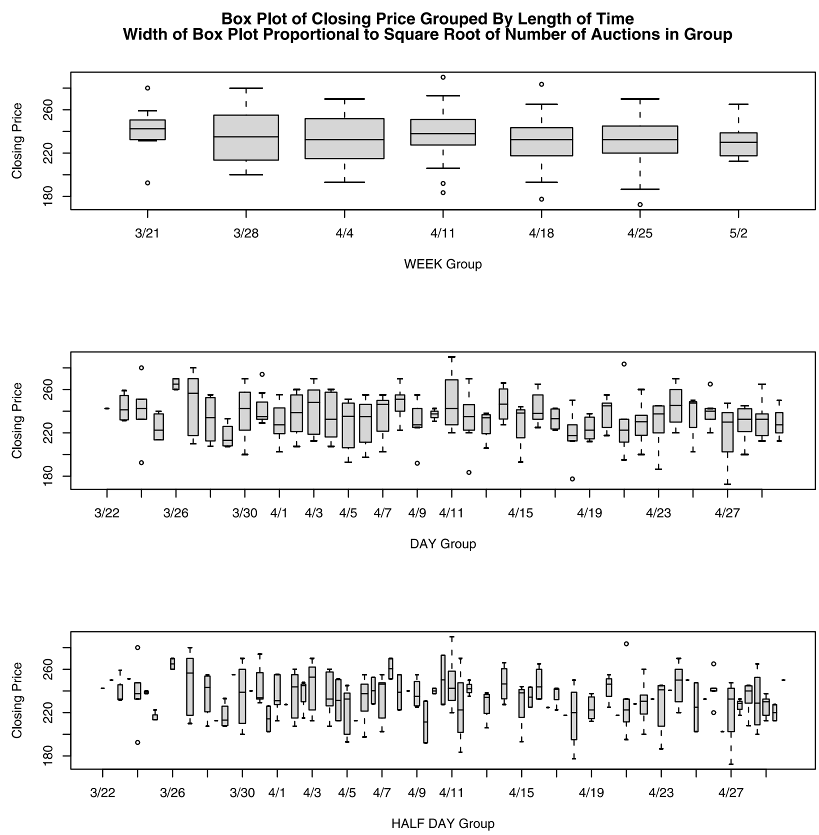

Figure 7. Grouped boxplots of closing price with width proportional to the square root of the number of auctions.

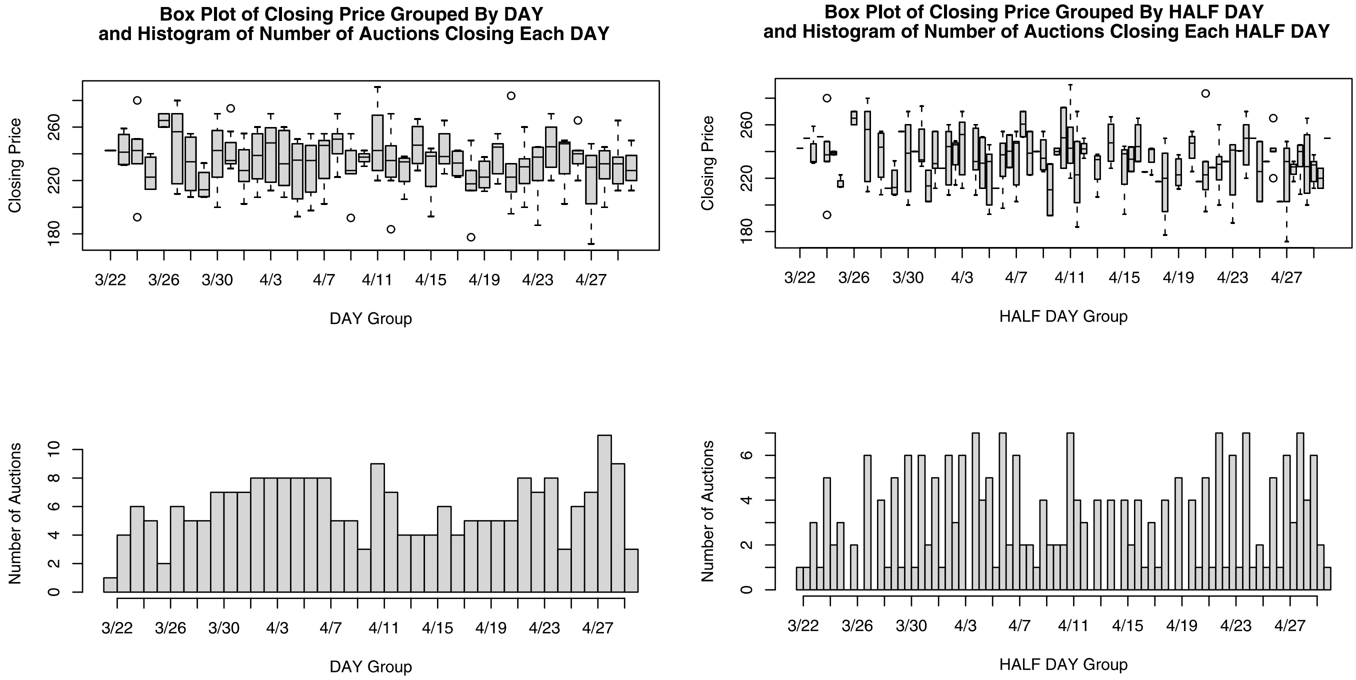

Figure 8. Closing price boxplots and auction frequency for day (left) and half day (right) groups.

always the case due to factors such as seller rating, shipping costs, experience of the bidders, and so on. that may also influence price. For instance, quite a few auctions end around April 3 and result in a high median price, much higher than the median price on the days before or after where the auction volume was comparable.

The histograms provided in the half-daily displays (right panel of Figure 8) show that most auctions close only during one half of the day, probably due to the fact that people sleep during the other half! However, there is also a significant number of auctions that close at night (e.g., around April 27) which could be attributed to time differences or simply to unfortunate and sloppy seller choices.

3.2.1 Comparing Medians of Adjacent Auctions

To better differentiate between adjacent medians, we use the test by Chambers, Cleveland, Kleiner, and Tukey (1983) that

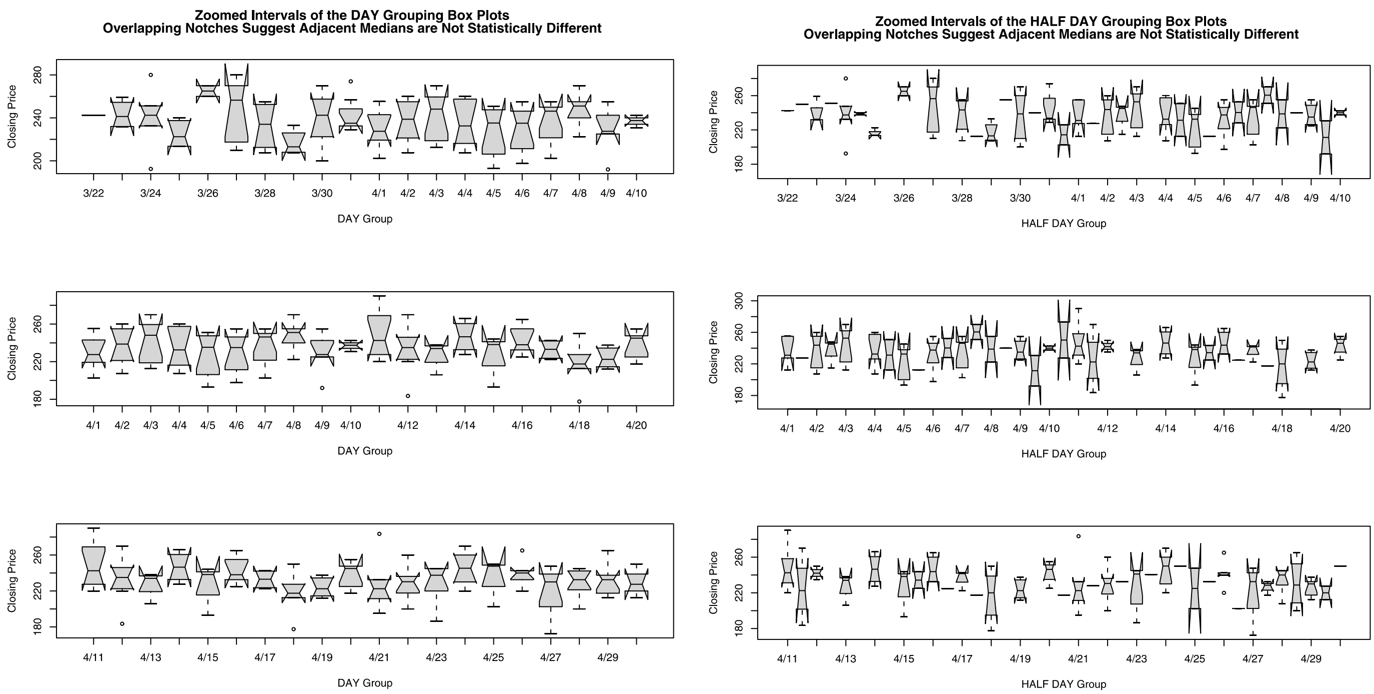

Figure 9. Closing price boxplots zoomed at different overlapping periods of calendar time for day groups (left) and half day groups (right).

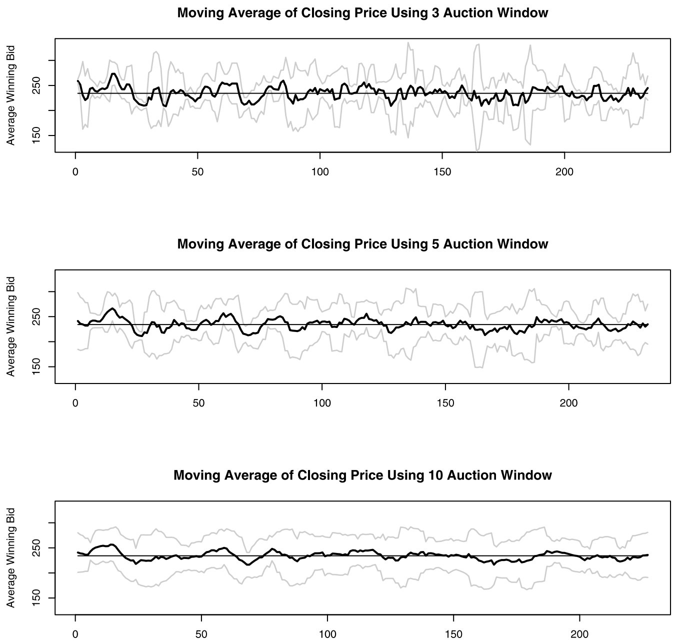

Figure 10. Closing price moving average based on different numbers of auctions.

compares medians of boxplots. The histogram function in R performs this test with the option "notch." The notches extend to $\pm 1.58\,\mathrm{IQR}/\sqrt{n}$ where IQR is the interquartile range. If the notches overlap, then the medians are not significantly different with 95% confidence. This test assumes asymptotic normality of the median and roughly equal sample sizes. Both assumptions seem reasonable with the auction data.

There are a total of 41 daily boxplots between March 22–April 30 in Figure 8. In order to de-clutter the graphical display, we partitioned the 41 plots into three groups (March 21–April 10, April 1–20, and April 11–30) and enhanced them with notches described above. The result can be seen in Figure 9. For instance, we can see that the notches for March 26 and 27 overlap; median prices for these two days are *not* significantly different. In contrast, the median prices of March 25 and 26 are significantly different. In general, while the median daily closing prices are variable, the differences are not statistically significant in most instances.

Figure 9 shows the notched boxplots for the half-daily displays. In this display, we find more statistical differences between adjacent medians. This suggests that the time of day can have an advantageous effect for the bidder or the seller. This finding also coincides with Snir (in press) who finds lower bids now (i.e., at night) if more auctions are expected in the future (i.e., during the day).

3.3 Moving Statistics Plots

Another way of visualizing time-trends in prices (or any other measure of interest) is by applying *moving statistics plots*, that is, via graphing the moving average or median. The particular statistic is computed over subsets of the data. We use the observation index for the purpose of grouping (the earliest closing auction is denoted "1," the next closing auction is denoted by "2," etc.). For example, using a moving window of 3 auctions for our 236 auctions in each window yields a series of 215 statistics.

Note that this method does not take into account the temporal distance between auctions. This means that moving statistics do not differentiate between auctions that are close in time and those that are far apart. One possibility is to use weights that reflect these temporal distances.

A moving statistic plot helps in understanding time-trends in the data, and unlike side-by-side boxplots, they present a smoother transition over time. Figure 10 shows a moving average plot for time-windows of 3, 5, and 10 auctions. The dark solid line shows the moving average, and the lighter lines correspond to upper and lower 95% confidence bands. The straight line through the center of the data shows the overall average closing price.

The increased variability in median price for the three- and five-auction time-windows compared to the relatively constant 10-day window suggests, again, that auctions are similar to the stock market: while there exists some variability in the short-run, prices tend to converge to the market value in the long-run. This also suggests that the smart (or lucky) bidder can take advantage of the variability to obtain a price that is lower than the market value.

3.4 Price Autocorrelation

We now examine the correlation between adjacent auctions. If the closing price in an auction affects the closing price in nearby auctions, we expect to find some degree of autocorrelation across the series of prices, especially if the groups are small enough. It is less likely to find autocorrelation in weekly prices because of the large price variability in auctions every week, and also possibly because the most popular auction duration is one week long. Narrowing the window to the day or half-day window may provide a better indication of autocorrelation. Ideally, we would like to look at hourly autocorrelation; unfortunately, this dataset is too sparse to do so. Because most auctions do not close in the middle of the night, it would be hard to find a dataset rich enough for this type of investigation. Therefore, a method for dealing with this type of missing values is necessary.

Even in the half-daily series we encounter missing data. Our solution is to impute missing values with the average of prices from one time period before and after the missing value. This "nearest neighbors" approach can be extended to multiple neighbors if it is reasonable to assume that the price is stable within a wider window of auctions. Another possibility is to impute the missing value using only information prior to the missing value, to reflect that fact that bidders do not have information about future prices.

Figure 11 shows price-autocorrelation based on the series of moving averages and moving medians for several window widths. It appears that there is no autocorrelation in the means or medians at any window. We suspect that the groups are too large and do not differentiate well between auctions that closed hours apart and those that closed minutes apart. Perhaps a richer

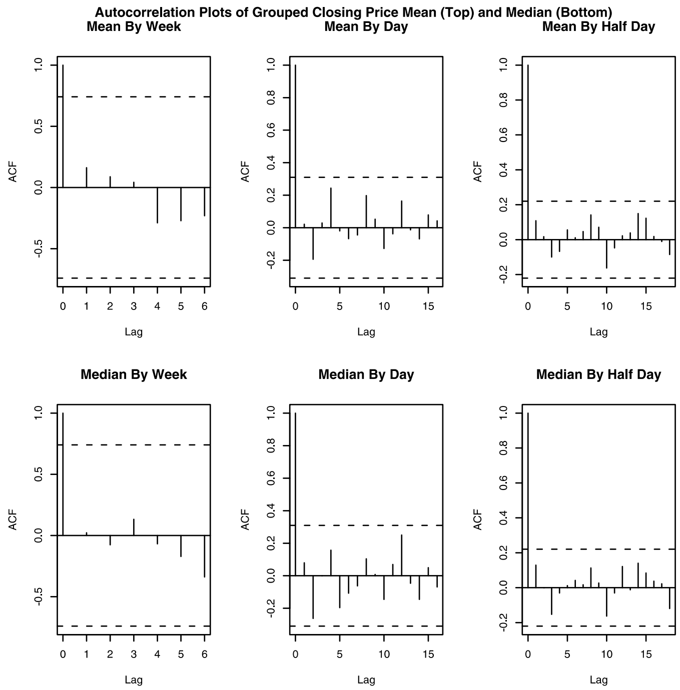

Figure 11. Autocorrelation for week (left), day (middle), and half day (right) groups for mean (top) and median (bottom).

dataset would reveal more, or a moving statistic that would account for the temporal distances between auctions. The general question is how to represent time-lags in unequally spaced observations (e.g., Jank and Shmueli 2006).

4. CONCLUSIONS

Concurrency is prevalent in the world of commerce. Whether it is stocks, sales promotions of competing products, or mortgage rates offered by lenders, action often results in reaction. Concurrency is particulary prevalent in the online environment since the openness of the Internet allows players to observe each other's moves in real time. In this article we consider concurrency of online auctions on eBay. We approach the study of concurrency from a visual point of view, and propose a series of visualizations that are suitable for the special structure of bid data. Considering the entire price evolution during an auction and its dynamics, we find several patterns and raise new research questions. We also find that online auctions tend to resemble the stock market in that most auctions tend to close around their market value, but still contain variability that allows bidders to purchase products significantly below market value. While some of this variability could be attributed to other factors such as seller rating and the winner's experience, it creates opportunities for an eBay bidder. We also find a difference in median closing prices at different times of the day. A bidder, for instance, may be able to obtain a lower value by taking part in auctions that close later in the evening when fewer people are taking part in an auction. These observations are all based on the exploratory analysis performed here and can serve as a basis for further studies of concurrency in online auctions.

[Received October 2005. Revised April 2006.]

REFERENCES

Allen, M. T., and Swisher, J. (2000), "An Analysis of the Price Formation Process at a HUD Auction," *Journal of Real Estate Research*, 20, 279–298.

Anwar, A., McMillan, R., and Zheng, M. (2006), "Bidding Behavior in Competing Auctions: Evidence from eBay," *European Economic Review*, 50, 307–322.

Aris, A., Shneiderman, B., Plaisant, C., Shmueli, G., and Jank, W. (2005), "Representing Unevenly-Spaced Time Series Data for Visualization and Interactive Exploration," in *Proceedings of the International Conference on Human-Computer Interaction (INTERACT 2005)*, LNCS 3585, pp. 835–846.

Bajari, P., and Hortacsu, A. (2003), "Winner's Curse, Reserve Price and Endogenous Entry: Empirical Insights from eBay," *RAND Journal of Economics*, 34, 329–355.

Chambers, J. M., Cleveland, W. S., Kleiner, B., and Tukey, P. A. (1983), *Graphical Methods for Data Analysis*, Belmont, CA: Wadsworth International Group.

Deltas, G. (1999), "Auction Size and Price Dynamics in Sequential Auctions," unpublished working paper, University of Illinois.

Etzion, H., Pinker, E., and Seidermann, A. (2004), "Analyzing the Simultaneous Use of Auctions and Posted Prices for On-line Selling," unpublished working paper No. CIS 03-01, Simon Business School.

Gallien, J. (2002), "Dynamic Mechanism Design for Online Commerce," unpublished working paper, MIT Sloan School of Management.

Ghani, R., and Simmons, H. (2004), "Predicting the End-price of Online Auctions," in *Proceedings of the International Workshop on Data Mining and Adaptive Modelling Methods for Economics and Management*.

Guerre, E., Perrigne, I., and Vuong, Q. (2000), "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68, 525–574.

Jank, W., and Shmueli, G. (2005), "Profiling Price Dynamics in Online Auctions Using Curve Clustering," Robert H. Smith School Research Paper No. RHS-06-004, University of Maryland. Available online at *http://ssrn.com/abstract=902893*.

——— (2006) "Modeling Concurrencty of Events in Online Auctions via Spatio-Temporal Semiparametric Models," unpublished working paper, Smith School of Business, University of Maryland.

Jank, W., Shmueli, G., Plaisant, C., and Shneiderman, B. (2006), "Visualizing Functional Data with an Application to eBays Online Auctions," in *Handbook on Computational Statistics on Data Visualization*, eds. Chen, Haerdle, and Unwin, Springer Verlag, Heidelberg.

Kauffman, R. J., and Wood, C. A. (2005), "The Effects of Shilling on Final Bid Prices in Online Auctions," *Electronic Commerce Research and Applications*, 4, 18–31.

Lucking-Reiley, D., Bryan, D., Prasad, N., and Reeves, D. (2005), "Pennies from eBay: the Determinants of Price in Online Auctions," unpublished working paper, University of Arizona.

McGill, R., Tukey, J. W., and Larsen, W. A. (1978), "Variations of Box Plots," *The American Statistician*, 38, 12–16.

Ramsey, J. O., and Silverman, B. W. (2005), *Functional Data Analysis*, New York; Springer-Verlag.

Roth, A. E., and Ockenfels, A. (2002), "Last-Minute Bidding and the Rules for Ending Second-price Auctions: Evidence from Ebay and Amazon Auctions on the Internet," *The American Economic Review*, 92, 1093–1103.

Shmueli, G., and Jank W., (2005), "Visualizing Online Auctions," *Journal of Computational and Graphical Statistics*, 14, 299–319.

Shmueli, G., Russo, R. P., and Jank, W. (2004), "Modelling Bid Arrivals in Online Auctions," unpublished working paper, Robert H. Smith School, Research Paper No. RHS-06-001, University of Maryland. Available online at *http://ssrn.com/abstract=902868*.

Snir, E. M., (in press), "Online Auction Enabling the Secondary Computer Market," *Information Technology and Management*.

Vakrat, Y., and Seidmann, A. (1999), "Can Online Auctions Beat Online Catalogs?," in *Proceedings of the Twentieth International Conference on Information Systems (ICIS)*, Charlotte, NC.

Zeithammer, R. (2005), "Forward-Looking Bidding in Online Auctions," *Journal of Marketing Research*, August.