# The Sequential Auction Problem on eBay: An Empirical Analysis and a Solution*

Adam I. Juda
Division of Engineering and Applied Sciences
Harvard University,
Harvard Business School
ajuda@hbs.edu

David C. Parkes
Division of Engineering and Applied Sciences
Harvard University
parkes@eecs.harvard.edu

## ABSTRACT

Bidders on eBay have no dominant bidding strategy when faced with multiple auctions each offering an item of interest. As seen through an analysis of 1,956 auctions on eBay for a Dell E193FP LCD monitor, some bidders win auctions at prices higher than those of other available auctions, while others never win an auction despite placing bids in losing efforts that are greater than the closing prices of other available auctions. These misqueues in strategic behavior hamper the efficiency of the system, and in so doing limit the revenue potential for sellers. This paper proposes a novel options-based extension to eBay's proxy-bidding system that resolves this strategic issue for buyers in commoditized markets. An empirical analysis of eBay provides a basis for computer simulations that investigate the market effects of the options-based scheme, and demonstrates that the options-based scheme provides greater efficiency than eBay, while also increasing seller revenue.

## Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

## General Terms

Algorithms, Design, Economics

## Keywords

Online Auctions, Options, Proxy Bidding, Sequential Auction Problem, eBay

## 1. INTRODUCTION

Electronic markets represent an application of information systems that has generated significant new trading opportunities while allowing for the dynamic pricing of goods. In addition to marketplaces such as eBay, electronic marketplaces are increasingly used for business-to-consumer auctions (e.g. to sell surplus inventory [19]).

Many authors have written about a future in which commerce is mediated by online, automated trading agents [10, 25, 1]. There is still little evidence of automated trading in e-markets, though. We believe that one leading place of resistance is in the lack of provably optimal bidding strategies for any but the simplest of market designs. Without this, we do not expect individual consumers, or firms, to be confident in placing their business in the "hands" of an automated agent.

One of the most common examples today of an electronic marketplace is eBay, where the gross merchandise volume (i.e., the sum of all successfully closed listings) during 2005 was $44B. Among items listed on eBay, many are essentially identical. This is especially true in the Consumer Electronics category [9], which accounted for roughly $3.5B of eBay's gross merchandise volume in 2005. This presence of essentially identical items can expose bidders, and sellers, to risks because of the sequential auction problem.

For example, Alice may want an LCD monitor, and could potentially bid in either a 1 o'clock or 3 o'clock eBay auction. While Alice would prefer to participate in whichever auction will have the lower winning price, she cannot determine beforehand which auction that may be, and could end up winning the "wrong" auction. This is a problem of *multiple copies*.

Another problem bidders may face is the *exposure problem*. As investigated by Bykowsky et al. [6], exposure problems exist when buyers desire a bundle of goods but may only participate in single-item auctions.[1] For example, if Alice values a video game console by itself for $200, a video game by itself for $30, and both a console and game for $250, Alice must determine how much of the $20 of synergy value she might include in her bid for the console alone. Both problems arise in eBay as a result of sequential auctions of single items coupled with patient bidders with substitutes or complementary valuations.

Why might the sequential auction problem be bad? Complex games may lead to bidders employing costly strategies and making mistakes. Potential bidders who do not wish to bear such costs may choose not to participate in the

*A preliminary version of this work appeared in the AMEC workshop in 2004.

[1]The exposure problem has been primarily investigated by Bykowsky et al. in the context of *simultaneous* single-item auctions. The problem is also a familiar one of online decision making.


*EC'06*, June 11–15, 2006, Ann Arbor, Michigan, USA.

market, inhibiting seller revenue opportunities. Additionally, among those bidders who do choose to participate, the mistakes made may lead to inefficient allocations, further limiting revenue opportunities.

We are interested in creating modifications to eBay-style markets that simplify the bidder problem, leading to simple equilibrium strategies, and preferably better efficiency and revenue properties.

## 1.1 Options + Proxies: A Proposed Solution

Retail stores have developed policies to assist their customers in addressing sequential purchasing problems. Return policies alleviate the exposure problem by allowing customers to return goods at the purchase price. Price matching alleviates the multiple copies problem by allowing buyers to receive from sellers after purchase the difference between the price paid for a good and a lower price found elsewhere for the same good [7, 15, 18]. Furthermore, price matching can reduce the impact of exactly when a seller brings an item to market, as the price will in part be set by others selling the same item. These two retail policies provide the basis for the scheme proposed in this paper.[2]

We extend the proxy bidding technology currently employed by eBay. Our *"super"-proxy* extension will take advantage of a new, *real options-based*, market infrastructure that enables simple, yet optimal, bidding strategies. The extensions are computationally simple, handle temporal issues, and retain seller autonomy in deciding when to enter the market and conduct individual auctions.

A seller sells an *option* for a good, which will ultimately lead to either a sale of the good or the return of the option. Buyers interact through a proxy agent, defining a value on all possible bundles of goods in which they have interest together with the latest time period in which they are willing to wait to receive the good(s). The proxy agents use this information to determine how much to bid for options, and follow a dominant bidding strategy across all relevant auctions. A proxy agent exercises options held when the buyer's patience has expired, choosing options that maximize a buyer's payoff given the reported valuation. All other options are returned to the market and not exercised. The options-based protocol makes truthful and immediate revelation to a proxy a dominant strategy for buyers, whatever the future auction dynamics.

We conduct an empirical analysis of eBay, collecting data on over four months of bids for Dell LCD screens (model E193FP) starting in the Summer of 2005. LCD screens are a high-ticket item, for which we demonstrate evidence of the sequential bidding problem. We first infer a conservative model for the arrival time, departure time and value of bidders on eBay for LCD screens during this period. This model is used to simulate the performance of the options-based infrastructure, in order to make direct comparisons to the actual performance of eBay in this market.

We also extend the work of Haile and Tamer [11] to estimate an upper bound on the distribution of value of eBay bidders, taking into account the sequential auction problem when making the adjustments. Using this estimate, one can approximate how much greater a bidder's true value is from the maximum bid they were observed to have placed on eBay. Based on this approximation, revenue generated in a simulation of the options-based scheme exceeds revenue on eBay for the comparable population and sequence of auctions by 14.8%, while the options-based scheme demonstrates itself as being 7.5% more efficient.

[2]Prior work has shown price matching as a potential mechanism for colluding firms to set monopoly prices. However, in our context, auction prices will be matched, which are not explicitly set by sellers but rather by buyers' bids.

## 1.2 Related Work

A number of authors [27, 13, 28, 29] have analyzed the multiple copies problem, often times in the context of categorizing or modeling sniping behavior for reasons other than those first brought forward by Ockenfels and Roth [20]. These papers perform equilibrium analysis in simpler settings, assuming bidders can participate in at most two auctions. Peters & Severinov [21] extend these models to allow buyers to consider an arbitrary number of auctions, and characterize a perfect Bayesian equilibrium. However, their model does not allow auctions to close at distinct times and does not consider the arrival and departure of bidders.

Previous work have developed a data-driven approach toward developing a taxonomy of strategies employed by bidders in practice when facing multi-unit auctions, but have not considered the sequential bidding problem [26, 2]. Previous work has also sought to provide agents with smarter bidding strategies [4, 3, 5, 1]. Unfortunately, it seems hard to design artificial agents with equilibrium bidding strategies, even for a simple simultaneous ascending price auction.

Iwasaki et al. [14] have considered the role of options in the context of a single, monolithic, auction design to help bidders with marginal-increasing values avoid exposure in a multi-unit, homogeneous item auction problem. In other contexts, options have been discussed for selling coal mine leases [23], or as *leveled commitment contracts* for use in a decentralized market place [24]. Most similar to our work, Gopal et al. [9] use options for reducing the risks of buyers and sellers in the sequential auction problem. However, their work uses costly options and does not remove the sequential bidding problem completely.

Work on *online* mechanisms and *online* auctions [17, 12, 22] considers agents that can dynamically arrive and depart across time. We leverage a recent price-based characterization by Hajiaghayi et al. [12] to provide a dominant strategy equilibrium for buyers within our options-based protocol. The special case for single-unit buyers is equivalent to the protocol of Hajiaghayi et al., albeit with an options-based interpretation.

Jiang and Leyton-Brown [16] use machine learning techniques for bid identification in online auctions.

# 2. EBAY AND THE DELL E193FP

The most common type of auction held on eBay is a single-item proxy auction. Auctions open at a given time and remain open for a set period of time (usually one week). Bidders bid for the item by giving a proxy a value ceiling. The proxy will bid on behalf of the bidder only as much as is necessary to maintain a winning position in the auction, up to the ceiling received from the bidder. Bidders may communicate with the proxy multiple times before an auction closes. In the event that a bidder's proxy has been outbid, a bidder may give the proxy a higher ceiling to use in the auction. eBay's proxy auction implements an incremental version of a Vickrey auction, with the item sold to the highest bidder for the second-highest bid plus a small increment.

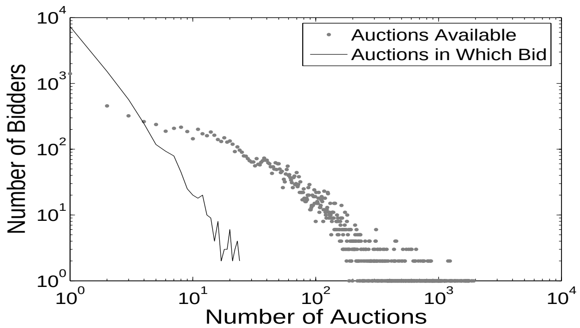


**Figure 1: Histogram of number of LCD auctions available to each bidder and number of LCD auctions in which a bidder participates.**

The market analyzed in this paper is that of a specific model of an LCD monitor, a 19" Dell LCD model E193FP. This market was selected for a variety of reasons including:

- The mean price of the monitor was \$240 (with standard deviation \$32), so we believe it reasonable to assume that bidders on the whole are only interested in acquiring one copy of the item on eBay.[3]
- The volume transacted is fairly high, at approximately 500 units sold per month.
- The item is not usually bundled with other items.
- The item is typically sold "as new," and so suitable for the price-matching of the options-based scheme.

Raw auction information was acquired via a PERL script. The script accesses the eBay search engine,[4] and returns all auctions containing the terms 'Dell' and 'LCD' that have closed within the past month.[5] Data was stored in a text file for post-processing. To isolate the auctions in the domain of interest, queries were made against the titles of eBay auctions that closed between 27 May, 2005 through 1 October, 2005.[6]

Figure 1 provides a general sense of how many LCD auctions occur while a bidder is interested in pursuing a monitor.[7] 8,746 bidders (86%) had more than one auction available between when they first placed a bid on eBay and the latest closing time of an auction in which they bid (with an average of 78 auctions available). Figure 1 also illustrates the number of auctions in which each bidder participates. Only 32.3% of bidders who had more than one auction available are observed to bid in more than one auction (bidding in 3.6 auctions on average). A simple regression analysis shows that bidders tend to submit maximal bids to an auction that are \$1.22 higher after spending twice as much time in the system, as well as bids that are \$0.27 higher in each subsequent auction.

Among the 508 bidders that won exactly one monitor and participated in multiple auctions, 201 (40%) paid more than \$10 more than the closing price of another auction *in which they bid*, paying on average \$35 more (standard deviation \$21) than the closing price of the cheapest auction in which they bid but did not win. Furthermore, among the 2,216 bidders that never won an item despite participating in multiple auctions, 421 (19%) placed a losing bid in one auction that was more than \$10 higher than the closing price of another auction *in which they bid*, submitting a losing bid on average \$34 more (standard deviation \$23) than the closing price of the cheapest auction in which they bid but did not win. Although these measures do not say a bidder that lost could have definitively won (because we only consider the final winning price and not the bid of the winner to her proxy), or a bidder that won could have secured a better price, this is at least indicative of some bidder mistakes.

# 3. MODELING THE SEQUENTIAL AUCTION PROBLEM

While the eBay analysis was for simple bidders who desire only a single item, let us now consider a more general scenario where people may desire multiple goods of different types, possessing general valuations over those goods.

Consider a world with buyers (sometimes called bidders) $B$ and $K$ different types of goods $G_1 ... G_K$. Let $T = \{0, 1, ...\}$ denote time periods. Let $L$ denote a bundle of goods, represented as a vector of size $K$, where $L_k \in \{0, 1\}$ denotes the quantity of good type $G_k$ in the bundle.[8] The type of a buyer $i \in B$ is $(a_i, d_i, v_i)$, with arrival time $a_i \in T$, departure time $d_i \in T$, and *private* valuation $v_i(L) \geq 0$ for each bundle of goods $L$ received between $a_i$ and $d_i$, and zero value otherwise. The arrival time models the period in which a buyer first realizes her demand and enters the market, while the departure time models the period in which a buyer loses interest in acquiring the good(s). In settings with general valuations, we need an additional assumption: an upper bound on the difference between a buyer's arrival and departure, denoted $\Delta_{\text{Max}}$. Buyers have quasi-linear utilities, so that the utility of buyer $i$ receiving bundle $L$ and paying $p$, in some period no later than $d_i$, is $u_i(L, p) = v_i(L) - p$. Each seller $j \in S$ brings a single item $k_j$ to the market, has no intrinsic value and wants to maximize revenue. Seller $j$ has an arrival time, $a_j$, which models the period in which she is first interested in listing the item, while the departure time, $d_j$, models the latest period in which she is willing to consider having an auction for the item close. A seller will receive payment by the end of the reported departure of the winning buyer.

[3] For reference, Dell's October 2005 mail order catalogue quotes the price of the monitor as being \$379 without a desktop purchase, and \$240 as part of a desktop purchase upgrade.

[4] http://search.ebay.com

[5] The search is not case-sensitive.

[6] Specifically, the query found all auctions where the title contained all of the following strings: 'Dell,' 'LCD' and 'E193FP,' while excluding all auctions that contained any of the following strings: 'Dimension,' 'GHZ,' 'desktop,' 'p4' and 'GB.' The exclusion terms were incorporated so that the only auctions analyzed would be those selling exclusively the LCD of interest. For example, the few bundled auctions selling both a Dell Dimension desktop and the E193FP LCD are excluded.

[7] As a reference, most auctions close on eBay between noon and midnight EDT, with almost two auctions for the Dell LCD monitor closing each hour on average during peak time periods. Bidders have an average observed patience of 3.9 days (with a standard deviation of 11.4 days).

[8] We extend notation whereby a single item $k$ of type $G_k$ refers to a vector $L : L_k = 1$.

We say an individual auction in a sequence is *locally strategyproof* (LSP) if truthful bidding is a dominant strategy for a buyer that can only bid in that auction. Consider the following example to see that LSP is insufficient for the existence of a dominant bidding strategy for buyers facing a sequence of auctions.

EXAMPLE 1. *Alice values one ton of Sand with one ton of Stone at* $2,000. *Bob holds a Vickrey auction for one ton of Sand on Monday and a Vickrey auction for one ton of Stone on Tuesday. Alice has no dominant bidding strategy because she needs to know the price for Stone on Tuesday to know her maximum willingness to pay for Sand on Monday.*

DEFINITION 1. The sequential auction problem. *Given a sequence of auctions, despite each auction being locally strategyproof, a bidder has no dominant bidding strategy.*

Consider a sequence of auctions. Generally, auctions selling the same item will be *uncertainly-ordered*, because a buyer will not know the ordering of closing prices among the auctions. Define the *interesting bundles* for a buyer as all bundles that could maximize the buyer's profit for some combination of auctions and bids of other buyers.[9] Within the interesting bundles, say that an item has *uncertain marginal value* if the marginal value of an item depends on the other goods held by the buyer.[10] Say that an item is *over-supplied* if there is more than one auction offering an item of that type. Say two bundles are *substitutes* if one of those bundles has the same value as the union of both bundles.[11]

PROPOSITION 1. *Given* locally strategyproof *single-item auctions, the sequential auction problem exists for a bidder if and only if either of the following two conditions is true: (1) within the set of interesting bundles (a) there are two bundles that are substitutes, (b) there is an item with uncertain marginal value, or (c) there is an item that is over-supplied; (2) a bidder faces competitors' bids that are conditioned on the bidder's past bids.*

PROOF. (Sketch.)($\Leftarrow$) A bidder does not have a dominant strategy when (a) she does not know which bundle among substitutes to pursue, (b) she faces the exposure problem, or (c) she faces the multiple copies problem. Additionally, a bidder does not have a dominant strategy when she does not how to optimally influence the bids of competitors.($\Rightarrow$) By contradiction. A bidder has a dominant strategy to bid its constant marginal value for a given item in each auction available when conditions (1) and (2) are both false. □

For example, the following buyers all face the sequential auction problem as a result of condition (a), (b) and (c) respectively: a buyer who values one ton of Sand for $1,000, or one ton of Stone for $2,000, but not both Sand and Stone; a buyer who values one ton of Sand for $1,000, one ton of Stone for $300, and one ton of Sand and one ton of Stone for $1,500, and can participate in an auction for Sand before an auction for Stone; a buyer who values one ton of Sand for $1,000 and can participate in many auctions selling Sand.

[9]Assume that the empty set is an interesting bundle.

[10]Formally, an item $k$ has *uncertain marginal value* if $|\{m : m = v_i(Q) - v_i(Q-k), \forall Q \subseteq L \in \text{InterestingBundle}, Q \supseteq k\}| > 1$.

[11]Formally, two bundles $A$ and $B$ are *substitutes* if $v_i(A \cup B) = \max(v_i(A), v_i(B))$, where $A \cup B = L$ where $L_k = \max(A_k, B_k)$.

# 4. "SUPER" PROXIES AND OPTIONS

The novel solution proposed in this work to resolve the sequential auction problem consists of two primary components: richer proxy agents, and options with price matching.

In finance, a real option is a right to acquire a real good at a certain price, called the *exercise price*. For instance, Alice may obtain from Bob the right to buy Sand from him at an exercise price of $1,000. An option provides the *right* to purchase a good at an exercise price but not the obligation. This flexibility allows buyers to put together a collection of options on goods and then decide which to exercise.

Options are typically sold at a price called the *option price*. However, options obtained at a non-zero option price cannot generally support a simple, dominant bidding strategy, as a buyer must compute the expected value of an option to justify the cost [8]. This computation requires a model of the future, which in our setting requires a model of the bidding strategies and the values of other bidders. This is the very kind of game-theoretic reasoning that we want to avoid.

Instead, we consider *costless options* with an option price of zero. This will require some care as buyers are weakly better off with a costless option than without one, whatever its exercise price. However, multiple bidders pursuing options with no intention of exercising them would cause the efficiency of an auction for options to unravel. This is the role of the **mandatory** *proxy agents*, which intermediate between buyers and the market. A proxy agent forces a link between the valuation function used to acquire options and the valuation used to exercise options. If a buyer tells her proxy an inflated value for an item, she runs the risk of having the proxy exercise options at a price greater than her value.

## 4.1 Buyer Proxies

### 4.1.1 Acquiring Options

After her arrival, a buyer submits her valuation $\hat{v}_i$ (perhaps untruthfully) to her proxy in some period $\hat{a}_i \geq a_i$, along with a claim about her departure time $\hat{d}_i \geq \hat{a}_i$. All transactions are intermediated via proxy agents. Each auction is modified to sell an option on that good to the highest bidding proxy, with an *initial* exercise price set to the second-highest bid received.[12]

When an option in which a buyer is interested becomes available for the first time, the proxy determines its bid by computing the buyer's *maximum marginal value* for the item, and then submits a bid in this amount. A proxy does not bid for an item when it already holds an option. The bid price is:

$$bid_i^t(k) = \max_L [\hat{v}_i(L+k) - \hat{v}_i(L)] \quad (1)$$

By having a proxy compute a buyer's maximum marginal value for an item and then bidding only that amount, a buyer's proxy will win any auction that *could possibly* be of benefit to the buyer and only lose those auctions that *could never* be of value to the buyer.

[12]The system can set a reserve price for each good, provided that the reserve is universal for all auctions selling the same item. Without a universal reserve price, price matching is not possible because of the additional restrictions on prices that individual sellers will accept.

| Buyer | Type | Monday | Tuesday |
|---|---|---|---|
| Molly | (Mon, Tues, \$8) | $6_{\text{Nancy}}$ | $6_{\text{Nancy}} \rightarrow 4_{\text{Polly}}$ |
| Nancy | (Mon, Tues, \$6) | - | $4_{\text{Polly}}$ |
| Polly | (Mon, Tues, \$4) | - | - |

**Table 1: Three-buyer example with each wanting a single item and one auction occurring on Monday and Tuesday. "$X_Y$" implies an option with exercise price $X$ and bookkeeping that a proxy has prevented $Y$ from currently possessing an option. "→" is the updating of exercise price and bookkeeping.**

When a proxy wins an auction for an option, the proxy will store in its local memory the identity (which may be a pseudonym) of the proxy not holding an option because of the proxy's win (i.e., the proxy that it 'bumped' from winning, if any). This information will be used for price matching.

#### 4.1.2 Pricing Options

Sellers agree by joining the market to allow the proxy representing a buyer to adjust the exercise price of an option that it holds downwards if the proxy discovers that it could have achieved a better price by waiting to bid in a later auction for an option on the same good. To assist in the implementation of the price matching scheme each proxy tracks future auctions for an option that it has already won and will determine who would be bidding in that auction had the proxy delayed its entry into the market until this later auction. The proxy will request price matching from the seller that granted it an option if the proxy discovers that it could have secured a lower price by waiting. To reiterate, the proxy does *not* acquire more than one option for any good. Rather, it reduces the exercise price on its already issued option if a better deal is found.

The proxy is able to discover these deals by asking each future auction to report the identities of the bidders in that auction together with their bids. This needs to be enforced by eBay, as the central authority. The highest bidder in this later auction, across those whose identity is **not** stored in the proxy's memory for the given item, is exactly the bidder against whom the proxy would be competing had it delayed its entry until this auction. If this high bid is lower than the current option price held, the proxy "price matches" down to this high bid price.

After price matching, one of two adjustments will be made by the proxy for bookkeeping purposes. If the winner of the auction is the bidder whose identity has been in the proxy's local memory, the proxy will replace that local information with the identity of the bidder whose bid it just price matched, as that is now the bidder the proxy has prevented from obtaining an option. If the auction winner's identity is not stored in the proxy's local memory the memory may be cleared. In this case, the proxy will simply price match against the bids of future auction winners on this item until the proxy departs.

EXAMPLE 2 (TABLE 1). *Molly's proxy wins the Monday auction, submitting a bid of \$8 and receiving an option for \$6. Molly's proxy adds Nancy to its local memory as Nancy's proxy would have won had Molly's proxy not bid. On Tuesday, only Nancy's and Polly's proxy bid (as Molly's proxy holds an option), with Nancy's proxy winning an option for \$4 and noting that it bumped Polly's proxy. At this time, Molly's proxy will price match its option down to \$4 and replace Nancy with Polly in its local memory as per the price match algorithm, as Polly would be holding an option had Molly never bid.*

| Buyer | Type | Monday | Tuesday |
|---|---|---|---|
| Truth: | | | |
| Molly | (Mon, Mon, \$8) | $6_{\text{Nancy}}$ | - |
| Nancy | (Mon, Tues, \$6) | - | $4_{\text{Polly}}$ |
| Polly | (Mon, Tues, \$4) | - | - |
| Misreport: | | | |
| Molly | (Mon, Mon, \$8) | - | - |
| Nancy | (Mon, Tues, $\widehat{\$10}$) | $8_{\text{Molly}}$ | $8_{\text{Molly}} \rightarrow 4_{\phi}$ |
| Polly | (Mon, Tues, \$4) | - | $0_{\phi}$ |
| Misreport & match low: | | | |
| Molly | (Mon, Mon, \$8) | - | - |
| Nancy | (Mon, Tues, $\widehat{\$10}$) | 8 | $8 \rightarrow 0$ |
| Polly | (Mon, Tues, \$4) | - | 0 |

**Table 2: Examples demonstrating why bookkeeping will lead to a truthful system whereas simply matching to the lowest winning price will not.**

#### 4.1.3 Exercising Options

At the reported departure time the proxy chooses which options to exercise. Therefore, a seller of an option must wait until period $\hat{d}_w$ for the option to be exercised and receive payment, where $w$ was the winner of the option.[13] For bidder $i$, in period $\hat{d}_i$, the proxy chooses the option(s) that maximize the (reported) utility of the buyer:

$$\theta_t^* = \underset{\theta \subseteq \Theta}{\operatorname{argmax}} (\hat{v}_i(\gamma(\theta)) - \pi(\theta)) \tag{2}$$

where $\Theta$ is the set of all options held, $\gamma(\theta)$ are the goods corresponding to a set of options, and $\pi(\theta)$ is the sum of exercise prices for a set of options. All other options are returned.[14] No options are exercised when no combination of options have positive utility.

#### 4.1.4 Why bookkeep and not match winning price?

One may believe that an alternative method for implementing a price matching scheme could be to simply have proxies match the lowest winning price they observe after winning an option. However, as demonstrated in Table 2, such a simple price matching scheme will not lead to a truthful system.

The first scenario in Table 2 demonstrates the outcome if all agents were to truthfully report their types. Molly

[13] While this appears restrictive on the seller, we believe it not significantly different than what sellers on eBay currently endure in practice. An auction on eBay closes at a specific time, but a seller must wait until a buyer relinquishes payment before being able to realize the revenue, an amount of time that could easily be days (if payment is via a money order sent through courier) to much longer (if a buyer is slow but not overtly delinquent in remitting her payment).

[14] Presumably, an option returned will result in the seller holding a new auction for an option on the item it still possesses. However, the system will not allow a seller to re-auction an option until $\Delta_{\text{Max}}$ after the option *had first been issued* in order to maintain a truthful mechanism.

would win the Monday auction and receive an option with an exercise price of \$6 (subsequently exercising that option at the end of Monday), and Nancy would win the Tuesday auction and receive an option with an exercise price of \$4 (subsequently exercising that option at the end of Tuesday).

The second scenario in Table 2 demonstrates the outcome if Nancy were to misreport her value for the good by reporting an inflated value of \$10, *using the proposed bookkeeping method*. Nancy would win the Monday auction and receive an option with an exercise price of \$8. On Tuesday, Polly would win the auction and receive an option with an exercise price of \$0. Nancy's proxy would observe that the highest bid submitted on Tuesday among those proxies not stored in local memory is Polly's bid of \$4, and so Nancy's proxy would price match the exercise price of its option down to \$4. Note that the exercise price Nancy's proxy has obtained at the end of Tuesday is the same as when she truthfully revealed her type to her proxy.

The third scenario in Table 2 demonstrates the outcome if Nancy were to misreport her value for the good by reporting an inflated value of \$10, *if the price matching scheme were for proxies to simply match their option price to the lowest winning price at any time while they are in the system*. Nancy would win the Monday auction and receive an option with an exercise price of \$8. On Tuesday, Polly would win the auction and receive an option with an exercise price of \$0. Nancy's proxy would observe that the lowest price on Tuesday was \$0, and so Nancy's proxy would price match the exercise price of its option down to \$0. Note that the exercise price Nancy's proxy has obtained at the end of Tuesday is lower than when she truthfully revealed her type to the proxy.

Therefore, a price matching policy of simply matching the lowest price paid may not elicit truthful information from buyers.

## 4.2 Complexity of Algorithm

An XOR-valuation of size $M$ for buyer $i$ is a set of $M$ terms, $< L^1, v_i^1 > ... < L^M, v_i^M >$, that maps distinct bundles to values, where $i$ is interested in acquiring at most one such bundle. For any bundle $S$, $v_i(S) = \max_{L^m \subseteq S}(v_i^m)$.

THEOREM 1. *Given an XOR-valuation which possesses $M$ terms, there is an $O(KM^2)$ algorithm for computing all maximum marginal values, where $K$ is the number of different item types in which a buyer may be interested.*

PROOF. For each item type, recall Equation 1 which defines the maximum marginal value of an item. For each bundle $L$ in the $M$-term valuation, $v_i(L + k)$ may be found by iterating over the $M$ terms. Therefore, the number of terms explored to determine the maximum marginal value for any item is $O(M^2)$, and so the total number of bundle comparisons to be performed to calculate all maximum marginal values is $O(KM^2)$. □

THEOREM 2. *The total memory required by a proxy for implementing price matching is $O(K)$, where $K$ is the number of distinct item types. The total work performed by a proxy to conduct price matching in each auction is $O(1)$.*

PROOF. By construction of the algorithm, the proxy stores one maximum marginal value for each item for bidding, of which there are $O(K)$; at most one buyer's identity for each item, of which there are $O(K)$; and one current option exercise price for each item, of which there are $O(K)$. For each auction, the proxy either submits a precomputed bid or price matches, both of which take $O(1)$ work. □

## 4.3 Truthful Bidding to the Proxy Agent

Proxies transform the market into a direct revelation mechanism, where each buyer $i$ interacts with the proxy only once,[15] and does so by declaring a bid, $b_i$, which is defined as an announcement of her type, $(\hat{a}_i, \hat{d}_i, \hat{v}_i)$, where the announcement may or may not be truthful. We denote all received bids other than $i$'s as $b_{-i}$. Given bids, $b = (b_i, b_{-i})$, the market determines allocations, $x_i(b)$, and payments, $p_i(b) \geq 0$, to each buyer (using an online algorithm).

A *dominant* strategy equilibrium for buyers requires that $v_i(x_i(b_i, b_{-i})) - p_i(b_i, b_{-i}) \geq v_i(x_i(b'_i, b_{-i})) - p_i(b'_i, b_{-i}), \forall b'_i \neq b_i, \forall b_{-i}$.

We now establish that it is a dominant strategy for a buyer to reveal her *true* valuation and *true* departure time to her proxy agent immediately upon arrival to the system. The proof builds on the price-based characterization of strategyproof single-item online auctions in Hajiaghayi et al. [12].

Define a monotonic and value-independent price function $ps_i(a_i, d_i, L, v_{-i})$ which can depend on the values of other agents $v_{-i}$. Price $ps_i(a_i, d_i, L, v_{-i})$ will represent the price available to agent $i$ for bundle $L$ in the mechanism if it announces arrival time $a_i$ and departure time $d_i$. The price is independent of the value $v_i$ of agent $i$, but can depend on $a_i$, $d_i$ and $L$ as long as it satisfies a monotonicity condition.

DEFINITION 2. *Price function $ps_i(a_i, d_i, L, v_{-i})$ is monotonic if $ps_i(a'_i, d'_i, L', v_{-i}) \leq ps_i(a_i, d_i, L, v_{-i})$ for all $a'_i \leq a_i$, all $d'_i \geq d_i$, all bundles $L' \subseteq L$ and all $v_{-i}$.*

LEMMA 1. *An online combinatorial auction will be strategyproof (with truthful reports of arrival, departure and value a dominant strategy) when there exists a monotonic and value-independent price function, $ps_i(a_i, d_i, L, v_{-i})$, such that for all $i$ and all $a_i, d_i \in T$ and all $v_i$, agent $i$ is allocated bundle $L^* = \text{argmax}_L\, [v_i(L) - ps_i(a_i, d_i, L, v_{-i})]$* ***in period*** *$d_i$ and makes payment $ps_i(a_i, d_i, L^*, v_{-i})$.*

PROOF. Agent $i$ cannot benefit from reporting a later departure $\hat{d}_i$ because the allocation is made in period $\hat{d}_i$ and the agent would have no value for this allocation. Agent $i$ cannot benefit from reporting a later arrival $\hat{a}_i \geq a_i$ or earlier departure $\hat{d}_i \leq d_i$ because of price monotonicity. Finally, the agent cannot benefit from reporting some $\hat{v}_i \neq v_i$ because its reported valuation does not change the prices it faces and the mechanism maximizes its utility given its reported valuation and given the prices. □

LEMMA 2. *At any given time, there is at most one buyer in the system whose proxy does not hold an option for a given item type because of buyer $i$'s presence in the system, and the identity of that buyer will be stored in $i$'s proxy's local memory at that time if such a buyer exists.*

PROOF. By induction. Consider the first proxy that a buyer prevents from winning an option. Either (a) the

[15]For analysis purposes, we view the mechanism as an opaque market so that the buyer cannot condition her bid on bids placed by others.

bumped proxy will leave the system having never won an option, or (b) the bumped proxy will win an auction in the future. If (a), the buyer's presence prevented exactly that one buyer from winning an option, but will have not prevented any other proxies from winning an option (as the buyer's proxy will not bid on additional options upon securing one), and will have had that bumped proxy's identity in its local memory by definition of the algorithm. If (b), the buyer has not prevented the bumped proxy from winning an option after all, but rather has prevented only the proxy that lost to the bumped proxy from winning (if any), whose identity will now be stored in the proxy's local memory by definition of the algorithm. For this new identity in the buyer's proxy's local memory, either scenario (a) or (b) will be true, ad infinitum. □

Given this, we show that the options-based infrastructure implements a price-based auction with a monotonic and value-independent price schedule to every agent.

THEOREM 3. *Truthful revelation of valuation, arrival and departure is a dominant strategy for a buyer in the options-based market.*

PROOF. First, define a simple agent-independent price function $p_i^k(t, v_{-i})$ as the highest bid by the proxies not holding an option on an item of type $G_k$ at time $t$, not including the proxy representing $i$ herself and not including any proxies that would have already won an option had $i$ never entered the system (i.e., whose identity is stored in $i$'s proxy's local memory)($\infty$ if no supply at $t$). This set of proxies is independent of any declaration $i$ makes to its proxy (as the set explicitly excludes the at most one proxy (see Lemma 2) that $i$ has prevented from holding an option), and each bid submitted by a proxy within this set is only a function of their own buyer's declared valuation (see Equation 1). Furthermore, $i$ cannot influence the supply she faces as any options returned by bidders due to a price set by $i$'s proxy's bid will be re-auctioned after $i$ has departed the system. Therefore, $p_i^k(t, v_{-i})$ is independent of $i$'s declaration to its proxy. Next, define $ps_i^k(\hat{a}_i, \hat{d}_i, v_{-i}) = \min_{\hat{a}_i \leq \tau \leq \hat{d}_i}[p_i^k(\tau, v_{-i})]$ (possibly $\infty$) as the minimum price over $p_i^k(t, v_{-i})$, which is clearly monotonic. By construction of price matching, this is exactly the price obtained by a proxy on any option that it holds at departure. Define $ps_i(\hat{a}_i, \hat{d}_i, L, v_{-i}) = \sum_{k=1}^{k=K} ps_i^k(\hat{a}_i, \hat{d}_i, v_{-i})L_k$, which is monotonic in $\hat{a}_i$, $\hat{d}_i$ and $L$ since $ps_i^k(\hat{a}_i, \hat{d}_i, v_{-i})$ is monotonic in $\hat{a}_i$ and $\hat{d}_i$ and (weakly) greater than zero for each $k$. Given the set of options held at $\hat{d}_i$, which may be a subset of those items with non-infinite prices, the proxy exercises options to maximize the reported utility. Left to show is that all bundles that could not be obtained with options held are priced sufficiently high as to not be preferred. For each such bundle, either there is an item priced at $\infty$ (in which case the bundle would not be desired) or there must be an item in that bundle for which the proxy does not hold an option that was available. In all auctions for such an item there must have been a distinct bidder with a bid greater than $bid_i^t(k)$, which subsequently results in $ps_i^k(\hat{a}_i, \hat{d}_i, v_{-i}) > bid_i^t(k)$, and so the bundle without $k$ would be preferred to the bundle. □

THEOREM 4. *The super proxy, options-based scheme is individually-rational for both buyers and sellers.*

| | $\overline{\text{Price}}$ | $\sigma$(Price) | $\overline{\text{Value}}$ | $\overline{\text{Surplus}}$ |
|---|---|---|---|---|
| eBay | \$240.24 | \$32 | \$244 | \$4 |
| Options | \$239.66 | \$12 | \$263 | \$23 |

**Table 3: Average price paid, standard deviation of prices paid, average bidder value among winners, and average winning bidder surplus on eBay for Dell E193FP LCD screens as well as the simulated options-based market using worst-case estimates of bidders' true value.**

PROOF. By construction, the proxy exercises the profit maximizing set of options obtained, or no options if no set of options derives non-negative surplus. Therefore, buyers are guaranteed non-negative surplus by participating in the scheme. For sellers, the price of each option is based on a non-negative bid or zero. □

# 5. EVALUATING THE OPTIONS / PROXY INFRASTRUCTURE

A goal of the empirical benchmarking and a reason to collect data from eBay is to try and build a realistic model of buyers from which to estimate seller revenue and other market effects under the options-based scheme.

We simulate a sequence of auctions that match the timing of the Dell LCD auctions on eBay.[16] When an auction successfully closes on eBay, we simulate a Vickrey auction for an option on the item sold in that period. Auctions that do not successfully close on eBay are not simulated. We estimate the arrival, departure and value of each bidder on eBay from their observed behavior.[17] Arrival is estimated as the first time that a bidder interacts with the eBay proxy, while departure is estimated as the latest closing time among eBay auctions in which a bidder participates.

We initially adopt a particularly conservative estimate for bidder value, estimating it as the highest bid a bidder was observed to make on eBay. Table 3 compares the distribution of closing prices on eBay and in the simulated options scheme. While the average revenue in both schemes is virtually the same (\$239.66 in the options scheme vs. \$240.24 on eBay), the winners in the options scheme tend to value the item won 7% more than the winners on eBay (\$263 in the options scheme vs. \$244 on eBay).

## 5.1 Bid Identification

We extend the work of Haile and Tamer [11] to sequential auctions to get a better view of underlying bidder values. Rather than assume for bidders an equilibrium behavior as in standard econometric techniques, Haile and Tamer do not attempt to model *how* bidders' true values get mapped into a bid in any given auction. Rather, in the context of repeated

[16]When running the simulations, the results of the first and final ten days of auctions are not recorded to reduce edge effects that come from viewing a discrete time window of a continuous process.

[17]For the 100 bidders that won multiple times on eBay, we have each one bid a constant marginal value for each additional item in each auction until the number of options held equals the total number of LCDs won on eBay, with each option available for price matching independently. This bidding strategy is not a dominant strategy (falling outside the type space possible for buyers on which the proof of truthfulness has been built), but is believed to be the most appropriate first order action for simulation.

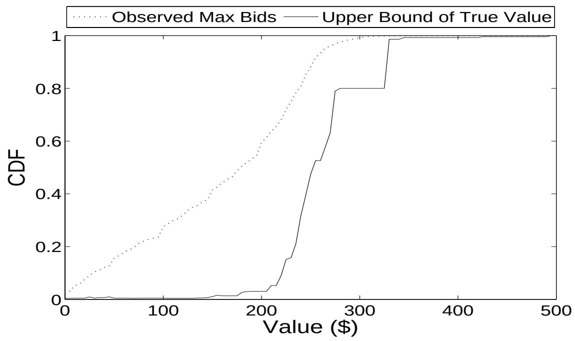


**Figure 2: CDF of maximum bids observed and upper bound estimate of the bidding population's distribution for maximum willingness to pay. The true population distribution lies below the estimated upper bound.**

single-item auctions with distinct bidder populations, Haile and Tamer make only the following two assumptions when estimating the distribution of true bidder values:

1. Bidders do not bid more than they are willing to pay.
2. Bidders do not allow an opponent to win at a price they are willing to beat.

From the first of their two assumptions, given the bids placed by each bidder in each auction, Haile and Tamer derive a method for estimating an upper bound of the bidding population's true value distribution (i.e., the bound that lies above the true value distribution). From the second of their two assumptions, given the winning price of each auction, Haile and Tamer derive a method for estimating a lower bound of the bidding population's true value distribution. It is only the upper-bound of the distribution that we utilize in our work.

Haile and Tamer assume that bidders only participate in a single auction, and require independence of the bidding population from auction to auction. Neither assumption is valid here: the former because bidders are known to bid in more than one auction, and the latter because the set of bidders in an auction is in all likelihood not a true i.i.d. sampling of the overall bidding population. In particular, those who win auctions are less likely to bid in successive auctions, while those who lose auctions are more likely to remain bidders in future auctions.

In applying their methods we make the following adjustments:

- Within a given auction, each individual bidder's true willingness to pay is assumed weakly greater than the maximum bid that bidder submits across all auctions for that item (either past or future).
- When estimating the upper bound of the value distribution, if a bidder bids in more than one auction, randomly select one of the auctions in which the bidder bid, and only utilize that one observation during the estimation.[18]

| | $\overline{\text{Price}}$ | $\sigma$(Price) | $\overline{\text{Value}}$ | $\overline{\text{Surplus}}$ |
|---|---|---|---|---|
| eBay | $240.24 | $32 | $281 | $40 |
| Options | $275.80 | $14 | $302 | $26 |

**Table 4: Average price paid, standard deviation of prices paid, average bidder value among winners, and average winning bidder surplus on eBay for Dell E193FP LCD screens as well as in the simulated options-based market using an adjusted Haile and Tamer estimate of bidders' true values being 15% higher than their maximum observed bid.**

Figure 2 provides the distribution of maximum bids placed by bidders on eBay as well as the estimated upper bound of the true value distribution of bidders based on the extended Haile and Tamer method.[19] As can be seen, the smallest relative gap between the two curves meaningfully occurs near the 80th percentile, where the upper bound is 1.17 times the maximum bid. Therefore, adopted as a less conservative model of bidder values is a uniform scaling factor of 1.15.

We now present results from this less conservative analysis. Table 4 shows the distribution of closing prices in auctions on eBay and in the simulated options scheme. The mean price in the options scheme is now significantly higher, 15% greater, than the prices on eBay ($276 in the options scheme vs. $240 on eBay), while the standard deviation of closing prices is lower among the options scheme auctions ($14 in the options scheme vs. $32 on eBay). Therefore, not only is the expected revenue stream higher, but the lower variance provides sellers a greater likelihood of realizing that higher revenue.

The efficiency of the options scheme remains higher than on eBay. The winners in the options scheme now have an average estimated value 7.5% higher at $302.

In an effort to better understand this efficiency, we formulated a mixed integer program (MIP) to determine a simple estimate of the allocative efficiency of eBay. The MIP computes the efficient value of the offline problem with full hindsight on all bids and all supply.[20] Using a scaling of 1.15, the total value allocated to eBay winners is estimated at $551,242, while the optimal value (from the MIP) is $593,301. This suggests an allocative efficiency of 92.9%: while the typical value of a winner on eBay is $281, an average value of $303 was possible.[21] Note the options-based

[18] In current work, we assume that removing duplicate bidders is sufficient to make the buying populations independent i.i.d. draws from auction to auction. If one believes that certain portions of the population are drawn to certain auctions, then further adjustments would be required in order to utilize these techniques.

[19] The estimation of the points in the curve is a minimization over many variables, many of which can have small-numbers bias. Consequently, Haile and Tamer suggest using a weighted average over all terms $y_i$ of $\sum_i y_i \frac{\exp(y_i \rho)}{\sum_j \exp(y_j \rho)}$ to approximate the minimum while reducing the small number effects. We used $\rho = -1000$ and removed observations of auctions with 17 bidders or more as they occurred very infrequently. However, some small numbers bias still demonstrated itself with the plateau in our upper bound estimate around a value of $300.

[20] Buyers who won more than one item on eBay are cloned so that they appear to be multiple bidders of identical type.

[21] As long as one believes that every bidder's true value is a constant factor $\alpha$ away from their observed maximum bid, the 92.9% efficiency calculation holds for any value of $\alpha$. In practice, this belief may not be reasonable. For example, if losing bidders tend to have true values close to their observed

scheme comes very close to achieving this level of efficiency [at 99.7% efficient in this estimate] even though it operates without the benefit of hindsight.

Finally, although the typical winning bidder surplus decreases between eBay and the options-based scheme, some surplus redistribution would be possible because the total market efficiency is improved.[22]

# 6. DISCUSSION

The biggest concern with our scheme is that proxy agents who may be interested in many different items may acquire many more options than they finally exercise. This can lead to efficiency loss. Notice that this is not an issue when bidders are only interested in a single item (as in our empirical study), or have linear-additive values on items.

To fix this, we would prefer to have proxy agents use more caution in acquiring options and use a more adaptive bidding strategy than that in Equation 1. For instance, if a proxy is already holding an option with an exercise price of $3 on some item for which it has value of $10, and it values some substitute item at $5, the proxy could reason that in no circumstance will it be useful to acquire an option on the second item.

We formulate a more sophisticated bidding strategy along these lines. Let $\Theta_t$ be the set of all options a proxy for bidder $i$ already possesses at time $t$. Let $\theta_t \subseteq \Theta_t$, be a subset of those options, the sum of whose exercise prices are $\pi(\theta_t)$, and the goods corresponding to those options being $\gamma(\theta_t)$. Let $\Pi(\theta_t) = \hat{v}_i(\gamma(\theta_t)) - \pi(\theta_t)$ be the (reported) available surplus associated with a set of options. Let $\theta_t^*$ be the set of options currently held that would maximize the buyer's surplus; i.e., $\theta_t^* = \text{argmax}_{\theta_t \subseteq \Theta_t} \Pi(\theta_t)$.

Let the *maximal willingness to pay* for an item $k$ represent a price above which the agent knows it would never exercise an option on the item *given the current options held*. This can be computed as follows:

$$bid_i^t(k) = \max_L [0, \min[\hat{v}_i(L+k) - \Pi(\theta_t^*), \hat{v}_i(L+k) - \hat{v}_i(L)]] \quad (3)$$

where $\hat{v}_i(L+k) - \Pi(\theta_t^*)$ considers surplus already held, $\hat{v}_i(L+k) - \hat{v}_i(L)$ considers the marginal value of a good, and taking the max[0, .] considers the overall use of pursuing the good.

However, and somewhat counter intuitively, we are not able to implement this bidding scheme without forfeiting truthfulness. The $\Pi(\theta_t^*)$ term in Equation 3 (i.e., the amount of guaranteed surplus bidder $i$ has already obtained) can be influenced by proxy $j$'s bid. Therefore, bidder $j$ may have the incentive to misrepresent her valuation to her proxy if she believes doing so will cause $i$ to bid differently in the future in a manner beneficial to $j$. Consider the following example where the proxy scheme is refined to bid the maximum willingness to pay.

Example 3. *Alice values either one ton of Sand or one ton of Stone for $2,000. Bob values either one ton of Sand or one ton of Stone for $1,500. All bidders have a patience of 2 days. On day one, a Sand auction is held, where Alice's proxy bids $2,000 and Bob's bids $1,500. Alice's proxy wins an option to purchase Sand for $1,500. On day two, a Stone auction is held, where Alice's proxy bids $1,500 [as she has already obtained a guaranteed $500 of surplus from winning a Sand option, and so reduces her Stone bid by this amount], and Bob's bids $1,500. Either Alice's proxy or Bob's proxy will win the Stone option. At the end of the second day, Alice's proxy holds an option with an exercise price of $1,500 to obtain a good valued for $2,000, and so obtains $500 in surplus.*

Now, consider what would have happened had Alice declared that she valued only Stone.

Example 4. *Alice declares valuing only Stone for $2,000. Bob values either one ton of Sand or one ton of Stone for $1,500. All bidders have a patience of 2 days. On day one, a Sand auction is held, where Bob's proxy bids $1,500. Bob's proxy wins an option to purchase Sand for $0. On day two, a Stone auction is held, where Alice's proxy bids $2,000, and Bob's bids $0 [as he has already obtained a guaranteed $1,500 of surplus from winning a Sand option, and so reduces his Stone bid by this amount]. Alice's proxy wins the Stone option for $0. At the end of the second day, Alice's proxy holds an option with an exercise price of $0 to obtain a good valued for $2,000, and so obtains $2,000 in surplus.*

By misrepresenting her valuation (i.e., excluding her value of Sand), Alice was able to secure higher surplus by guiding Bob's bid for Stone to $0. An area of immediate further work by the authors is to develop a more sophisticated proxy agent that can allow for bidding of maximum willingness to pay (Equation 3) while maintaining truthfulness.

An additional, practical, concern with our proxy scheme is that we assume an available, trusted, and well understood method to characterize goods (and presumably the quality of goods). We envision this happening in practice by sellers defining a classification for their item upon entering the market, for instance via a UPC code. Just as in eBay, this would allow an opportunity for sellers to improve revenue by overstating the quality of their item ("new" vs. "like new"), and raises the issue of how well a reputation scheme could address this.

# 7. CONCLUSIONS

We introduced a new sales channel, consisting of an options-based and proxied auction protocol, to address the sequential auction problem that exists when bidders face multiple auctions for substitutes and complements goods. Our scheme provides bidders with a simple, dominant and truthful bidding strategy even though the market remains open and dynamic.

In addition to exploring more sophisticated proxies that bid in terms of maximum willingness to pay, future work should aim to better model seller incentives and resolve the strategic problems facing sellers. For instance, does the options scheme change seller incentives from what they currently are on eBay?


## Acknowledgments

We would like to thank Pai-Ling Yin. Helpful comments have been received from William Simpson, attendees at Har-


maximum bids while eBay winners have true values much greater than their observed maximum bids then downward bias is introduced in the efficiency calculation at present.

[22]The increase in eBay winner surplus between Tables 3 and 4 is to be expected as the $\alpha$ scaling strictly increases the estimated value of the eBay winners while holding the prices at which they won constant.

vard University's EconCS and ITM seminars, and anonymous reviewers. Thank you to Aaron L. Roth and Kang-Xing Jin for technical support. All errors and omissions remain our own.

## 8. REFERENCES


[1] P. Anthony and N. R. Jennings. Developing a bidding agent for multiple heterogeneous auctions. *ACM Trans. On Internet Technology*, 2003.

[2] R. Bapna, P. Goes, A. Gupta, and Y. Jin. User heterogeneity and its impact on electronic auction market design: An empirical exploration. *MIS Quarterly*, 28(1):21–43, 2004.

[3] D. Bertsimas, J. Hawkins, and G. Perakis. Optimal bidding in on-line auctions. Working Paper, 2002.

[4] C. Boutilier, M. Goldszmidt, and B. Sabata. Sequential auctions for the allocation of resources with complementarities. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 527–534, 1999.

[5] A. Byde, C. Preist, and N. R. Jennings. Decision procedures for multiple auctions. In *Proc. 1st Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS-02)*, 2002.

[6] M. M. Bykowsky, R. J. Cull, and J. O. Ledyard. Mutually destructive bidding: The FCC auction design problem. *Journal of Regulatory Economics*, 17(3):205–228, 2000.

[7] Y. Chen, C. Narasimhan, and Z. J. Zhang. Consumer heterogeneity and competitive price-matching guarantees. *Marketing Science*, 20(3):300–314, 2001.

[8] A. K. Dixit and R. S. Pindyck. *Investment under Uncertainty*. Princeton University Press, 1994.

[9] R. Gopal, S. Thompson, Y. A. Tung, and A. B. Whinston. Managing risks in multiple online auctions: An options approach. *Decision Sciences*, 36(3):397–425, 2005.

[10] A. Greenwald and J. O. Kephart. Shopbots and pricebots. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 506–511, 1999.

[11] P. A. Haile and E. Tamer. Inference with an incomplete model of English auctions. *Journal of Political Economy*, 11(1), 2003.

[12] M. T. Hajiaghayi, R. Kleinberg, M. Mahdian, and D. C. Parkes. Online auctions with re-usable goods. In *Proc. ACM Conf. on Electronic Commerce*, 2005.

[13] K. Hendricks, I. Onur, and T. Wiseman. Preemption and delay in eBay auctions. *University of Texas at Austin Working Paper*, 2005.

[14] A. Iwasaki, M. Yokoo, and K. Terada. A robust open ascending-price multi-unit auction protocol against false-name bids. *Decision Support Systems*, 39:23–40, 2005.

[15] E. G. James D. Hess. Price-matching policies: An empirical case. *Managerial and Decision Economics*, 12(4):305–315, 1991.

[16] A. X. Jiang and K. Leyton-Brown. Estimating bidders' valuation distributions in online auctions. In *Workshop on Game Theory and Decision Theory (GTDT) at IJCAI*, 2005.

[17] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC-00)*, 2000.

[18] Y. J. Lin. Price matching in a model of equilibrium price dispersion. *Southern Economic Journal*, 55(1):57–69, 1988.

[19] D. Lucking-Reiley and D. F. Spulber. Business-to-business electronic commerce. *Journal of Economic Perspectives*, 15(1):55–68, 2001.

[20] A. Ockenfels and A. Roth. 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):1093–1103, 2002.

[21] M. Peters and S. Severinov. Internet auctions with many traders. *Journal of Economic Theory (Forthcoming)*, 2005.

[22] R. Porter. Mechanism design for online real-time scheduling. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 61–70. ACM Press, 2004.

[23] M. H. Rothkopf and R. Engelbrecht-Wiggans. Innovative approaches to competitive mineral leasing. *Resources and Energy*, 14:233–248, 1992.

[24] T. Sandholm and V. Lesser. Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35:212–270, 2001.

[25] T. W. Sandholm and V. R. Lesser. Issues in automated negotiation and electronic commerce: Extending the Contract Net framework. In *Proc. 1st International Conference on Multi-Agent Systems (ICMAS-95)*, pages 328–335, 1995.

[26] H. S. Shah, N. R. Joshi, A. Sureka, and P. R. Wurman. Mining for bidding strategies on eBay. Lecture Notes on Artificial Intelligence, 2003.

[27] M. Stryszowska. Late and multiple bidding in competing second price Internet auctions. *EuroConference on Auctions and Market Design: Theory, Evidence and Applications*, 2003.

[28] J. T.-Y. Wang. Is last minute bidding bad? *UCLA Working Paper*, 2003.

[29] R. Zeithammer. An equilibrium model of a dynamic auction marketplace. *Working Paper, University of Chicago*, 2005.