Artificial Intelligence ••• (••••) •••–•••

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

An options-based solution to the sequential auction problem ☆

Adam I. Juda [1], David C. Parkes *,[2]

School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA

ARTICLE INFO

Article history:
Received 2 October 2007
Received in revised form 1 January 2009
Accepted 4 January 2009
Available online xxxx

Keywords:
Coordination problems
Preferences
Dynamic auctions
Proxy agents
Options
Strategyproofness
Electronic markets

ABSTRACT

The *sequential auction problem* is commonplace in open, electronic marketplaces such as eBay. This is the problem where a buyer has no dominant strategy in bidding across multiple auctions when the buyer would have a simple, truth-revealing strategy if there was but a single auction event. Our model allows for multiple, distinct goods and market dynamics with buyers and sellers that arrive over time. Sellers each bring a single unit of a good to the market while buyers can have values on bundles of goods. We model each individual auction as a second-price (Vickrey) auction and propose an *options-based, proxied* solution to provide price and winner-determination coordination across auctions. While still allowing for temporally uncoordinated market participation, this options-based approach solves the sequential auction problem and provides truthful bidding as a weakly dominant strategy for buyers. An empirical study suggests that this coordination can enable a significant efficiency and revenue improvement over the current eBay market design, and highlights the effect on performance of complex buyer valuations (buyers with *substitutes* and *complements* valuations) and varying the market liquidity.

1. Introduction

Electronic markets generate significant new trading opportunities and expand the opportunity for the dynamic pricing of goods, and lead to improved market efficiency in many settings [11,12,64]. Electronic markets find application not only for person-to-person transactions (e.g., auctions), but also increasingly for business-to-consumer auctions such as selling surplus inventory [39] and business-to-business sourcing events [59].

But despite the new efficiencies offered by electronic markets, for example by enabling the application of optimization to decision making in many parts of the supply chain, the role of *automated trading agents* – although long envisioned by artificial intelligence researchers [1,26,60] – remains more fiction than reality. (One notable exception is the significant role of automated trading for financial securities [23].) One major impediment to the adoption of automated trading agents is that users often have insufficient trust of software agents that would work on their behalf [16,37]. Users may even have higher expectations for automated agents than human agents in regard to what constitutes acceptable behavior [62]. For example, the proposal by the London International Financial Futures Exchange (Liffe) to introduce automated trading was

☆ Significantly expanded from earlier versions of this paper that appeared in the Seventh ACM Conference on Electronic Commerce, 2006 and in the 6th Workshop on Agent-Mediated Electronic Commerce, 2004.

* Corresponding author.
E-mail addresses: adamjuda@post.harvard.edu (A.I. Juda), parkes@eecs.harvard.edu (D.C. Parkes).

[1] This work developed while a PhD Candidate in the Information, Technology and Management Program at the Harvard Business School and School of Engineering and Applied Sciences, Harvard University.

[2] Supported in part by NSF grant IIS-0238147.

0004-3702/$ – see front matter
doi:10.1016/j.artint.2009.01.002

the cause of much debate.[3] A key area of concern was that users could not be certain that the automated systems would behave "optimally" in all situations.

One important, if unsurprising, aspect in the adoption of automated agents is that these agents avoid, and be seen to avoid, making mistakes [47]. In the context of electronic markets, an insight offered by *mechanism design*, the subdiscipline of microeconomics that seeks to design protocols to achieve system-wide objectives with self-interested agents, is that one can sometimes design protocols that *simplify* the strategic considerations of bidding agents. Classic solutions, such as those offered by the Vickrey–Clarke–Groves (VCG) mechanism [20,27,67], provide this simplicity via the property of *strategyproofness* (which brings truthful bidding into a dominant strategy equilibrium), and in addition provide market efficiency. But we argue that they are often not applicable in practice because they require too much temporal coordination on the part of participants. Electronic markets such as eBay[4] allow a would-be seller to decide when to sell goods in the marketplace, and buyers can visit the marketplace at times of their own choosing. The VCG mechanism on the other hand requires that buyers and sellers be grouped into a single, coordinated auction with all bids placed at the same time and all goods sold at the same time.

In fact, the individual auctions on eBay are very similar to *single-item* second-price (Vickrey) auctions. The most significant difference is that eBay provides, via the use of a mandatory proxy agent that bids on behalf of a buyer, a "staged Vickrey auction" such that a buyer can effectively increase her bid price at any time until an auction closes.[5] Electronic markets such as eBay do not provide a large-scale combinatorial auction, structured as a VCG mechanism (e.g., for all LCD monitors in the market), arguably because this would require too much temporal coordination on the part of buyers and sellers. A second issue would be determining how to define the scope for such an event to include a suitable domain of goods likely to subsume most of those of interest to a set of potential buyers. A third issue is that at some point the computational cost for running such a large, coordinated mechanism would also get prohibitive given that winner-determination for combinatorial auctions is **NP**-hard [56]. Problems of communication complexity and preference elicitation could also become a concern [57,63].

The absence of such a large-scale coordination mechanism in markets such as eBay leads to strategic complexity for participants. Despite a (weak) strategic equivalence between an individual auction on eBay and the Vickrey auction [40,65], there are many reasons for buyers not to truthfully bid their value in any one auction. One reason follows from the auction being staged rather than sealed-bid. Because some bidders may be "followers" – bidding up when others do – it can be rational to delay until the last minute and "snipe" to avoid driving up competitors' bid prices [4,46]. The *sequential auction problem* provides another reason. This relates to issues that arise when composing strategies across a sequence of auctions.

For example, suppose that multiple copies of essentially identical items are offered for sale sequentially. For example, Alice may want an LCD monitor, and could potentially bid in either a 1 o'clock or 3 o'clock auction. Alice would prefer to participate in the auction that will have the lower winning price, but she cannot determine beforehand which auction that will be. As a result, she could end up winning in the "wrong" auction, that is the auction with the higher price. A related example of the sequential auction problem is familiar from the *exposure problem* studied in *simultaneous* ascending price auctions [14], which also exists in our setting when a buyer desires a bundle of goods but must participate in auctions on individual items. 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, she must determine how much of the \$20 of synergy value to include in her bid when bidding for the console alone. If Alice incorporates some of the synergy value (e.g., by placing a bid of \$210 and paying her bid), she may incur a loss if she can not subsequently win the video game for less than \$40.[6]

The main technical question addressed in this paper is: *can one design a marketplace for temporally uncoordinated buyers and sellers, and distinct goods, in which buyers have a simple, dominant bidding strategy*? As a solution we propose a *real-options based* market infrastructure, coupled with *proxy bidding*, that enables simple, yet optimal, bidding strategies while retaining the dynamic arrivals and departures that are a defining feature of electronic markets such as eBay. Our main assumptions (in increasing order of strength) are that:

- Each buyer has an arrival time and a departure time in the market, and is indifferent between buying items at any time before her departure and with zero value after her departure.
- Buyers may have general valuations on bundles of distinct goods, but are interested in at most one unit of each of these goods.

[3] The Financial Times said at the time that "*Electronic trading is the biggest single issue to face the futures community today and the industry has long confronted a philosophical split on its merits*", in an article "Liffe's new automated trading system has sparked a debate on automated trading", November 30, 1989.

[4] www.ebay.com.

[5] While an eBay auction is open, a buyer provides her automated proxy agent with a bid ceiling. While the ceiling the agent has received is greater than the winning price and the agent is not winning, the agent will submit a bid some amount, ϵ, above the current winning price (where $\epsilon > 0$ is set by eBay anywhere from cents to dollars depending on the value of the item). Therefore, when an auction ends, the winning buyer will pay a price ϵ above the highest ceiling another buyer submitted, and the outcome is nearly identical to the outcome of the Vickrey auction.

[6] A third reason for strategic complexity in markets such as eBay can be that the quality of items may be uncertain and buyers may adjust their belief about the value of the item based on others' bids; this is the so-called *interdependent* values model of auction theory. We do not consider interdependent value domains in this paper. This makes our results applicable instead to markets in which buyers know their value for goods and can determine this without seeing the bids from others. An example of such a domain is provided by our empirical study on the eBay market for a Dell LCD monitor.

- Each seller offers a single unit of a good, has an arrival time and departure time in the market, has no intrinsic value for the good, and is willing to wait until the departure time of a buyer that wins no later than her departure to receive payment.
- Sellers are non-strategic and truthfully report arrival and departure to the market, which defines the interval of time during which they are willing to sell the good.

We will not need to assume that buyers can only participate under one identity or prevent buyers from re-entering the market; this provides robustness to the false-name bidding considered by Yokoo and colleagues [70]. On the sell side of the market, we do not allow sellers to have a hard limit on the latest period by which they require payment, although this will exist because we will focus on markets in which a maximum patience (departure-arrival) can be assumed for buyers.

In general, the options-based market will perform best when many goods can readily be categorized into few equivalence classes such that every buyer is indifferent across goods in the class. While in a worst case there could be as many equivalence classes as auctions in a market place, many items listed on eBay are essentially identical to other items, and especially in categories such as *Consumer Electronics*, where the sum of all successfully closed listings during 2005 was U.S. \$3.5B (of U.S. \$44B in total for all of eBay) [25]. This category is the focus of our empirical analysis, in which we consider auctions for 19″ Dell LCD monitors (Model E193FP) conducted on eBay during the summer of 2005. Moreover, individual buyers may of course view goods in different classes as equivalent and no two buyers need agree on the actual value that they assign to a good in any particular class.[7]

In our options-based solution, a seller auctions an option for her good rather than auctioning the good directly. The option will ultimately either lead to a sale or require the seller to return to the market and offer another option on the same good if so interested. By participation in the framework, sellers agree to allow proxy agents to *price-match* their goods against others of equal type, with the payment a seller finally receives defined in terms of the *minimal price that the winning bidder could have bid and still traded with some seller in some auction for a good of equal type during the bidder's arrival-departure interval*. As noted above, all sellers are assumed to behave non-strategically and to truthfully report their temporal constraints in the marketplace. Furthermore, we assume that sellers do not have an intrinsic value for the good.

All buyers in our framework must interact through a mandatory proxy agent, and do so by reporting a value on all possible bundles of goods of interest along with a departure time. While such an enumeration may seem daunting at first glance, there are several reasons not to view this as a major concern in consumer markets. First, a very common purchasing scenario is for a buyer to want a single item or to be indifferent among only a few different items. Second, Cantillon and Pesendorfer [15] and Sandholm [59] provide empirical support that buyers can manage to construct bids in large combinatorial settings. Third, a number of expressive, concise bidding languages have been developed for combinatorial auctions [9,17,45]. It is also possible to allow buyers to provide lower-bounds on values, and increase these bounds over time just as with the mandatory proxy bidding agents of eBay (see Section 6 for further discussion on this point).

A buyer's proxy agent uses the reported information about value and departure time to determine how to bid for options and also to determine which options held at the buyer's reported departure time to exercise. The options that maximize the buyer's surplus given the reported valuation are exercised and all other options returned to the sellers. The options-based protocol is useful because it makes *truthful and immediate revelation to a proxy a dominant strategy for buyers, whatever the future auction dynamics*. Thus it can be seen as a method to generalize eBay's existing proxy scheme to handle the sequential auction problem in suitable categories of goods, while extending to embrace dynamic, combinatorial auctions.

In addressing questions about market efficiency and revenue we perform an empirical analysis using data on eBay auctions for 19″ Dell LCD monitors (Model E193FP) sold from 27 May, 2005 through 1 October, 2005. A conservative estimate is that an improvement in efficiency and revenue of around 4% and 9% respectively would be enabled through an options-based scheme. This estimate is generated on the basis of non-parametric estimation of the true value of buyers for items, generalizing a method due to Haile and Tamer [28] to sequential auctions. The eBay analysis also informs an extensive set of simulation experiments, in which we explore the effect of *substitutes* ("I want A or B") and *complements* ("I only want A if I also get B") valuations on the efficiency of the options-based scheme and also consider the impact of market liquidity. Buyer populations with substitutes valuations can hamper the efficiency of the marketplace because of hold-up problems in which a buyer's proxy holds a number of options that ultimately go unexercised but were unavailable to other buyers. However we find that this effect is mitigated when individual buyers have negatively correlated values across items. We also find that the buyer-to-seller ratio (a measure of *liquidity*) plays a critical role in market efficiency in the context of substitutes valuations. Efficiency first decreases and then increases as the buyer-to-seller ratio increases and the market becomes more competitive. For a low buyer-to-seller ratio the market remains efficient with substitutes valuations because there is plenty of supply. Market efficiency is also high for relatively large buyer-to-seller ratios (above 4:1 in our simulations) and substitutes valuations because increased competition segments the market; buyers tend to be competitive on only a small number of goods. In the context of *complements* valuations we find that market efficiency is fairly insensitive to positive or negative correlation in value across items and remains reasonably high.

[7] In the absence of a third party logistics partner, such as Amazon, that offers fulfillment and commits to the quality of a good (e.g., new, and "in box") it is likely that sellers could improve revenue in the short-term by overstating the quality of their item and misleading buyers in the marketplace. However, and just as on eBay, a well-functioning reputation system should mitigate this concern [54].

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

We also provide a *worst-case* analysis for the practically important special case in which buyers are interested in only one kind of good. This analysis is a slight extension of that due to Hajiaghayi et al. [29]. In our analysis, we parameterize the competitive ratio in terms of the maximal spread of upper and lower valuations across all buyers. The main focus of our evaluation remains empirical, and we do not provide any worst-case analysis for the general, combinatorial case.

By providing a system in which buyers possess a simple, dominant strategy, the proxy-based solution arguably reduces the participation costs of buyers. While the impact of this improvement is hard to estimate, this can be expected to both improve buyer loyalty to sellers and also make the market more appealing for new entrants. These two effects ought to preserve and enhance the health of the market and maintain seller revenue in the long term.

In outline, we first discuss related work, and then use Section 2 to introduce the model and define and characterize the sequential auction problem. Section 3 describes the options-based scheme, giving examples and a complexity analysis. Section 4 provides a strategic analysis and a worst-case efficiency analysis. Section 5 presents an experimental study, first on eBay data (the LCD market) and then extending – in simulation – to consider substitutes and complements preferences and the effect of market liquidity. Section 6 discusses the challenges of allowing proxy agents to bid less aggressively, thus mitigating hold-up. Section 7 concludes.

1.1. Related work

A number of authors have analyzed sequential auctions selling the same item with buyers interested in buying just a single item. This "multiple copies problem" is often studied in the context of explaining sniping behavior; see also Ockenfels and Roth [46], who give a collusion-based explanation for sniping. From the perspective of developing models for the multiple copies problem, Stryszowska [66] models the problem as one of a dynamic multi-unit auction, allowing for explanations of sniping as well as for bidding multiple times within an auction. Hendricks et al. [32] demonstrate that sniping is a symmetric equilibrium in the absence of a "BuyItNow" opportunity (wherein a buyer can choose to buy an item at any point for some fixed price). Wang [68] demonstrates, using a two-period model, how sniping in the first period is a unique equilibrium and Zeithammer [71] provides an equilibrium model for strategic sellers and forward-looking buyers. Notably, none of this prior work considers buyers that are able to participate in more than two auctions, while we consider settings in which buyers may participate in an arbitrary number of auctions. Peters and Severinov [52] also allow this more general capability and characterize a perfect Bayesian equilibrium where sellers set a reserve price equal to their true costs. These authors consider neither buyers entering at random times nor auctions closing at different times, both of which are addressed in our work. While these papers provide a Bayesian–Nash analysis of models that approximate current eBay-like markets, we study the dominant strategy equilibrium in an options-based variation on current markets.

Problems of the same kind as the sequential auction problem were previously observed by Wellman and Wurman [69] in the context of boundaries between multiple mechanisms, and later discussed by Parkes [48] and Ng et al. [43,44] in the context of "strategyproof computing". The problem has often been identified in the context of simultaneous ascending price auctions, where it is termed the *exposure problem* [14]. Previous work addressing the exposure problem has considered two different directions. First, one can change the mechanism and define an expressive bidding language and a strategyproof mechanism, as seen in work on combinatorial auctions [56]. Second, one can attempt to provide automated bidding agents with sophisticated strategies, as seen for example in the work of Boutilier et al. [8], Byde et al. [13], Anthony and Jennings [1], Reeves et al. [53], and Gerding et al. [24]. Unfortunately, it seems hard to design artificial agents with equilibrium bidding strategies, even for a simultaneous ascending price auction (i.e., without dynamic arrivals of new sellers) and all these papers make significant assumptions.

Iwasaki et al. [34] have previously considered the use of options in the context of a single, monolithic, auction design to help buyers with increasing marginal values avoid exposure in a multi-unit, homogeneous item auction. Sandholm and Lesser [61] have considered options in the form of *leveled commitment contracts* for facilitating multi-way recontracting in a completely decentralized market place. Rothkopf and Engelbrecht-Wiggans [55] discuss the advantages associated with the use of options for selling coal mine leases. To the best of our knowledge, ours is the first work to study the role of options as a method to enable dominant strategies in the context of dynamic auctions. Gopal et al. [25] have considered the use of options for reducing exposure to risk in the context of the sequential auction problem. Our work differs in a number of ways, including how the options are priced, which buyers obtain options, and in how much risk remains with buyers once options are used. Buyers still face risk and have no dominant strategy in the method of Gopal et al. [25].

The technical contribution of this paper is related to *online mechanism design* [30,36,49,50]. In online mechanism design (online MD), one seeks an incentive mechanism for a dynamic environment in which agents arrive and depart and in which there is uncertainty about the future. In the analysis of our protocol, we slightly generalize the price-based characterization of Hajiaghayi et al. [29] to establish a dominant strategy equilibrium for buyers, creating a truthful online combinatorial auction from an uncoordinated sequence of single item auctions. The options-based scheme extends the earlier protocol in Hajiaghayi et al. [29] to combinatorial settings, although without providing any worst-case analysis. The mechanism presented here reduces to this earlier mechanism in an environment in which everyone is buying and selling a single unit of the same kind of good, although reinterpreted here within a decentralized architecture.

A dynamic VCG mechanism [18,49–51] could, in principle, be used in this environment because we consider only one-sided private information (and ignore strategic considerations on the sell side). But on the other hand, these mechanisms require *optimal* allocation policies (which may be intractable in domains of interest) together with a probabilistic model of agent arrivals and valuations, and provide within-period ex post incentive-compatibility rather than dominant-strategy incentive-compatibility [6,18]. One would also need to consider how to divide the revenue collected by the mechanism across multiple sellers. A dynamic extension of the *expected externality mechanism* [2,21] is available for dynamic two-sided markets [3], but would again require an optimal policy and provide only a weak guarantee about the relationship between a buyer's payment and her bid price. Bredin et al. [10] develop a framework for constructing dynamic double auctions with dominant-strategy incentive-compatibility on both sides of the market, but their methods are applicable only in the limited setting in which all agents trade only a single unit of the same good.

It is interesting to note that some of the strategic difficulties that buyers face in uncoordinated electronic auctions such as eBay are also faced by consumers acquiring items in the retail sector and it is interesting that retail stores have developed policies that can be interpreted as assisting customers in this regard. *Return policies* alleviate the exposure problem by allowing customers to return goods at the purchase price while *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 [19,38]. A concern that is discussed in the academic literature in regard to these practices in the retail sector is that they can be *anti-competitive*, with sellers using them as a commitment device for avoiding price competition [33]. We do not foresee this issue in the context of proxied, sequential auctions as proposed in this paper because the prices are not set by sellers but rather determined by competition on the buy side.

There has been some follow-up work to an earlier version of this paper. Mous et al. [42] present work on a variation of our scheme in which options are priced, and adopt a decision-theoretic analysis in exploring the effect on the sequential auction problem. This stands in contrast to our game-theoretic analysis, but has interest because it provides for less aggressive options accumulation by buyers.

2. Preliminaries: The sequential auction problem

In this section we introduce the formal model and define and characterize the sequential auction problem which motivates our work.

2.1. The model

In our domain, there are K different kinds of goods (often referred to as *items*), denoted with set G and $G_1, \ldots, G_K$ for the individual kinds of goods, a set B of buyers (perhaps unbounded), a set S of sellers (perhaps unbounded), and $T = \{0, 1, \ldots\}$ discrete time periods. Each buyer $i \in B$ has a utility function parameterized with *type* $\theta_i = (a_i, d_i, v_i) \in \Theta$, where Θ is the set of all possible types, defining her *arrival time* $a_i \in T$, *departure time* $d_i \geqslant a_i \in T$, and *valuation* $v_i(L) \geqslant 0$ for every possible bundle of goods $L \subseteq G$. There are multiple copies of each good G_k for sale but no buyer demands more than a single unit of each good and we write $L_k \in \{0, 1\}$ to denote whether or not bundle L contains a unit of good G_k.[8] We assume *free disposal* and *normalization*, with $v_i(L) \geqslant v_i(S)$ for $L \supseteq S$ and $v_i(\emptyset) = 0$.

The semantics of arrival and departure are such that buyer i has value $v_i(L) \geqslant 0$ for a bundle of goods $L \subseteq G$ that is allocated (potentially in multiple pieces) across periods $[a_i, \ldots, d_i]$ but has no value for goods allocated outside of this time interval. Buyers have *quasi-linear* utilities, so that the utility of buyer i for receiving bundle L and paying $p \in \mathbb{R}_{\geqslant 0}$, in some period no later than d_i, is $u_i(L, p) = v_i(L) - p$. For the sake of analysis it is convenient to assume the existence of a *maximal patience*, such that $d_i - a_i \leqslant \Delta_{\max}$, for some constant $\Delta_{\max}$.

We motivate the semantics of the arrival time by associating this with the period in which a buyer first realizes her demand for the good(s), or as the period at which a buyer first realizes that the market exists. The departure time models the period in which a buyer loses interest in acquiring the good(s) from this marketplace. For example, a buyer may lose interest in items whose value is realized at a specific past moment in time (e.g., Saturday night movie tickets), or because she simply wishes to take advantage of an outside opportunity to acquire the item (e.g., a buyer deciding to acquire an item at a posted price on a certain date). The model is restrictive in that it precludes a buyer having probabilistic beliefs about her value before arrival.[9] The model also requires that a buyer has constant value during the arrival-departure interval. This seems reasonable when this interval is small in relation to the time over which the good(s) will be used; e.g., an LCD monitor that a buyer plans on using for 3 years provides a buyer with roughly equivalent value if held for 1000 days or 998 days.

Each seller $j \in S$ brings a single unit of one kind of good, denoted $k_j \in G$ to the market and is assumed to have no intrinsic value for the good. Seller j has an *arrival time*, $a_j \in T$, which models the period in which she is first interested in listing the item, and a *departure time*, $d_j \geqslant a_j \in T$, which defines the latest period in which she is willing to consider having an auction for the item close. By listing a good for sale until d_j, a seller is indicating her willingness to receive payment by the end of the reported departure of the winning buyer in an auction closing at d_j, while preferring to exit the market

[8] The static version of this model was earlier referred to by Bartal et al. [5] as the "*k-duplicates*" combinatorial auction, with in their terminology k-units of each good and each buyer restricted to buying at most one unit of each good. We adopt k to reflect the number of kinds of goods rather than the number of units of each good.

[9] One reviewer suggested that another justification could be that the arrival time simply represents a constraint on the earliest period at which a bidder is able to participate.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

without her item being sold if there is no winner by that time. A departure time of d_j therefore guarantees that a successful seller will receive payment by $d_j + \Delta_{\max}$, where $\Delta_{\max}$ is the maximal patience of buyers. We will not be concerned with strategic behavior by sellers and focus our analysis on removing the sequential auction problem for buyers whatever the strategy of sellers.

The options-based framework provides a *direct-revelation* (online) mechanism in which each buyer i interacts with the market only once by declaring a *bid*, $b_i \in \Theta$, which is a (perhaps untruthful) claim about her type. Denoting the bids from all buyers as $b = (b_1, \ldots, b_N)$, a direct-revelation mechanism determines an *allocation* $x_i(b) \subseteq G$ and *payment*, $p_i(b) \geqslant 0$, to each buyer. The outcome also depends on the sell-side, i.e., the goods that are brought to market, but we leave this dependence silent in the notation. Since this is an *online* setting, where the bids are reported over time, the allocation and payment functions must be *online computable*; i.e., if a buyer is to be allocated an item for sale in period t then this must be known based on information available up to and including period t.

We adopt the standard assumption of *limited misreports* [29,30,36], with buyers unable to bid before their true arrival period. This is equivalent to requiring that reported type, $\hat{\theta}_i \neq \theta_i$, with $\hat{\theta}_i = (\hat{a}_i, \hat{d}_i, \hat{v}_i)$, satisfies $\hat{a}_i \geqslant a_i$. Given that the arrival time models the period in which a buyer first realizes her demand and enters the market, this assumption simply asserts that a buyer will not enter a market for goods for which she currently has no perceived value. Given this assumption, then strategy $\mathbf{b}_i^*(\theta_i) \in \Theta$ is a dominant-strategy equilibrium in mechanism (x, p), when

$$v_i\big(x_i\big(\mathbf{b}_i^*(\theta_i), b_{-i}\big)\big) - p_i\big(\mathbf{b}_i^*(\theta_i), b_{-i}\big) \geqslant v_i\big(x_i(b_i', b_{-i})\big) - p_i(b_i', b_{-i}), \quad \forall b_i' \in Y(\theta_i),\ \forall b_{-i} \in \Theta_{-i},\ \forall \theta_i \in \Theta, \tag{2.1}$$

for every buyer $i \in B$, where $b_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_N)$, and $\Theta_{-i} = \Pi_{\neq i}\Theta$, i.e., the joint type space of the other buyers. Here, we adopt $Y(\theta_i) \subseteq \Theta$ to denote the limited reports available to agent i, i.e., with $\hat{\theta}_i \in Y(\theta_i) \Rightarrow \hat{a}_i \geqslant a_i$.

Given this assumption of limited misreports, which we make for the rest of our analysis, a mechanism is said to be *strategyproof* (for the buy-side) when reporting the *true* type is a dominant strategy equilibrium for buyers. Truthful reporting by buyers is required to be a dominant strategy for all realizations of supply. Moreover, and even though we are not concerned with strategic sellers here, this also provides robustness for buyers against strategic seller behavior.

Our concern here is with online markets that are structured as a sequence of separate auctions. For this reason, it is helpful to adopt the terminology *locally strategyproof*, in describing an auction in which truthful bidding is a dominant strategy *if a buyer was restricted to bidding in that one auction* [48]. This would be true, for example, if the auction is for a unique item that is only available for auction once in a buyer's lifetime (e.g., a piece of artwork for which there is no substitute), or if a buyer is very impatient. In Section 5 we present empirical evidence to suggest that this is generally not a good assumption in markets such as eBay, and it is this observation that leads to the sequential auction problem.

One metric for the performance of our mechanism will be efficiency. The efficiency is a measure of the total value of buyers for an allocation of goods. For the empirical analysis, we will generally adopt one of the following two metrics for efficiency:

- average buyer value,
- total buyer value normalized by the value of an online benchmark.

Average buyer value provides a suitable efficiency metric when buyers are each interested in one unit of one good, i.e. the "single item demand" setting. No normalization is required when comparing the options-based scheme with the efficiency of eBay in this setting, because all goods are sold under both mechanisms and thus the total number of winning buyers is the same.

But when considering buyers with substitutes valuations, we normalize to an online benchmark because the total number of winning buyers varies between our simulation of eBay and the options-based scheme. The benchmark is explained in Section 5.2. For the case of complements valuations, it will be sufficient for our analysis (which is focused on the effect of correlations on values across items) to consider the total buyer value. We will adopt analogous metrics for seller revenue and buyer surplus; see Section 5.

For the purpose of theoretical analysis, we will also adopt a *worst-case* analysis for the single item demand setting, where we bound the worst-case efficiency of the options-based mechanism as a fraction of the value of the best offline solution, i.e., the best omniscient solution.

2.2. The sequential auction problem

The sequential auction problem describes the strategic problem that can face a buyer even though she faces a sequence of locally strategyproof auctions. Consider the following two motivating examples:

Example 1. Alice values acquiring one ton of sand before Wednesday for \$1000. Bob will hold a Vickrey auction for one ton of sand on Monday, and another such auction on Tuesday. Alice has no dominant bidding strategy because she cannot predict whether the price of the Tuesday auction will be greater or less than the price of the Monday auction, and she needs to know this price when deciding on an optimal bidding strategy on Monday.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

Example 2. Alice values one ton of sand with one ton of stone at $2000. Bob holds a Vickrey auction for one ton of sand on Monday. Charlie holds 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 in order to know how much to bid for sand on Monday. If Alice bids too high on Monday, she may be left with one ton of sand but no ability to buy the one ton of stone required to complete her construction project. If Alice bids too low on Monday, she might forfeit the opportunity to buy both the sand and stone, for example if the price of stone on Tuesday is low.

Definition 1 *(sequential auction problem).* The sequential auction problem exists when a buyer has no dominant bidding strategy in a sequence of auctions, despite each auction being locally strategyproof.

There are a variety of ways in which an absence of a dominant strategy may arise. First, consider a buyer who values one ton of sand for $1000, one ton of stone for $2000 and both for $2000 (i.e., substitutes valuations). Suppose the buyer faces one auction for sand followed by one for stone and that the possible prices for either goods are unconstrained. She faces the sequential auction problem because there are future prices for which it is optimal to purchase sand and some for which it is optimal to purchase stone. In general, buyers with substitutes valuations face the sequential auction problem. As a second example, consider a buyer who values one ton of sand for $1000 and one ton of sand and one ton of stone together for $1500. Suppose that the buyer faces an auction for sand and another for stone and that either bundle $\{\text{sand}\}$ or $\{\text{sand}, \text{stone}\}$ may be the utility maximizing bundle given possible prices. She faces the sequential auction problem because the good that first goes to auction has uncertain marginal value; this value is either $1000 or $1500 if the good is sand and $0 or $500 if the good is stone. In general, buyers who are interested in goods whose marginal values are dependent on the acquisition of other goods face the sequential auction problem.

As another example, consider a buyer who values one ton of sand for $1000 and faces two auctions for sand, either of which may have the lowest price. She faces the sequential auction problem because even though the value of the sand is $1000 and constant in both auctions, either auction may have the best price. In general, buyers who are interested in fewer instances of a good than there are auctions selling that good face the sequential auction problem. For a final example, consider a buyer who values sand for $1000, stone for $500 and sand and stone together for $1500 (i.e., a linear valuation) and suppose just two auctions, one for sand and then one for stone. Suppose one other buyer competes in both auctions. If the competitor may bid $1,000,000 for stone if the buyer bids more than $300 for sand, and $10 otherwise, the buyer is better off bidding below $300 for sand, even if this involves losing that auction, because she will then receive the stone for a payoff of $1490. In general, buyers whose competitors condition future bids on the buyers' past bids face the sequential auction problem.

3. An options-based scheme

In what follows we focus exclusively on domains in which the underlying auctions are Vickrey auctions for individual items (i.e., second price, sealed-bid auctions). Vickrey auctions are selected not only because of convenience, but also because they nicely model eBay auctions.

The solution that we propose, in resolving the sequential auction problem in this context, consists of two primary components: the use of *real options* to allow buyers to secure the lowest possible prices and the use of *mandatory proxy agents* to prevent the abuse of these options through costless hoarding. A real option is a right to acquire a real good at a certain price, called the *exercise price*; see Dixit and Pindyck [22]. For instance, Alice may obtain from Bob the right to buy sand from him at an exercise price of $1000. An option gives the *right* to purchase a good at an exercise price but does not imply an *obligation*. We will see that this flexibility makes options useful in addressing the sequential auction problem. Proxy agents acting on behalf of buyers can put together a collection of options, and then decide which options (perhaps none) to exercise.

While the buyer of an option has the *right* but not the obligation to purchase the good, the seller must honor the contract if the option is exercised. For this reason, options are typically sold at a premium called the *option price*. Several factors are often considered when a seller tries to determine how to price an option with a particular exercise price, including the relationship between the exercise price and the perceived value of the good available in the option, the volatility of value the good may experience over time, and the length of time over which a buyer can decide to exercise the option. Real options are often difficult to price as the metrics for determining a price are difficult to quantify. However, among *traded* options (i.e., options for traded securities such as stock), much progress has been made in determining the prices of options, with one of the most celebrated being the formula of Black and Scholes [7].

The problem with options with a non-zero option price in our setting is that they cannot support a simple, dominant bidding strategy because a buyer would need to compute the expected value of an option to justify the cost. But this expected value requires a probabilistic model of the future, which in turn requires the buyer to model the bidding strategies and values of other buyers. This is the very pattern of reasoning that we are trying to avoid in designing the options-based marketplace! For this reason, we adopt *costless options*, which always have an option price of zero. The exercise price is set competitively in the marketplace.

The traditional issue with costless options is that buyers are always weakly better off with a costless option than without one, whatever the exercise price. A buyer need exercise only those option(s) that result in a gain of surplus and bears no cost by not exercising an option. But having buyers that pursue options that they have no intention of exercising would cause market efficiency to unravel. To prevent this kind of hoarding of options, we adopt mandatory proxy agents. These proxy agents provide buyers with an *obligation* to exercise those options that maximize their reported utility – as defined by reported valuation $\hat{v}_i$ – given the exercise price. The proxy agents also act to restrict buyers to acquiring only options that they might credibly choose to exercise.[10] Only proxy agents place bids in the underlying auctions.

In Section 3.1 we give details about the proxied, options-based solution. Section 3.2 provides a number of detailed examples and Section 3.3 provides a complexity analysis. We delay until future sections a formal proof of strategyproofness and any analysis of the efficiency and revenue properties.

3.1. The bidding proxy and price-matching rules

In our framework we first modify each individual Vickrey auction to sell a real option for the underlying good to the highest bidder with an *initial* exercise price equal to the second-highest bid price received.[11] Each option is costless, and is set to expire at the end of the winning proxy's patience. Proxy agents bid in these auctions. In opting into the options-based protocol, a seller gives the winning proxy agent the right to reduce the exercise price on the option issued by the seller given evidence that a lower price would have been available had the proxy waited and bid instead in some future auction. This is what we mean by "price matching."

Buyers must compete in the market by submitting a bid to their proxy agent. For buyer $i \in B$, this bid occurs in some reported arrival time, $\hat{a}_i \geqslant a_i$, and is a claim about her valuation $\hat{v}_i$ (perhaps untruthful) for different bundles of goods and also about her departure time $\hat{d}_i \geqslant \hat{a}_i$. All transactions are intermediated by proxy agents. In what follows, we describe the three steps that are followed by a proxy agent: (a) acquiring options, (b) setting the exercise price on options via seller-sanctioned price matching, and (c) exercising options. This completely defines the options-based mechanism.

Step one: Acquiring options. When an option for a good in which a buyer is interested is available in an auction and the proxy does not hold an option, the proxy submits a bid equal to the buyer's (reported) *maximum marginal value* for the item. A proxy does not bid for an item on which it already holds an option. The maximum marginal value for an item G_k, given reported valuation $\hat{v}_i$, is defined as:

$$bid_i(k) = \max_{L \subseteq G}\left[\hat{v}_i\left(L \cup \{G_k\}\right) - \hat{v}_i(L)\right]. \tag{3.1}$$

By bidding this value, a proxy will compete for any option that *could possibly* be of benefit to the buyer and choose not to bid only on those options that *could never* be of value to the buyer. Note that this is a static determination, made entirely in terms of the valuation of a buyer and without considering the prices for which the proxy already holds options. In bidding up to the maximal possible value of some item, the proxy is considering the case that all goods in bundle L that go together with this good will be available and for a price of \$0.[12]

Step two: Setting the exercise price. Rather than acquire more than one option on the same kind of good, proxy agents are authorized by sellers to adjust the exercise price of an option that they win downwards (from the initially set price of the second-highest bid). Such an adjustment is made whenever a proxy discovers that it could have achieved a better exercise price for an option on the same kind of good by waiting to bid in a later auction. A proxy is able to identify such a missed opportunity by storing locally, for each good on which it holds an option, *the identity of the active bidder* (*if any*) *that would have already won an option had the proxy itself not won an option*. Initially, when a proxy wins an option it stores in its local memory the identity (which may be a pseudonym) of the proxy agent that it "bumped" from winning, if such a proxy exists (i.e., the second-highest bidder).

To see how price matching works, fix some good G_k on which Proxy A holds an option. Proxy A now monitors each future auction for the same kind of good and determines what the buyer population would be had Proxy A delayed its own bid until that auction. To make this determination, Proxy A requests from the auction the identities of the buyer proxies (if any) and their bids.[13] Proxy A identifies the highest bid across those proxies whose identity is not stored in A's local memory. This is the bid against which Proxy A would be competing had A delayed its entry until this auction. If the bid price of Proxy B is lower than Proxy A's current exercise price, Proxy A *price matches* down to Proxy B's bid price, since this is exactly the price that Proxy A could have achieved by delaying its bid until this later auction. If there is no such bid from another proxy then the price is matched down to zero (or a reserve price, if any).

[10] Unlike in the financial markets, our solution does not permit option holders to sell their options to others. Therefore, proxies have no incentives to hoard costless options speculating on future option reselling opportunities.

[11] The system can also set a reserve price for each kind of good, provided that the same reserve is adopted for all auctions selling the same good. Without such an invariant reserve price, price matching would not be possible as a seller might be required to match a price below their personal reserve price. Exploring how limiting this may be in scenarios where sellers actually have different reservation values is an area of future work.

[12] If the proxy has knowledge that some items will not be for sale, or can lower-bound the possible price on other items (e.g., because of a market-wide reservation price adopted by sellers – see Footnote 11), then the marginal value can be modified downwards to preclude such bundles or adjust downward by lower bounds on the price of items. Care must be taken, though. We return to this issue in Section 6.

[13] In a marketplace such as eBay, this information could be provided (again in pseudonymous form) by the market infrastructure. Moreover, the only information that is minimally required is the highest bid price across all buyer proxies except one stated by the proxy, and the identity of the highest proxy if it was not the winner of the auction.

Table 1
A 3 buyer example in which each buyer wants a single item and one auction occurs on each of Monday and Tuesday. "X_Y" indicates an option with exercise price X and bookkeeping that a proxy has prevented Y from currently possessing an option. "$\rightarrow$" indicates the updating of exercise price, together with bookkeeping for an option already held.

Buyer	Reported type	Monday	Tuesday
Molly	(Monday, Tuesday, \$8)	6_{Nancy}	$6_{\text{Nancy}} \rightarrow 4_{\text{Polly}}$
Nancy	(Monday, Tuesday, \$6)	–	4_{Polly}
Polly	(Monday, Tuesday, \$4)	–	–

Formally, the exercise price $p^t(k)$ in period t on good G_k on which the proxy holds an option is set to $\min_{\hat{a}_i \leqslant \tau \leqslant t}\{p^{\tau}_{b_{-i}}(k)\}$ where $p^{\tau}_{b_{-i}}(k)$ is the maximal bid price across other proxies in period τ on good G_k, but excluding any bid from a proxy whose identity is stored in i's local memory in period τ. If there are no such bids from other proxies, then $p^{\tau}_{b_{-i}}(k)$ is set to zero (or a reserve price, if any).

After price matching, one of two adjustments is made by Proxy A for bookkeeping purposes. First, if Proxy B was the winner of this latest auction, Proxy A's local memory as it relates to this good can be cleared. This is because the proxy earlier bumped by Proxy A must no longer be bidding (otherwise it would have won this latest auction before Proxy B), and thus Proxy A's earlier win no longer affects the set of active bidders for this good going forward. On the other hand, if Proxy B is the second-highest bidder in this latest auction (having been outbid by the proxy stored in Proxy A's local memory), Proxy B will now be stored in Proxy A's local memory, as Proxy B would have won the good without the presence of Proxy A in the market (because the proxy stored in Proxy A's local memory would not have competed against Proxy B in this latest auction had Proxy A never been around to bump it from the auction that Proxy A won).

Step three: Exercising options. At the reported departure time, $\hat{d}_i$, the proxy for buyer i chooses which options (if any) to exercise. The option(s) that are exercised, L^*_i, are those that maximize the reported utility of the buyer given the final exercise price on each good:

$$L^*_i \in \operatorname*{argmax}_{L \subseteq \mathcal{O}} \left[\hat{v}_i\big(\gamma(L)\big) - p(L)\right], \tag{3.2}$$

where $\mathcal{O}$ is the set of all options held by the proxy, $\gamma(L) \subseteq G$ are the goods that correspond to some subset $L \subseteq \mathcal{O}$ of these options, and $p(L) = \sum_{k \in L} p(k)$ is the total exercise price for this set of options where $p(k)$ is the exercise price on the option corresponding to good G_k (as determined via price-matching). All other options are returned. No options are exercised when there is no bundle of options with (weakly) positive utility.

Remark: Re-posting seller options. An auction for the good brought to market by a seller will first occur at the arrival period of the seller. If at some point later the buyer that wins this auction returns the option unexercised and the time period is before the seller's departure then it would be ideal to be able to initiate another auction for an option on the seller's good. However, the system prevents a seller from re-auctioning an option until the maximal patience, $\Delta_{\max}$, after the option was first allocated. Recall that $\Delta_{\max}$ defines the maximal patience (departure-arrival) over all possible buyers in the market.[14] This maintains a truthful mechanism by preventing a buyer from acquiring an option with a view to holding it, returning it unexercised, and later re-entering the market when the option is again auctioned and achieving a lower price.[15]

In the absence of *strong identities*, where the market definitively knows the identity of each market participant, the market prevents buyers from affecting future supply in a useful way by waiting a sufficiently long amount of time before reauctioning a returned option.[16] However, in the presence of strong identities, the market can explicitly prevent a proxy's buyer who has returned an option from bidding on that option for $\Delta_{\max}$ into the future. Consequently, a seller can reauction an option as soon as it is returned, increasing the likelihood of selling the item before her departure.

3.2. Examples of proxy behavior

We first provide an example to illustrate the price-matching logic that is followed by proxies (also illustrated as Table 1):

Example 3. Consider three buyers, all of whom enter the market on Monday and depart the market after Tuesday. Molly values an item for \$8, Nancy for \$6 and Polly for \$4. On Monday, an auction occurs where all three proxies bid, with Molly's

[14] In practice one might choose to make a tradeoff here, in which that gain from allowing earlier reposting of an item can be traded for the cost of losing strategyproofness for the most patient buyers.

[15] Here is an example where Alice has a useful manipulation of this kind: Alice values an Apple for \$5 from Monday to Wednesday. Bob values the Apple for \$8 but only on Monday. Consider a seller with an Apple and high patience. If Alice is truthful then Bob wins on Monday for \$5 and exercises the option. But if Alice claims to value an Apple together with a Banana at \$10 only from Monday to Tuesday then Alice wins the Apple option for \$8, but returns it at the end of Tuesday. On Wednesday, the seller re-posts and an Apple auction occurs with Alice returning and bidding \$5 for an Apple alone and winning the Apple for \$0.

[16] One common technique that is used at present to achieve strong, or almost strong, identities in electronic markets is to require a unique cell phone number or credit card number of every registration.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

proxy winning with the highest bid of \$8 and receiving an option for \$6. Molly's proxy adds Nancy to its local memory. On Tuesday, another auction occurs where only Nancy's and Polly's proxies bid, 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 (because Nancy is already in memory) and replace Nancy with Polly in its local memory for bookkeeping purposes, as Polly would be holding an option had Molly delayed her bid past this round.

To illustrate how the options-based scheme handles the exposure problem, consider the following example where Alice desires a bundle of two goods:

Example 4. Alice values one ton of sand and one ton of stone together for \$3000 (but has no value for either by itself). Bob values one ton of sand for \$800. Charlie values one ton of stone for \$2000. All buyers have a patience of 2 days. On day one, a stone auction is held, where Alice's proxy bids \$3000 and Charlie's bids \$2000. Alice's proxy wins an option to purchase stone for \$2000. On day two, a sand auction is held, where Alice's proxy bids \$3000 and Bob's bids \$800. Alice's proxy wins an option to purchase sand for \$800. At the end of the second day, Alice's proxy holds an option to buy stone for \$2000 and sand for \$800 and exercises both options spending a total of \$2800.

As an illustration of how the options-based scheme handles substitutes values consider the following example:

Example 5. Alice values either one ton of coarse sand for \$1000, or one ton of fine sand for \$800 (but only \$1000 for both). Bob values coarse sand for \$800. Charlie values fine sand for \$900. On day one, a coarse sand auction is held where Alice's proxy bids \$1000 and Bob's proxy bids \$800, resulting in Alice's proxy winning an option for the coarse sand with an exercise price of \$800. On day two, a fine sand auction is held where Alice's proxy bids \$800 and Charlie's proxy bids \$900, resulting in Charlie's proxy winning an option for the fine sand with an exercise price of \$800. At the end of day two, Alice's proxy exercises its coarse sand option and Charlie's proxy exercises its fine sand option.

3.3. Complexity analysis

In providing a complexity analysis for the problem facing proxy agents, we consider the particular case of a valuation function that is described in the *exclusive-or* (XOR) bidding language [45]. An XOR valuation of size M defines a set of M bundle-value pairs (or *atomic terms*), $\{(L^1, v_i(L^1)), \ldots, (L^M, v_i(L^M))\}$, and defines valuation

$$v_i(S) = \max_{L^m \subseteq S, m \in \{1,\ldots,M\}} \left[v_i\left(L^m\right)\right]$$

for any bundle S, where L^m is one of the atomic terms. This is equivalent to saying that buyer i is interested in buying *at most one* bundle. We have the following two immediate results:

Theorem 1. *Given an XOR valuation of size M, there is an $O(K_i M^2)$ algorithm for computing the maximum marginal value on each interesting good for buyer $i \in B$, where $K_i = |\bigcup_{m\in\{1,\ldots,M\}} L^m|$ is the number of different items in which the buyer is interested.*

Proof. For each item, recall Eq. (3.1), which defines the maximum marginal value of an item. For each bundle L^m in the M-term valuation, and any item G_k, $v_i(L^m \cup \{G_k\})$ can be identified by considering each of the M terms in sequence. Therefore, the number of terms explored to determine the maximum marginal value for any item G_k is $O(M^2)$, and the total number of bundle comparisons to be performed to calculate the maximum marginal value on every item is $O(K_i M^2)$. □

Theorem 2. *The total memory required by a proxy to implement price matching is $O(K_i(\log(V) + \log(B)))$, where*

$$K_i = \left| \bigcup_{m\in\{1,\ldots,M\}} L^m \right|$$

is the number of different items in which the buyer is interested, V is the maximum value a buyer possesses for a bundle, and B is the maximum number of buyers within the system at any given time. The total work performed by a proxy in updating the state in each auction is $O(1)$.

Proof. The proxy stores one maximum marginal value for each item of interest, i.e., $O(\log(V))$ information for each of $O(K_i)$ items. The proxy also stores at most one buyer's identity for each item of interest, i.e., $O(\log(B))$ information for each of $O(K_i)$ items.[17] The proxy also stores one exercise price for each item of interest, i.e., $O(\log(V))$ information for

[17] As Lemma 2 implies, while more than B buyers may be present over the entire time in which a buyer participates in the marketplace, only the B buyers in the market at any given time actually matter to a proxy, and so only $\log(B)$ information is required to store the identities, assuming proper recycling of information when buyers leave the market.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

each of $O(K_i)$ items. For each auction, the proxy either submits a precomputed bid or price matches, both of which take constant work. □

The computation required to determine the most profitable bundle at the time of departure is $O(MK_i)$ arithmetic operations, as for each bundle one can determine the potential profit on that bundle by iterating over the options held on items in that bundle, of which there are at most K_i.[18]

4. Theoretical analysis of the options-based scheme

In this section, we establish that the options-based scheme supports truthful bidding as a dominant bidding strategy for buyers. We also develop a worst-case, competitive analysis for allocative efficiency in the practical case in which all buyers and sellers want to trade one unit of an identical good. The competitive analysis generalizes earlier analysis due to Hajiaghayi et al. [29] to include a bound on the maximum ratio of minimum to maximum values in the buyer population. This is a useful modification because values will typically fall into natural bounds in practical settings.[19]

4.1. Strategic analysis: Establishing truthfulness

The proxies transform the market into a direct revelation mechanism, where each buyer i makes a claim about her type to her proxy agent in some period $\hat{a}_i \geqslant a_i$. The dominant bidding strategy for each buyer is to report her *true* valuation and *true* departure time to her proxy agent immediately upon arrival to the system, that is, to bid her true type $\theta_i = (a_i, d_i, v_i)$. We assume that neither buyers nor sellers receive any information about prices, or previous bids of other agents, before submitting a bid to their proxy agents.[20] A brief discussion about the effect on equilibrium properties of allowing buyers and sellers to see information about the market before bidding or deciding when to list goods is given at the end of this section.

In establishing strategyproofness, we provide a slight generalization of an existing characterization of strategyproof online auctions [29], to allow for combinatorial online auctions. For this, define a *value-independent price function*, $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L) \geqslant 0$, on all $L \subseteq G$, which can depend on the bids of other agents b_{-i} and the reported arrival $\hat{a}_i$ and departure $\hat{d}_i$ of buyer i. The price function is *value-independent* because it does not depend on the reported valuation $\hat{v}_i$ of the buyer. The price function will depend on the realization of supply, but this dependence is suppressed to keep the notation simple.

Definition 2 *(monotonic prices).* A value-independent price function is monotonic in arrival and departure if $ps_{a'_i,d'_i,b_{-i}}(L) \geqslant ps_{a_i,d_i,b_{-i}}(L)$, for all buyers i, all $a'_i \geqslant a_i$, all $d'_i \leqslant d_i$, all bids b_{-i} by other buyers, all realizations of supply, all a_i, d_i and all $L \subseteq G$.

Monotonic prices increase with a tighter arrival-departure interval, fixing the bids submitted by other buyers and the realization of supply within the market.

Lemma 1. *An online combinatorial auction is strategyproof if there exists a monotonic, value-independent price function, $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L)$, and given that for every buyer i, all reports $\hat{\theta}_i = (\hat{a}_i, \hat{d}_i, \hat{v}_i)$, all bids b_{-i} from other buyers, and all realizations of supply, the buyer is allocated a bundle $L^*_i \in \operatorname{argmax}_{L\subseteq G}[v_i(L) - ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L)]$ in period $\hat{d}_i$ and makes payment $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L^*_i)$.*

Proof. Fix some a'_i and d'_i. A buyer should report true valuation function, $\hat{v}_i = v_i$, because the prices she faces are independent of her report and by being truthful bundle L^*_i maximizes her true utility. This in place, fix $\hat{v}_i = v_i$. Now, it is never useful to bid $\hat{d}_i > d_i$ because the buyer will not receive her allocation until her true departure (and have zero value). By limited misreports, the buyer cannot report $\hat{a}_i < a_i$. Reporting $\hat{a}_i > a_i$ or $\hat{d}_i < d_i$ (weakly) increases the price on every bundle L by monotonicity. □

---

[18] We can also consider the impact of alternate bidding languages, such as additive-OR valuations or its generalizations (e.g., OR* [45]), which can be exponentially more concise than the XOR language, and might therefore be desirable in some markets. First notice that the time and space complexity of price matching is unchanged. On the other hand, it is well known that the *valuation problem*, i.e., finding $v_i(S)$ for some bundle S, is NP-hard for OR and OR* because one must solve a maximal weighted set packing problem, with the atomics providing the elements of the set. Thus, the problem of finding the most profitable bundle upon departure is also NP-hard, with the value of each atomic adjusted according to the option prices. The maximum marginal value problem is also NP-hard, via a reduction from the valuation problem to the marginal value problem, in which a new item k' is introduced to all atomics, such that $v_i(S) = 0$ while $v_i(S \cup G_{k'})$ provides the value on bundle S for the original bid.

[19] The worst-case analysis that we present is limited to this "single-item" environment in which each agent is buying or selling one copy of an identical good. As a purely computational problem, the so-called k-duplicates combinatorial auction winner determination problem is NP-hard to approximate to within a worst-case factor of $O(K^{\frac{1}{n+1}-\epsilon})$, with K types of goods, n copies of each good, and where each agent interested in at most one good [5]. In the special case of $n = 1$ (i.e., with one copy of each good in the market) this reduces to the (tight) lower-bound of $O(K^{\frac{1}{2}-\epsilon})$ [31,58]. The lower-bound in the k-duplicates setting with a large number of each item is more forgiving, leaving more opportunity for developing methods with good worst-case properties.

[20] This is in contrast to proxy agents, who must receive information about bids from other proxies in order to perform price matching.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

Observe that it is not necessary for the prices to be monotonic with respect to subset inclusion of bundle L because the assumption made in the lemma is that the bundle allocated is one that maximizes payoff across all bundles; if a larger bundle has a smaller price then it would be selected.

Price-matching and bookkeeping are the methods by which a monotonic, value-independent price function is constructed in the options-based auction. This is established in the following lemmas. The first lemma is useful in establishing the correctness of the bookkeeping algorithm:

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

Proof. Fix some item and the proxy for buyer i. The proof is by induction on the sequence of auctions for the item while buyer i's proxy is present in the market. The correctness of the information in the proxy's local memory is easy to establish in the base case before the proxy has participated in any auction. Now consider the first auction for this item in which the proxy wins and suppose it prevents another proxy from winning an option on the good (which is the interesting case). Consider now two cases: (a) the bumped proxy will leave the system having never won an option on the item, or (b) the bumped proxy will win an auction on this item in the future. In case (a), while this bumped proxy is still present then proxy i's presence prevented exactly that one proxy but no other proxies from winning an option. This is because the presence of the bumped proxy in the market does not preclude any other proxy from winning (because the bumped proxy is losing anyway). The identity of this bumped proxy remains in proxy i's local memory because no price matching will have occurred on this item because each winning proxy in each subsequent auction must have submitted a bid higher than the bumped proxy's bid (else the bumped proxy would have won) and therefore higher than proxy i's exercise price, which is initialized to the bid price of the bumped proxy. Eventually the bumped proxy leaves and proxy i no longer has any effect on the bid dynamics for this item. At some point some other proxy may win an option for this item and the proxy's local memory is cleared. In case (b), while the bumped proxy is not yet winning then this is as in case (a). In period t in which the bumped proxy wins then proxy j, with the highest other bid in that auction (if any), would have won without i's presence. Proxy i necessarily price matches in this case – because the exercise price it could have achieved is the bid price of proxy j and less than that of the winning (bumped) proxy and thus its current exercise price – and then updates its local memory to contain the identity of proxy j. Proxy j is the new proxy that does not hold an option because of buyer i's presence in the market. Case (a) or (b) now holds again for this new bumped proxy, proxy j, and the proof continues until proxy i finally departs the market. □

In establishing strategyproofness, it is sufficient to consider the special case of a price function $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L)$ that is linear in the items $G_k \in L$, so that $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L) = \sum_{G_k \in L} ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(k)$. This will be the case because the price on a bundle in the options-based scheme is constructed by adding the options price on the constituent goods of the bundle. The price function $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(k)$ on items is itself monotonic with respect arrival and departure, with prices associated with the options prices, because of price-matching on individual options. This implies that the price function $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L)$ is monotonic. The following easy lemma is stated without proof:

Lemma 3. *The bundle of goods that maximizes a buyer's reported valuation given a linear, agent-independent and monotonic price function, defined in terms of the sum of prices on individual goods, will never contain an item that is priced above her maximum marginal value.*

Theorem 3. *Truthful bidding is a dominant-strategy equilibrium for buyers in the options-based, proxied market and for all sell-side strategies.*

Proof. Fix buyer i. First, define an agent-independent price, $p^t_{b_{-i}}(k)$, on item G_k in period t as the highest bid by the proxies $\neq i$ not holding an option on item G_k at time t (∞ if there is no supply at t), and not including any proxy that would have already won an option had i never entered the system (i.e., whose identity is stored in the i's proxy's local memory for item G_k by Lemma 2). Conditioned on $t \geqslant \hat{a}_i$, this price is well-defined and independent of any report $\hat{\theta}_i$ buyer i makes to her proxy because the set explicitly excludes the one proxy (see Lemma 2) that i prevents from holding an option by its presence, and comprises exactly those bids that would be made without i present. (For this, we exploit the fact that the bid values are equal to maximal marginal values and independent of earlier bids by buyer i's proxy.) Furthermore, i cannot influence the supply of item G_k because any options returned by other buyers due to a price set by i's proxy's bid will be re-auctioned (if at all) after i has departed the system. Moreover, the market is opaque and thus seller's decisions about when to list an item are independent of the bid of buyer i. Now define $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(k) = \min_{\hat{a}_i \leqslant \tau \leqslant \hat{d}_i} p^\tau_{b_{-i}}(k)$ (possibly ∞), which is well defined upon $\hat{d}_i$ and is monotonic in arrival and departure because it is defined as the minimal value over its domain. Conditioned on holding an option on G_k upon departure, this is exactly the exercise price obtained by buyer i's proxy. Now define $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L) = \sum_{k:\ G_k \in L} ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(k)$, which is monotonic in $\hat{a}_i$ and $\hat{d}_i$ because $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(k)$ is monotonic and remains value-independent. (Note that this price is ∞ when there was never any supply of item G_k.) Given the options held by a proxy at $\hat{d}_i$, which may be a subset of those items G_k with prices $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(k) < \infty$, the proxy exercises options to maximize utility based on reported valuation, $\hat{v}_i$. We show that the proxy would not want to select any options on items that are not available because the prices on missing options would be too high. For this, consider such a bundle L', that is interesting based on $\hat{v}_i$ but has one or more missing options. One possibility is that there is an item $G_k \in L'$ for which an option was never for sale, in which case $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(k) = \infty$ and $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L') = \infty$ and L' is not utility maximizing. On the other hand, in the case that every item in L' was available for sale in interval $[a_i, \ldots, d_i]$, we know that $p^t_{b_{-i}}(k)$ was at least the maximal marginal value in every such period (else the buyer's proxy would have won) and the bundle cannot maximize utility by Lemma 3. □

Note that a proxy acquires an option on every item that could possibly be in the buyer's utility-maximizing bundle at the final prices. For any option that a proxy agent does not explicitly hold upon departure, either the item was never for sale, or the competition was such that the price was always so high that the buyer would not want to exercise the option even if the clearing price on the other options was zero.

Remark 1. The options-based scheme satisfies *voluntary participation* (or "individual rationality") for both buyers and sellers, meaning that every participant has non-negative utility in equilibrium. For buyers, this follows because the proxy exercises a utility maximizing set of options and will exercise no options if all bundles have negative utility. For sellers without any intrinsic value for the good, voluntary participation follows because the prices on options remain non-negative.

Remark 2. The options-based scheme is robust against buyers that can participate under multiple identities or through re-entry with the same identity. This is because the price function $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(L)$ is *linear* on every bundle of goods $L \subseteq G$, being defined as the sum of the prices, $ps_{\hat{a}_i,\hat{d}_i,b_{-i}}(k)$ introduced in the proof. So long as the prices that a buyer faces remain unchanged when she participates multiple times, she cannot gain by winning multiple bundles because of linearity of prices and because all alternate bids must necessarily have (weakly) tighter arrival-departure intervals and therefore higher prices. In fact, the prices may increase – but cannot decrease – when a buyer participates multiple times because prices in the underlying auctions are weakly increasing in more participants (a property of the Vickrey auction), and note that the supply available to a buyer is unaffected by its strategy because of the delay that is required of a seller before reposting an item for auction.

Remark 3. In a practical market such as eBay, it is often desirable to allow buyers to see information about the market before bidding; e.g., about recent prices and recent bids. The effect of providing information about the market is a weakening of the equilibrium from a dominant strategy equilibrium to an *ex post* Nash equilibrium, in which truthful bidding would remain a best-response for every buyer (and for all possible types of other buyers and all possible market dynamics), as long as other buyers are rational and bid truthfully. Consider the strategic problem facing buyer i. The earlier analysis holds, but only as long as the bids of other buyers and therefore the prices faced by this buyer, continue to be independent of her own bid. This is true in the *ex post* Nash equilibrium because the other buyers follow the truthful-revealing, equilibrium strategy. But this agent-independence of prices is no longer true whatever the strategy of the other buyers, because another buyer might follow a "crazy" strategy and bid extremely high if she sees a bid of a particular value from buyer i, and thus buyer i should deviate from bidding truthfully to prevent the triggering of this action by the other buyer.[21] A similar consideration applies on the sell-side, with buy-side incentive compatibility retained as long as seller's strategies are unaffected by market information that is affected by active buyers, and including when sellers are non-strategic.

4.2. Efficiency analysis

In this section, we provide a worst-case, competitive analysis in the practical case in which all participants are buying or selling one unit of the same kind of good. An online mechanism is said to be *k-competitive with respect to efficiency* if it is guaranteed to achieve an allocation with value at least $1/k$ of that achieved in an optimal offline, or *omniscient* allocation. The omniscient allocation maximizes total value given perfect hindsight about the arrival and departure and values of market participants. For example, a mechanism that is 2-competitive for efficiency will achieve a total value that is at least half of the total value of the optimal, offline allocation, for all possible agent populations.

In this special case, in which all buyers and sellers trade at most one unit of the same good, our problem is the same as that considered by Hajiaghayi et al. [29], and the options-based scheme presented here is equivalent to their mecha-

21 The same kind of phenomenon occurs in indirect mechanisms such as ascending-price VCG mechanisms, where straightforward bidding is an *ex post* Nash equilibrium but not a dominant strategy equilibrium [41].

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

nism.[22] These authors show that their mechanism, and thus also the options-based scheme, is 2-competitive for efficiency in this setting and also prove a tight lower-bound to show that no truthful, online mechanism can provide better than 2-competitiveness and the options-based protocol has the best possible, worst-case efficiency.

We provide a slight generalization of the analysis and parameterize the competitive analysis with a lower bound, $0 \leqslant \alpha \leqslant 1$, on the ratio of the minimum to maximum value for the item in the buyer population; i.e., $w_i/w_j \geqslant \alpha$ for all buyers $i, j \in B$, with values w_i, w_j on the item for buyers i and j respectively. For $\alpha = 0$ the relative values are unbounded and w_i can be arbitrarily smaller than w_j. For $\alpha = 1$ the values must be identical across all buyers because if $w_i \neq w_j$ then either $w_i/w_j < 1$ or $w_j/w_i < 1$. This parameter provides for a tighter worst-case analysis in settings in which such a bound on the relative values of buyers is available. The result relies on an additional assumption when $\alpha > 0$, which is that there is at least one buyer without an option present in every auction. This is enough to guarantee that every item is sold, which is required for the analysis.[23]

Theorem 4. *When every buyer and seller is interested in trading one unit of the same item, given a lower-bound $0 \leqslant \alpha \leqslant 1$ on the ratio of minimum to maximum values in the buyer population, and given at least one buyer in every auction not holding an option if $\alpha > 0$, then the options-based scheme is $\frac{2}{1+\alpha}$-competitive for efficiency.*

Proof. Let *OFF* and *ON* denote the winners in the offline and online solutions respectively. We seek a lower bound $V_{\mathrm{ON}}/V_{\mathrm{OFF}} = \sum_{i \in ON} v_i / \sum_{i \in OFF} v_i \geqslant \frac{1+\alpha}{2}$ for all possible inputs. For any input, we can place an upper bound on V_{OFF} in terms of V_{ON} through a charging argument, as follows. Following Hajiaghayi et al. [29], consider some buyer $i \in OFF$. We "charge" her value to a buyer in *ON*. If $i \in ON$ then we charge the value to herself. Otherwise, let *auc* be the auction that i wins offline. Since i never wins in the options-based market, she was present in the market when *auc* closed, and so the options-based scheme must have picked a winner $j \in ON$ whose value is (weakly) greater than the value of i. We charge the value of i to j. It is not hard to see that this charging scheme charges each agent j in the options-based market at most twice, each time for a value less than the value of j. Let $ONCE \subseteq ON$ and $TWICE \subseteq ON$ denote the online winners that are charged once and twice respectively. We have that $V_{\mathrm{OFF}} \leqslant \sum_{j \in ONCE} v_j + 2\sum_{j \in TWICE} v_j$. For $\alpha = 0$ this gives

$$\frac{V_{\mathrm{ON}}}{V_{\mathrm{OFF}}} \geqslant \frac{(\sum_{j \in ONCE} v_j + \sum_{j \in TWICE} v_j)}{(\sum_{j \in ONCE} v_j + 2\sum_{j \in TWICE} v_j)} \geqslant 1/2,$$

with the worst-case occurring for $ONCE = \emptyset$. Consider now $\alpha > 0$, and let $B' \subseteq OFF$ denote the subset of *OFF* that are matched to the *TWICE* set. Let $K = |B'|$ and note that K must be even so that $K/2$ is an integer. Now, for $\alpha > 0$ we know that $n = |OFF| = |ON|$, because all items are sold in the online solution by assumption. Then, we have $|ONCE| = n - K$, $|TWICE| = K/2$, and an additional $n - (n - K) - K/2 = K/2$ winners in *ON*. Let $\overline{V}$ denote the maximal value across all winners in set *OFF*. We have $V_{\mathrm{ON}} \geqslant \sum_{j \in ONCE} v_j + \sum_{j \in TWICE} v_j + (K/2)\alpha\overline{V}$, and therefore

$$\frac{V_{\mathrm{ON}}}{V_{\mathrm{OFF}}} \geqslant \frac{\sum_{j \in ONCE} v_j + \sum_{j \in TWICE} v_j + (K/2)\alpha\overline{V}}{\sum_{j \in ONCE} v_j + 2\sum_{j \in TWICE}} \tag{4.1}$$

$$\geqslant \frac{(n-K)\alpha\overline{V} + (K/2)\overline{V} + (K/2)\alpha\overline{V}}{(n-K)\alpha\overline{V} + K\overline{V}} \tag{4.2}$$

$$= \frac{\alpha(n-K) + (K/2)(1+\alpha)}{\alpha(n-K) + K} \tag{4.3}$$

$$\geqslant \frac{K/2(1+\alpha)}{K} = \frac{1+\alpha}{2}, \tag{4.4}$$

where the second inequality follows by substituting the smallest possible values of agents in *ONCE* (counted equally in numerator and the denominator) and the largest possible values of agents in *TWICE* (counted twice in the denominator). The final inequality follows by analysis of the rate of change of the numerator and the denominator with respect to K for any $\alpha \in (0, 1]$, which is always more positive for the denominator than the numerator and therefore a valid lower bound is achieved by setting $K = n$. □

The assumption that there is "at least one buyer in every auction not holding an option" implies a requirement on both the input and the algorithm. But we can also provide a *sufficient* condition on the input for this property to hold, namely that for all auctions, that either a new bid arrived since the last auction or there is a bid i present at time t (i.e., with $\hat{a}_i \leqslant t \leqslant \hat{d}_i$) for which at least one higher-value bid had arrived prior to all previous auctions to occur in periods $\{\hat{a}_i, \ldots, t-1\}$.

[22] The mechanism of Hajiaghayi et al. [29] combines a greedy matching algorithm, in which the next item is allocated to the agent with the highest value that is currently unmatched, with a "critical value" payment in which successful buyers pay the smallest value that they could have reported and still been successfully matched.

[23] The empirical analysis that we present in Section 5 is performed for a sequence of auctions that complete in a sale on eBay, and in our simulations it will be the case that every item is sold in the options-based scheme.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

When $\alpha = 0$, that is with no bound on buyer values, we recover the competitive ratio of 2. The competitiveness goes to 1 as $\alpha \to 1$, as the values become more homogeneous. In the analysis of the eBay LCD market, we find the values of all buyers are bounded between \$200 and \$300, in which case $\alpha = 0.67$ and the competitive ratio is around 1.2, which implies that the total value of the allocation made by the options-based scheme is guaranteed to be within 17% of the value of the best-possible offline solution.

5. Empirical analysis

In this section, we present the results of an experimental study of the *average-case* performance of the options-based scheme for both efficiency and seller revenue. This study is in two parts. We first report results from an analysis of data collected from eBay on all auctions for a 19″ Dell LCD monitor sold during the summer of 2005. From this data we derive a population of buyers and sellers, including estimates of the (true) arrival and departure times, and (true) values of buyers. We adapt a non-parametric approach to estimate the values of buyers (Haile and Tamer [28], extended to dynamic auctions in Juda [35]), and couple this with bootstrapping to provide robustness. We estimate that the options-based scheme would provide a 4% improvement in efficiency and a 9% improvement in seller revenue over the status quo, and thus a significant improvement in total surplus in the market.

We also report the results from additional simulations designed to understand the performance of the options-based scheme in environments in which buyers have *substitutes* valuations or *complements* valuations. These simulations are inspired by the eBay data but are not directly performed in terms of this data because we do not attempt to identify the preferences of buyers with more complex valuations.[24] Buyer populations with substitutes valuations can hamper efficiency, although efficiency remains high when the valuations of a given buyer for the different items are either negatively correlated or uncorrelated. Buyer populations with complements valuations tend to achieve consistent efficiencies for different within-buyer correlation on the value of items. Finally, we study the role of liquidity by varying the buyer-to-seller ratio and find that for buyer-to-seller ratios that are typical in the eBay marketplace, the efficiency remains high even when buyers have substitutes valuations over many different items.

As discussed in the introduction, for the purpose of our first set of results in regard to the Dell LCD market it is sufficient to compare the average value and revenue per successful buyer, because there is no hold-up problem and the same number of buyers win in both eBay and the options-based scheme. For the purpose of substitutes valuations we normalize with respect to an online benchmark and for the purpose of complements valuations it will be sufficient to consider the total value (and revenue) achieved in the options-based scheme.

5.1. An eBay market for LCD screens

We collected data from eBay on all auctions for a 19″ Dell LCD monitor (Model E193FP) sold from 27 May, 2005 through 1 October, 2005, of which there were 1956 instances.[25] Assuming that each pseudonym represents a unique buyer, we observe 10,151 distinct bidders participating in these auctions. Given the data, our aim is to simulate a sequence of auctions for options that match the timing of auctions on eBay and with the true, underlying value of the buyers as identified from the eBay data. We define eBay seller revenue as the actual closing prices in the data. For eBay efficiency we compute the efficiency implied by the allocation on eBay and the estimate of the true values that we make for each winner.

For each auction on eBay that closes with a sale, we simulate a Vickrey auction for an option on the item. Auctions on eBay in which the item goes unsold are not considered within our simulation.[26] The sequence in our simulation is the same as the sequence on eBay with an auction scheduled to occur when it first *opens* on eBay. An auction that opened at 1:00:00pm on day 1 would be simulated before an auction that opened at 1:00:01pm on day 1. We schedule auctions at the time an auction opens rather than closes on eBay to allow for the possibility of re-posting an item that goes unsold. (Although this is only relevant with more general valuations because all options are exercised in the current context.)

We estimate the arrival, departure and value of each buyer on eBay from their observed behavior.[27] Arrival is estimated as the *first time that a buyer interacts with the eBay proxy*, while departure is estimated as the *latest closing time among eBay*

[24] When buyers have substitute preferences, it would be difficult to determine the entire set of substitutes in which a buyer may be interested, as a buyer on eBay may have never bid on all substitutes. When buyers have complements preferences, it would be difficult to determine based solely on their bidding behavior the extent to which they already possess the complementary goods. For example, if we were to observe a buyer bidding on a left shoe, there is no way to know definitively if the buyer already possesses a right shoe, or only intends to start bidding on a right shoe once a left shoe has been acquired.

[25] Specifically, search queries found all auctions where the auction title contained all of the following terms: 'Dell', 'LCD' and 'E193FP', while excluding all auctions that contained any of the following terms: 'Dimension', 'GHZ', 'desktop', 'p4' and 'GB.' The exclusion terms exist 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. Further information on the fields for each auction and how those fields were processed is provided in Juda [35].

[26] Unsuccessful auctions are likely either completely unseen by the buyer population or reserve auctions where the reserve price was not met. In either scenario, we consider these auctions too unique to include in the simulation. For example, if unsuccessful auctions were modeled in the options-based scheme, then more items would be sold in the options scheme than on eBay, making it significantly more difficult to compare the total seller revenue generated between the two markets.

[27] The relatively few buyers observed to have won multiple items on eBay in practice are simulated as multiple buyers with identical arrival, departure and value. At most one of these identical buyers will participate in any given auction.

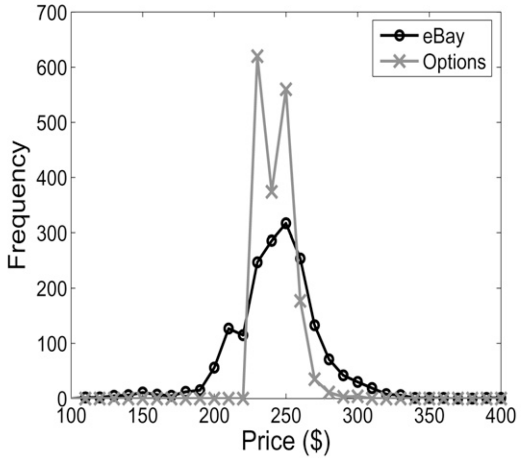

(a) The PDF of closing prices of Dell E193FP LCD screens using a worst-case estimate of buyer value. While the average closing price on eBay and the options scheme are comparable, the variance is significantly lower in the options scheme.

	Options	eBay
$\overline{Price}$	\$239.66	\$240.24
stddev(*Price*)	\$12	\$32
$\overline{Value}$	\$263	\$244
$\overline{BuyerSurplus}$	\$23	\$4

(b) The average price paid per good, average buyer value among winners, and average winning buyer surplus on eBay for Dell E193FP LCD screens as well as the simulated options-based market using worst-case estimates of buyer values. (Note: All items are sold in both markets.)

Fig. 1. Comparisons between empirically observed results on eBay and simulation results of the options scheme using a worst-case estimate of buyers' true valuations.

auctions in which a buyer participates. Both are clearly conservative, but adopted in the interest of simplicity. We note that less conservative estimates of these timing constraints would improve the performance of the options-based auction in simulation because there would be greater competition. We are careful about the end effects of the first few days of data and the last few days of data.[28]

In our first experiments, we adopt a conservative estimate of the true value of a buyer on eBay, estimating this simply as the *highest bid this buyer was observed to have placed in any LCD auction.* Fig. 1(a) shows the distribution of closing prices both on eBay and in the simulated options scheme under this "worst-case" assumption. The closing price in the simulation is defined as the final exercise price (i.e., after price matching). The seller revenue is similar between the two schemes while the estimated effect on efficiency suggests an advantage for the options-based scheme (as the winners in the options-based scheme possess a higher average value for the item than on eBay).[29] See Fig. 1(b). While the winning buyers on eBay have an estimated mean value of \$244, the options scheme's winning buyers have a mean value of \$263 (a 7% increase). Consumer surplus, which measures the average buyer utility over all winning buyers, increases in the options-based scheme from \$4 to \$23.[30] We notice also that the standard deviation of prices is significantly reduced in the options-based scheme.

Fig. 2(a) shows a very distinct difference between eBay and the options scheme. We plot the closing price of an auction against the patience of the auction's winner. In the options scheme, not only do buyers with a larger patience generally pay lower prices than buyers with smaller patience but the variance of price paid decreases with patience. Quantifying this effect, we ran linear regressions of price versus the log of patience in seconds for both eBay and the options scheme. While we cannot say that patience is correlated with price on eBay (with a 95% confidence interval of the patience coefficient ranging from -0.517 to 0.642), patience is negatively correlated with price in the options-based scheme (with a 95% confidence interval of the patience coefficient ranging from -1.995 to -0.938).

In our second set of experiments, we adopt a less conservative estimate of the true value of a buyer by using the non-parametric methods of Haile and Tamer [28], generalized to apply to dynamic auctions.[31] Fig. 3(a) shows the distribution of actual closing prices in eBay and in the options scheme as simulated with this new, less conservative estimate. The mean

[28] When running the simulations, the final ten days worth of observed auctions are not simulated. Ten days also is past the 90th percentile of the distribution of buyer patience. Buyers bidding toward the end of the window are likely to bid higher in the future on eBay than their current bids (this is suggested by the regression analysis in Juda [35]), which results in low estimates for buyer value among these buyers. Similarly, we must allow for a start effect in which the initial seller revenue in simulation will be artificially high because no buyers in the initial period will be marked as already traded when in fact they would have if the simulation had started earlier. For this, the options simulation starts from the initial period but seller revenue is only accounted after the first ten days.

[29] Note that in markets where buyers are only interested in a single item, every item is sold in the options based scheme. Consequently, an increase in average price or winners' average value does also imply an increase in total revenue and total allocative value (and hence efficiency).

[30] This improvement in consumer surplus occurs because of the redistribution effect of an improved allocation and represents the increase in value to winning buyers. But these are lower bounds on consumer surplus given the conservative estimate of true value and thus overstate the relative improvement in consumer surplus provided by the options-based scheme. Note that buyer surplus on eBay is above zero even when conservatively estimating buyer values, as there are buyers who had submitted losing bids in other auctions that were higher than the winning prices in auctions won.

[31] See Juda [35] for more information. In particular, we are able to estimate that the true values of buyers on eBay are 15% greater than their observed maximum bids. This estimate is based on identifying a multiplicative factor that separates the distribution of observed bid values from a conservative estimate on the distribution of underlying values that can be derived through analysis of order statistics and simple, reasonable assumptions.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

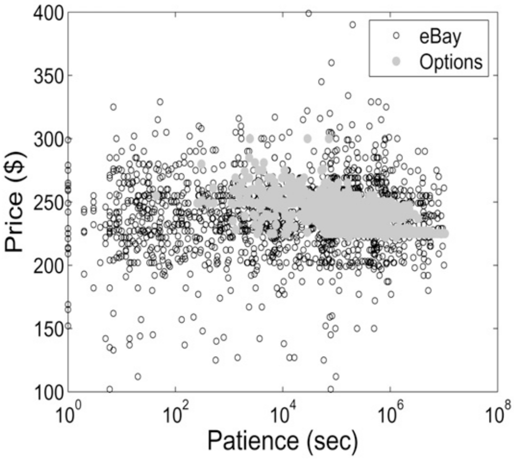

(a) A scatter plot of the price paid by winners against the patience of the buyer (in secs) for Dell E193FP LCD screens using worst-case estimates of buyer values.

	Options	eBay
Intercept	271.6	286.7
log(*Patience*)	−1.466*	0.063
	(0.270)	(0.295)

(b) Regressions of price versus the log of patience.
*indicates statistically significant at least at the 5% significance level. Standard error below coefficients.

Fig. 2. Price versus patience.

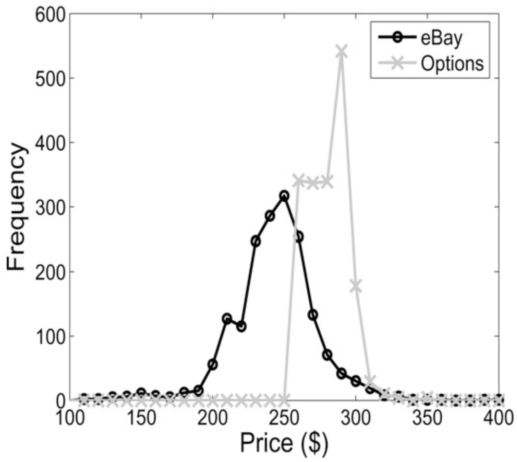

(a) The PDF of closing prices of Dell E193FP LCD screens using a less conservative estimate of buyer value.

	Options	eBay
$\overline{\mathit{Price}}$	\$275.80	**\$240.24**
stddev(*Price*)	\$14	**\$32**
$\overline{\mathit{Value}}$	\$302	\$281
$\overline{\mathit{BuyerSurplus}}$	\$26	\$40

(b) The average price paid per good, average buyer value among winners, and average winning buyer surplus on eBay and under the simulated options market, for Dell E193FP LCD screens and a less conservative estimate of buyer values. (Note: All items are sold in both markets. Numbers in bold identical to those in Fig. 1(b).)

Fig. 3. Comparisons between empirically observed results on eBay and simulation results of the options scheme using a less conservative estimate of buyers' true valuations.

price in the options scheme is now significantly higher than the mean price on eBay (\$240 on eBay, \$276 in the options scheme). The standard deviation on closing prices in the options scheme is also significantly less, being \$32 on eBay vs. \$14 in the options scheme. The estimated efficiency of the options-based scheme again remains higher than that on eBay because, while the winning buyers on eBay are estimated to have a mean (true) value of \$281, the winners in the options scheme are estimated to have a mean value that is 7.5% higher (at \$302).[32] It bears emphasis that in reporting these results, the *same* value estimates that are adopted in the simulation of the options scheme are also adopted in estimating the total value of the allocation implemented on eBay.

Comparing Fig. 1(b) and Fig. 3(b), we see a variety of changes between the two value estimations. Within the options based schemes, as the participants are estimated as having a higher value in Fig. 3(b), but submitting their true value as they have a dominant strategy to do so, average price and winning buyer value are appropriately higher. Alternatively, within eBay, a higher estimate of value does not impact the actual set of empirically observed results. Therefore, a higher estimated value per winning buyer coupled with identical prices results in a significant change in buyer surplus.

Bootstrapping. While the simulation of the options-based market suggests better efficiency and seller revenue than on eBay, a reasonable concern may be that the performance of the options-based market is influenced by specific details of

[32] The consumer surplus for buyers in the options scheme is below that on eBay, but this could be easily addressed by the use of fees or other methods to redistribute surplus. The significant effect is that both revenue and value have increased, representing greater *total* surplus.

Table 2
Average price paid per good, buyer value per winner, and buyer surplus per winner, over 50 bootstrapped simulations of the options-based scheme with less conservative estimate of buyer values (with standard deviations in parentheses). The table also provides the ratio of the performance of the options-based scheme (averaged over these bootstrapped simulations) and eBay.

	Options	$\frac{Bootstrap}{Options}$	Bootstrap	$\frac{Bootstrap}{eBay}$	eBay
$\overline{Price}$	\$275.80	0.95	\$261.89 (\$1.37)	1.09	\$240.24
$\overline{Value}$	\$302	0.97	\$292 (\$0.86)	1.04	\$281
$\overline{BuyerSurplus}$	\$26	1.14	\$30 (\$0.95)	0.74	\$40

the particular buyer population considered. To alleviate these concerns, we also perform a set of bootstrapped simulations. Rather than using the 10,151 unique buyers observed to have participated on eBay in the simulation we instead simulate the options-based market using 10,151 buyers where each buyer is drawn uniformly with replacement among all buyers observed to have bid on eBay. This creates a buyer population that is similar to that observed on eBay, while providing insight into how sensitive the results are to the exact combination of buyers observed. Table 2 provides the average results of 50 bootstrapped simulations, together with the performance of the options-based scheme without bootstrapping and the results for the actual allocation in the eBay market. The results support an estimate of efficiency that is 4% greater than on eBay and seller revenue that is 9% greater than on eBay.

5.2. Simulation: Substitutes preferences

While the options-based market appears effective when the bidding population has simple preferences and wants only a single good, we now examine the efficacy of the system when buyers have substitutes preferences. To see why substitutes preferences can be problematic consider the following:

Example 6. Alice values one of either an apple or banana by Tuesday for \$10. Bob values one of either an apple or banana by Tuesday for \$8. On Sunday, an apple auction is held where Alice's proxy wins an option for the apple for \$8. On Monday, a banana auction is held where Alice's proxy wins an option for the banana for \$8. At the end of Tuesday, Alice's proxy exercises one of her two options, returning the other option, while Bob leaves the market having acquired nothing.

Clearly, it would have been more efficient for Bob to have won an option for only one of the pieces of the fruit. However, the scheme has Alice's proxy holding both options. We refer to this as the *holdup problem*.[33]

We first consider a market in which buyers have substitutes preferences over two items. Inspired by the observed population of eBay, we consider a market with a 120-day time period where each buyer's value is distributed normally over a Gaussian distribution for Monitor A with mean \$265 and standard deviation \$45 and for Monitor B with mean \$240 and standard deviation \$20. The value for the bundle of A and B is the maximum of the two values. 5000 buyers arrive uniformly over the 120-day time period, and with patience distributed according to a Normal distribution with a mean of 3.9 days and a standard deviation 11.4 days (as was observed to be the mean and standard deviation of patience among the eBay buyers), and truncated to be non-negative and rounded to the nearest day. We model 2000 sellers that enter the market uniformly over the 120-day time period, with the patience of each seller distributed by a Normal distribution with a mean of 7 days and a standard deviation 1 day. Each seller offers one of Monitor A or Monitor B with equal probability. In the simulation, an auction is scheduled for a seller for an option on her good immediately upon arrival and the good is re-posted when a seller's option is returned unexercised before her departure.[34]

In the experiments, we consider both a positive and a negative correlation between the value that a buyer has for Monitor A and Monitor B. Positive correlations might exist if some buyers possess generally higher valuations across all items than other buyers. Negative correlations suggest that buyers have strong "tastes" for each item and it is likely that either one or the other item will appeal to a particular buyer but not both.[35]

Table 3 shows summary statistics for the performance of the market in this setting. Rather than provide a comparison to eBay, we compare the value of the options-based allocation to the value of a greedy online benchmark. The greedy online

[33] To provide a counterpoint, note that if buyers have linear-additive values on individual items then they will always exercise every option they win and no options will be returned and go unsold. This is because the maximum marginal value of each item is exactly the value a buyer will ultimately realize in exercising that option, whatever the details of the other goods that it wins.

[34] For this to be possible in practice without introducing new opportunities for buyer manipulation we would need to prevent a buyer's proxy from rebidding on the same item within period $\Delta_{\max}$ (as discussed in Section 3.1). This was not done within the simulation but would have a negligible effect on the results.

[35] To model positive correlation, if a buyer's valuation for Monitor A is x standard deviations above the mean, her valuation for Monitor B is set to x standard deviation above the mean (cf. for x standard deviations below the mean). Alternatively, to model negative correlation, if a buyer's valuation for Monitor A is x standard deviations above the mean, her valuation for Monitor B is set to x standard deviations below the mean (cf. for x standard deviations below the mean).

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

Table 3
Market performance (averaged over 30 instances) with 5000 buyers and 2000 sellers in a 120 day marketplace and buyers with substitutes preferences on 2 items.

Buyers' values	Items sold (%)	Total value	Buyer surplus (% value)	Sellers returned (%)	Value % of online benchmark
− Correlated	100.0	569,665	7.6	0.02	96.5
Uncorrelated	87.1	500,571	8.5	18.3	85.2
+ Correlated	61.4	362,704	9.2	41.7	63.3

Table 4
Market performance (averaged over 30 instances) with 5000 buyers and 3000 sellers in a 120 day marketplace. Buyers possess substitutes preferences over 3 items.

Buyers' values	Items sold (%)	Total value	Buyer surplus (% value)	Sellers returned (%)	Value % of online benchmark
Uncorrelated	75.0	632,504	7.5	33.8	74.5
+ Correlated	44.4	388,191	8.6	58.1	47.2

benchmark is computed as follows. In each period, it looks at all buyers and sellers in the market and computes a tentative, value-maximizing allocation of goods among the population using a mixed-integer program. For each pair of buyers and sellers that are matched in this allocation, if either one departs in this period then the trade between them is committed. Otherwise, both buyers and sellers carry over into the next period and continue to be available for provisional allocation.

The "sellers returned" statistic in Table 3 indicates the average number of sellers whom have an option returned unsold. One can understand by comparing the fraction of items sold and the fraction of "sellers returned" the number of items that are sold successfully on second (and later) attempts; e.g., for uncorrelated values 18.3% of items are initially unsold but only $100 - 87.1 = 12.9\%$ of items are unsold eventually.

When values across the two items are negatively correlated, the market in effect breaks itself up into two disjoint markets, one for the first and one for the second item, because buyers do not typically possess a sufficiently high value on both of the items to be competitive on both. Because of this, all items are typically sold and the estimated efficiency (based on the percentage as related to the greedy online benchmark) is high and there is only a slight holdup problem. On the other hand, when values across the two items are either uncorrelated or positively correlated, buyers are more likely to hold options on both items, thus causing holdup problems and blocking lower-valued buyers who may never hold an option.

Similar results are found when we consider substitute preferences over three items (where the distribution of value for the third item is normal with mean \$265 and standard deviation \$45), and scaling the number of sellers to 3000 from 2000 to keep the same number of each item supplied on average (see Table 4). When the values across the three items are positively correlated for a buyer, fewer items are sold and the market efficiency falls (as defined with respect to the greedy online benchmark).[36]

As a third experiment with substitutes preferences, we investigate the extent to which the *number* of items for which each buyer has some value can affect the performance of the market. For this, we consider a market with 1000 buyers and 1000 sellers and 10 different kinds of items. Buyers possess substitutes preferences on between 1 and 10 items, always with uncorrelated values across items.[37] The value on a bundle of items is defined as

$$v_i(L) = \max_{G_k \in L} v_i(k), \tag{5.1}$$

i.e., these are pure substitutes valuations. Fig. 4 illustrates the average market performance of the options-based scheme. As the number of items in which a buyer is interested increases (the "size" of a buyer's valuation), the additional competition reduces buyer surplus. In addition, efficiency (as defined with respect to the greedy online benchmark) falls because the holdup problem gets worse. However, it is interesting that the number of items sold remains fairly constant at around 53%. In Section 5.4 we will see that the effect of the size of a buyer's substitutes valuation on efficiency depends on the buyer-to-seller ratio and is significantly improved for higher buy-side competition.

5.3. Simulation: Complements preferences

We now consider buyers with complements valuations for items. While a buyer has a value for each individual item, in this simulation she is also interested in acquiring both items. The synergy (or lack thereof) of acquiring both items is provided via a Gaussian distributed *multiplicative factor*, $\beta \in (-1, 1)$, of the sum of the values of the individual components

[36] There being no simple way to define negatively correlated values on three items we just present results for uncorrelated and positively correlated values.
[37] The specific distributions from which valuations are drawn are as follows: $v_i(1) \sim N(260, 45)$, $v_i(2) \sim N(240, 10)$, $v_i(3) \sim N(250, 5)$, $v_i(4) \sim N(280, 5)$, $v_i(5) \sim N(260, 20)$, $v_i(6) \sim N(230, 55)$, $v_i(7) \sim N(220, 60)$, $v_i(8) \sim N(245, 5)$, $v_i(9) \sim N(260, 5)$, $v_i(10) \sim N(290, 5)$.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

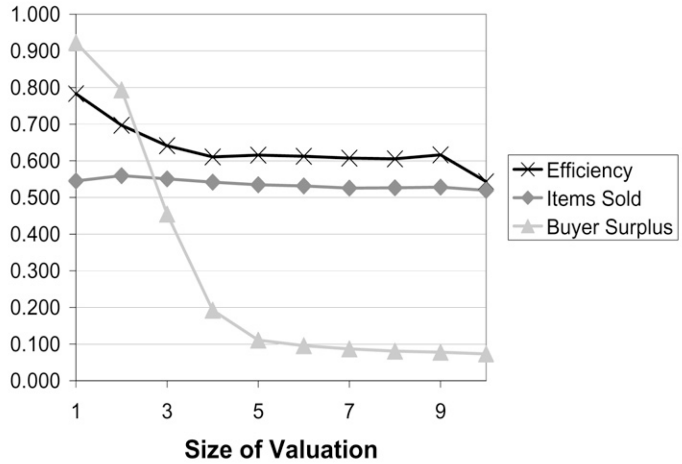

Fig. 4. Market performance (averaged over 30 runs) with 1000 buyers and 1000 sellers in a 120 day marketplace with 10 different kinds of goods being offered. Buyers have substitutes preferences over a varying number of items.

Table 5
Market performance with 5000 buyers and 2000 sellers in a 120 day marketplace. Buyers possess complements preferences over 2 items.

Buyers' values	Items sold (%)	Number of buyers	Average bundle size	Total value	Buyer surplus (% value)	Sellers returned (%)
− Correlated	80.0	1253	1.28	477,388	6.4	31.4
Uncorrelated	80.7	1139	1.42	487,841	7.5	29.7
+ Correlated	77.3	926	1.67	477,257	9.0	35.1

of the bundle, such that the value for two items, $v_i(\{A\} \cup \{B\}) = (1 + \beta)\ (v_i(\{A\}) + v_i(\{B\}))$, and we define $\beta \sim N(0, 0.1)$. We also consider correlation (both negative and positive) across the single-item values on different items. We simulate a market with 5000 buyers and 2000 sellers and 120 days, adopting the same set-up as in the substitutes experiments.

Table 5 shows summary statistics for the performance of the market. The average bundle size is the average number of items allocated to winning buyers. This increases from uncorrelated to correlated values across items. This time the correlation seems unimportant, and no matter whether values across items are positively or negatively correlated, we see that around 80% of the items are sold. The number of items sold is relatively high because buyers who hold multiple options can generally exercise both options because of their complements valuations. The number of items sold is not limited as much because of holdup with complements as with substitutes preferences. An individual buyer's surplus tends to be higher when values are positively correlated because, conditioned on winning at all, she is likely to have higher value on both items and thus higher surplus (while continuing to benefit from price matching). In a simulation in which the number of sellers is increased to 3000 and buyers have complements valuations on *three* items, but otherwise unchanged, the number of items sold remains at around 80% of supply for both uncorrelated and positively correlated across-item values.

5.4. Simulation: Market liquidity

For our final study we consider buyers with substitute preferences and uncorrelated values across items and vary the liquidity in the market. For this, we vary the *buyer-to-seller* ratio, which is the ratio of the total number of buyers to sellers in the market, and fix the number of sellers at 500. Fig. 5 plots the efficiency (calculated as the ratio of total value of goods allocated in the options-based scheme to a greedy online allocation that approximates the total possible realizable value) against buyer-to-seller ratio. We adopt this online benchmark because we are interested in understanding how our solution to the strategic problem (through options, proxies and price-matching) affects performance in relation to a non-strategyproof online algorithm.

What is particularly interesting is that efficiency initially tends to decrease with increasing buyer-to-seller ratio (for all numbers of items other than one) but then increases again.[38] For extremely low buyer-to-seller ratios there is little competition in the market and buyers do not hold up each other too badly and the efficiency is high even with substitutes valuations on many items. As the buyer-to-seller ratio increases, efficiency falls as more buyers are blocked because of the holdup problem. At some point, though, efficiency begins to increase again because it becomes less likely that any single

[38] For the problem instance where buyers are only interested in a single item, efficiency is always very close to 100% for all buyer-to-seller ratios. Note also that while there is a data point above 100%. This is not a spurious result because the online benchmark is not guaranteed to be optimal, and can on occasion be out-performed by the options-based scheme.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002

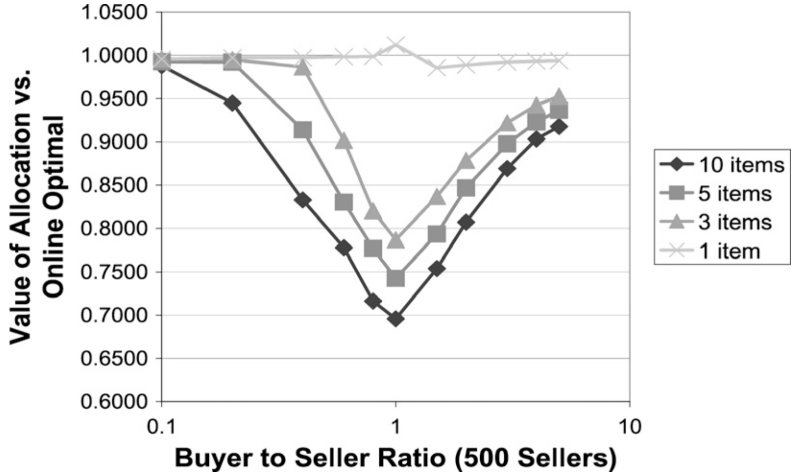

Fig. 5. Efficiency of the options-based scheme at various buyer-to-seller ratios (averaged over 30 runs). Buyers have substitutes preferences on different numbers of items. Number of sellers fixed at 500. Number of days fixed at 120.

buyer will be competitive on *multiple* items. The effect is to separate the market across different types of goods, with each buyer tending to only win on one or two goods and therefore causing only a slight holdup problem and providing higher efficiency. On eBay, for example, we see a buyer-to-seller ratio well above 5:1 in the LCD market and might expect efficiency to remain high even when buyers have large substitutes valuations.

6. Discussion: Improving market efficiency

The empirical study in the previous section not withstanding, two factors that limit the market efficiency of the options-based scheme are:

(1) Proxy agents hold onto options that they will likely not exercise.
(2) Proxy agents bid their maximum marginal value for options.

Regarding the first point, notice that while proxy agents exercise every option that they hold when items have constant marginal value (e.g., when a buyer wants a single item, or has a linear valuation function), a number of options will be returned in general. This occurs quite frequently in our experiments, although a number of these items are ultimately sold upon reposting in a new auction. One simple improvement that can be adopted is to allow a seller who wishes to leave the marketplace but has an issued option to leave (with the good) if the proxy already knows that the option will definitely go unexercised. In so doing, sellers need only remain in the market while there is some possibility that their option will be exercised. However, it is difficult to provide sellers with additional flexibility, for example to allow a seller to offer multiple options for the same unit of a good, without compromising the strategyproofness of the market for buyers. If faced with options that do not provide an *exclusive* right to exercise and receive a good (for instance if two options are issued on the same unit and both proxies seek to exercise them), then proxy agents would have an incentive to bid for multiple options on the same good.

Regarding the second point, notice that proxy agents may be submitting excessively large values when bidding maximum marginal values. Even if a proxy is already guaranteed utility z on a bundle of options (based on their current exercise price), the proxy will still go ahead and bid for options on items that could not possibly bring utility of more than z at *any* exercise price, including zero. The bid price is not adaptive to options and exercise prices already secured. We would wish to provide a proxy with a less aggressive bidding strategy and prevent a proxy from acquiring options on items that are currently priced too high to be exercised and especially those that will never be exercised. A natural candidate for a less aggressive bidding strategy – which would still provably acquire all options that could possibly be useful – is to bid the *maximum willingness to pay* given its current allocation of options and the current exercise prices on these options. That is, the proxy should factor in its current state in deciding how much to bid.

For example, suppose that Alice values exactly one piece of fruit (either an apple at \$10, a banana at \$5, or an orange at \$5). If Alice's proxy already holds an option for an apple with an exercise price of \$8, she might only bid \$3 (instead of \$5) in future auctions for bananas and oranges because securing an option for a banana or an orange at a price above \$3 would only be dominated by the apple option. Similarly, if Alice's proxy already holds an option for an apple with an exercise price of \$2, she should not bid at all for bananas or oranges, as the maximum surplus possible from acquiring a banana or an orange is guaranteed to be less than the surplus of $\$10 - 2 = \$8$ already guaranteed for the apple.

More formally, let $\mathcal{O}^t$ denote the set of options held by a proxy at time t, $\gamma(L) \subseteq G$ denote the goods that correspond to some subset $L \subseteq \mathcal{O}^t$ of these options, and $p^t(L) = \sum_{k \in L} p^t(k)$ denote the total exercise price in period t for this set of options where $p^t(k)$ is the current exercise price on the option for good G_k. Let $\hat{u}_i^t(L) = \hat{v}_i(\gamma(L)) - p^t(L)$ denote the reported utility of the buyer for options, L. Let L^{t*} denote the set of options that maximizes this reported utility. Given this, define the *maximal willingness to pay* for an option on an item k given the current state of the proxy as:

$$wtp_i^t(k) = \max_S \big[0, \min\big[\hat{v}_i\big(S \cup \{G_k\}\big) - \hat{u}_i^t\big(L^{t*}\big), \hat{v}_i\big(S \cup \{G_k\}\big) - \hat{v}_i(S)\big]\big], \tag{6.1}$$

where $\hat{v}_i(S \cup \{G_k\}) - \hat{u}_i^t(L^{t*})$ considers the utility already guaranteed with the current options, and $\hat{v}_i(S \cup \{G_k\}) - \hat{v}_i(S)$ is the maximal marginal value of good G_k. This expression calculates the greatest amount a buyer will possibly be willing to spend on an item given the current options held and with uncertainty as to what future items will appear in auctions and about future option prices (assume all $S \cup \{G_k\}$ are free while the prices on items in L^{t*} remain the same). However, this scheme cannot be implemented in the proxied architecture without forfeiting truthfulness:

Example 7. Both Alice and Bob have substitutes valuations. Alice values either one ton of sand for \$2000, one ton of stone for \$1900 and both for \$2000. Bob values either one ton of sand for \$1800, one ton of stone for \$1500 and both for \$1800. Both buyers have a patience of 2 days. On day one, a sand auction is held, where Alice's proxy bids \$2000 and Bob's bids \$1800. Alice's proxy wins an option to purchase sand for \$1800. On day two, a stone auction is held, where Alice's proxy bids \$1700 (as she has already obtained a guaranteed \$200 of surplus from winning a sand option, and so reduces her stone bid by this amount), and Bob's bids \$1500. Alice's proxy wins an option to purchase stone for \$1500. At the end of the second day, Alice's proxy would exercise the option she holds for stone with an exercise price of \$1500 to obtain a good valued for \$1900, and so obtains \$400 in surplus. Now consider what would happen if Alice instead lies, declares that she values only stone, and for \$1900. On day one, a sand auction is held, where Bob's proxy bids \$1800. Bob's proxy wins an option to purchase sand for \$0 (because Alice's proxy stays out). On day two, a stone auction is held, where Alice's proxy bids \$1900, and Bob's bids \$0 (as he has already obtained at least \$1800 of surplus from winning the sand option, and so is not interested in winning the stone option). Alice's proxy wins the stone option with an exercise price of \$0, achieving \$1900 in surplus.

By misrepresenting her valuation, Alice was able to secure higher surplus by giving more surplus to Bob and therefore reducing the competition that she faced in a future auction. The basic tenet of strategyproof mechanisms requires that the prices faced by a bidder, such as Alice, should be independent of her strategy. This has been compromised because there is now some potential for Alice to influence Bob's bids in the future and in turn the price that she will face. The technical problem is that the value, $\hat{u}_i^t(L^{t*})$, in Eq. (6.1) is the amount of surplus already guaranteed to buyer i, and this now depends on the proxy bids of some other buyer. A slight modification to the option-based scheme, presented in Juda [35], can address this problem in a restricted ("set-valued") class of valuation domains, namely those in which there are two kinds of goods A and B, and each buyer is either indifferent between A and B or interested in the bundle AB. However, we do not know of a remedy to this problem that will allow proxy agents to bid adaptively in more general valuation domains and leave this for future work.

We see a number of additional directions for future work. These include:

(a) *Allowing buyers to return options to the market early*. We ask whether a scheme can be developed in which a buyer can return an option as soon as her proxy determines that the option will never be exercised given the other options it holds. Such a return would reduce the holdup problem and improve efficiency.
(b) *Allowing buyers to demand multiple units of a particular item*. This seems quite challenging because a naive solution will leave proxies facing a tradeoff between competing to acquire additional options on a particular good or choosing not to compete, perhaps eliciting greater opportunities for price matching on options that they already hold.
(c) *Allowing sellers to have values on bundles rather than items*. Scenarios may exist where sellers are interested in selling multiple items in a single lot. The amount of information required to perform bookkeeping and track the minimal price possible across auctions is greater than in the current system; in particular, it seems that all bundles in which a bumped buyer is interested would need to be stored.
(d) *Allowing a buyer's value to vary with time*. For example, a buyer's valuation may decrease over time. The difficulty in retaining strategyproofness is that a naive solution will leave a proxy with the need to decide if it should delay the exercising of its option, risking the degradation of value but possibly gaining a lower exercise price. However, the system can be extended to enable buyers to increase their declared value for an item, express a value for new items, or further delay their departure time, while retaining strategyproofness for buyers that have certain valuations, as in our standard model. The effect would be to allow for new demand information to be immediately incorporated within the market should such a need develop. See Juda [35] for further discussion about these extensions.
(e) *Mitigating strategic problems facing sellers.* While our results suggest that sellers can increase their revenue over the eBay protocol when all sellers participate and are straightforward – reporting true arrival and departure information – such behavior may not always be in the best interest of each seller. For example, a seller may do better by delaying her entry into the market if she believes the market is currently "cold" but will soon become "hot" such that there is more competition and higher prices.

(f) *Exploring hierarchical options-based schemes*. The use of options in this paper shifts the sequential auction problem across auctions in a site to a similar problem across different auction sites (e.g., eBay and Yahoo!). It is interesting to consider whether a hierarchical options-based scheme can be developed to help to mitigate the strategic complexity that buyers likely face when bidding across sites.
(g) *Exploring interactions between sites*. When multiple auction sites are competing for buyers and sellers to conduct business within their market, one can imagine one site going to great lengths to drive traffic toward them. Within the options-based scheme, one potentially malicious strategy would be for a site to inject bad buyers or sellers into the other site's market. While strong identities can dramatically reduce the ability to inject malicious market participants, exploring how robust the options-based scheme is to the injection of malicious participants would be interesting.

Other future areas of work, as well as some progress on extensions of the options-based scheme, are found in Juda [35].

7. Conclusions

We have proposed a novel, options-based method for resolving the strategic difficulties faced by buyers in uncoordinated electronic marketplaces. Our solution to the sequential auction problem is to require that buyers submit bids to *mandatory* bidding proxies which then bid for options on goods and exercise options to maximize reported buyer utility. By allowing proxy agents to match the exercise prices of options to the lowest price that would have been possible through careful timing of the proxies' bid, it becomes a dominant strategy for buyers to report their true value and true temporal constraints to the market. The proxy agents act across multiple auctions and bring simple, truthful bidding into an equilibrium while retaining the dynamic arrivals and departures of open, Internet marketplaces and operating without batching auctions. An empirical analysis that is informed by data collected on eBay suggests that the options-based scheme can provide an improvement in efficiency and revenue over eBay of around 4% and 9% respectively. A series of experiments to examine the effect of a holdup problem that can exist when buyers have general valuations shows that this is mitigated when competition is either very low or high, or when individual buyers have negatively correlated values across items. Efficiency also remains relatively high when buyers have complements valuations, where efficiency is augmented when sellers can re-post goods for sale when options go unexercised.

Acknowledgements

The first author would like to thank Pai-Ling Yin and Barbara J. Grosz for useful guidance at various stages of this work. Thanks also to Aaron L. Roth and Kang-Xing Jin for technical support. This work has benefited from presentations at INFORMS 2006, ACMEC 2006, the 6th Workshop on Agent Mediated Electronic Commerce (2004), Harvard Business School, The Wharton School of the University of Pennsylvania, Columbia University Graduate School of Business, U.C. Berkeley School of Information, University of Connecticut School of Business, and Harvard University's EconCS and ITM seminars. The authors also would like to acknowledge the constructive and careful comments of the three anonymous reviewers.

References

[1] P. Anthony, N.R. Jennings, Developing a bidding agent for multiple heterogeneous auctions, ACM Transactions on Internet Technology 3 (3) (2003) 185–217.
[2] K.J. Arrow, The property rights doctrine and demand revelation under incomplete information, in: M. Boskin (Ed.), Economics and Human Welfare, Academic Press, New York, 1979.
[3] S. Athey, I. Segal, An efficient dynamic mechanism, Tech. Rep., Harvard University and Stanford University, 2007.
[4] R. Bapna, P. Goes, A. Gupta, Y. Jin, User heterogeneity and its impact on electronic auction market design: An empirical exploration, MIS Quarterly 28 (1) (2004) 21–43.
[5] Y. Bartal, R. Gonen, N. Nisan, Incentive compatible multi unit combinatorial auctions, in: Proceedings of Ninth Conference of Theoretical Aspects of Rationality and Knowledge (TARK'03), 2003.
[6] D. Bergemann, J. Välimäki, Efficient dynamic auctions, Tech. Rep. Cowles Foundation Discussion Paper No. 1584, Yale University, 2006.
[7] F. Black, M. Scholes, The pricing of options and corporate liabilities, Journal of Political Economy 81 (3) (1973) 637–654.
[8] C. Boutilier, M. Goldszmidt, B. Sabata, Sequential auctions for the allocation of resources with complementarities, in: Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI'99), 1999.
[9] C. Boutilier, H.H. Hoos, Bidding languages for combinatorial auctions, in: Proc. 18th International Joint Conference on Artificial Intelligence (IJCAI'01), 2001.
[10] J. Bredin, D.C. Parkes, Q. Duong, Chain: A dynamic double auction framework for matching patient agents, Journal of Artificial Intelligence Research 30 (2007) 133–179.
[11] E. Brynjolfsson, B. Kahin, Understanding the Digital Economy, MIT Press, 2000.
[12] E. Brynjolfsson, M.D. Smith, Frictionless commerce? A comparison of Internet and conventional retailers, Management Science 46 (4) (2000) 563–585.
[13] A. Byde, C. Preist, N.R. Jennings, Decision procedures for multiple auctions, in: Proc. 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'02), 2002.
[14] M.M. Bykowsky, R.J. Cull, J.O. Ledyard, Mutually destructive bidding: The FCC auction design problem, Journal of Regulatory Economics 17 (3) (2000) 205–228.
[15] E. Cantillon, M. Pesendorfer, Auctioning bus routes, in: P. Cramton, Y. Shoham, R. Steinberg (Eds.), Combinatorial Auctions, MIT Press, 2006, pp. 573–592. Chapter 22.
[16] J. Cassell, T. Bickmore, Negotiated collusion: Modeling social language and its relationship effects in intelligent agents, User Modeling and Adaptive Interfaces 12 (2002) 1–44.

[17] R. Cavallo, D.C. Parkes, A.I. Juda, A. Kirsch, A. Kulesza, S. Lahaie, B. Lubin, L. Michael, J. Shneidman, TBBL: A Tree-Based Bidding Language for iterative combinatorial exchanges, in: Multidisciplinary Workshop on Advances in Preference Handling (IJCAI'05), 2005.
[18] R. Cavallo, D.C. Parkes, S. Singh, Efficient online mechanisms for persistent, periodically inaccessible self-interested agents, Tech. Rep., Harvard University, 2007.
[19] Y. Chen, C. Narasimhan, Z.J. Zhang, Consumer heterogeneity and competitive price-matching guarantees, Marketing Science 20 (3) (2001) 300–314.
[20] E.H. Clarke, Multipart pricing of public goods, Public Choice 11 (1971) 17–33.
[21] C. d'Aspremont, L. Gerard-Varet, Incentives and incomplete information, Journal of Public Economics 11 (1979) 25–45.
[22] A.K. Dixit, R.S. Pindyck, Investment under Uncertainty, Princeton University Press, 1994.
[23] I. Domowitz, B. Steil, Automation, trading costs, and the structure of the securities trading industry, in: Brookings-Wharton Papers on Financial Services, 1999.
[24] E.H. Gerding, R.K. Dash, D.C.K. Yuen, N.R. Jennings, Bidding optimally in concurrent second-price auctions of perfectly substitutable goods, in: Proc. 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07), 2007.
[25] R. Gopal, S. Thompson, Y.A. Tung, A.B. Whinston, Managing risks in multiple online auctions: An options approach, Decision Sciences 36 (3) (2005) 397–425.
[26] A. Greenwald, J.O. Kephart, Shopbots and pricebots, in: Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI'99), 1999.
[27] T. Groves, Incentives in teams, Econometrica 41 (1973) 617–631.
[28] P.A. Haile, E. Tamer, Inference with an incomplete model of English auctions, Journal of Political Economy 111 (1) (2003) 1–51.
[29] M.T. Hajiaghayi, R. Kleinberg, M. Mahdian, D.C. Parkes, Online auctions with re-usable goods, in: Proc. 6th ACM Conference on Electronic Commerce (EC'05), 2005.
[30] M.T. Hajiaghayi, R. Kleinberg, D.C. Parkes, Adaptive limited-supply online auctions, in: Proc. 5th ACM Conference on Electronic Commerce (EC'04), 2004.
[31] J. Hastad, Clique is hard to approximate within $n^{1-\epsilon}$, Acta Mathematica.
[32] K. Hendricks, I. Onur, T. Wiseman, Preemption and delay in eBay auctions, Working Paper, University of Texas, Austin, 2005.
[33] J.D. Hess, E. Gerstner, Price-matching policies: An empirical case, Managerial and Decision Economics 12 (4) (1991) 305–315.
[34] A. Iwasaki, M. Yokoo, K. Terada, A robust open ascending-price multi-unit auction protocol against false-name bids, Decision Support Systems 39 (2005) 23–40.
[35] A.I. Juda, Coordination and costly preference elicitation in electronic markets, doctoral thesis, Harvard University, 2007.
[36] R. Lavi, N. Nisan, Competitive analysis of incentive compatible on-line auctions, in: Proc. 2nd ACM Conference on Electronic Commerce (EC'00), 2000.
[37] J. Lee, K. See, Trust in automation: Designing for appropriate reliance, Human Factors 46 (2004) 50–80.
[38] Y.J. Lin, Price matching in a model of equilibrium price dispersion, Southern Economic Journal 55 (1) (1988) 57–69.
[39] D. Lucking-Reiley, D.F. Spulber, Business-to-business electronic commerce, Journal of Economic Perspectives 15 (1) (2001) 55–68.
[40] P. Milgrom, Putting Auction Theory to Work, MIT Press, 2003.
[41] D. Mishra, D.C. Parkes, Ascending price Vickrey auctions for general valuations, Journal of Economic Theory 132 (2007) 335–366.
[42] L. Mous, V. Robu, H.L. Poutre, Using priced options to solve the exposure problem in sequential auctions, in: Tenth International Workshop on Agent Mediated Electronic Commerce (AMEC 2008), Estoril, Portugal, 2008.
[43] C. Ng, D.C. Parkes, M. Seltzer, Strategyproof Computing: Systems infrastructures for self-interested parties, in: 1st Workshop on Economics of Peer-to-Peer Systems, Berkeley, CA, 2003.
[44] C. Ng, D.C. Parkes, M. Seltzer, Virtual Worlds: Fast and Strategyproof Auctions for Dynamic Resource Allocation, in: Proc. 4th ACM Conference on Electronic Commerce (EC'03), 2003.
[45] N. Nisan, Bidding languages for combinatorial auctions, in: P. Cramton, Y. Shoham, R. Steinberg (Eds.), Combinatorial Auctions, MIT Press, 2006. Chapter 9.
[46] A. Ockenfels, 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) (2002) 1093–1103.
[47] R. Parasuraman, C.A. Miller, Trust and etiquette in high-criticality automated systems, Communications of the ACM 47 (4) (2004) 51–55.
[48] D.C. Parkes, Five AI challenges in Strategyproof Computing, in: Proc. Workshop on Autonomic Computing (IJCAI'03), 2003.
[49] D.C. Parkes, Online mechanisms, in: N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani (Eds.), Algorithmic Game Theory, Cambridge University Press, 2007. Chapter 16.
[50] D.C. Parkes, S. Singh, An MDP-based approach to online mechanism design, in: Proc. 17th Annual Conference on Neural Information Processing Systems (NIPS'03), 2003.
[51] D.C. Parkes, S. Singh, D. Yanovsky, Approximately efficient online mechanism design, in: Proc. 18th Annual Conf. on Neural Information Processing Systems (NIPS'04), 2004.
[52] M. Peters, S. Severinov, Internet auctions with many traders, Journal of Economic Theory 130 (1) (2005) 220–245.
[53] D.M. Reeves, M.P. Wellman, J.K. MacKie-Mason, A. Osepayshvili, Exploring bidding strategies for market-based scheduling, Decision Support Systems (2005) 67–85.
[54] P. Resnick, R. Zeckhauser, E. Friedman, K. Kuwabara, Reputation mechanisms, Communications of the ACM 43 (2000) 45–48.
[55] M.H. Rothkopf, R. Engelbrecht-Wiggans, Innovative approaches to competitive mineral leasing, Resources and Energy 14 (1992) 233–248.
[56] M.H. Rothkopf, A. Pekeč, R.M. Harstad, Computationally manageable combinatorial auctions, Management Science 44 (8) (1998) 1131–1147.
[57] T. Sandholm, C. Boutilier, Preference elicitation in combinatorial auctions, in: P. Cramton, Y. Shoham, R. Steinberg (Eds.), Combinatorial Auctions, MIT Press, 2006. Chapter 10.
[58] T.W. Sandholm, An algorithm for optimal winner determination in combinatorial auctions, in: Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99), 1999.
[59] T.W. Sandholm, Expressive commerce and its application to sourcing: How we conducted $35 billion of generalized combinatorial auctions, in: Proc. of the Innovative Applications of Artificial Intelligence conference (IAAI), 2006.
[60] T.W. Sandholm, 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), 1995.
[61] T.W. Sandholm, V.R. Lesser, Leveled commitment contracts and strategic breach, Games and Economic Behavior 35 (2001) 212–270.
[62] A.G. Sanfey, J.K. Rilling, J.A. Aronson, L.E. Nystrom, J.D. Cohen, The neural basis of economic decision-making in the Ultimatum Game, Science 300 (2003) 1755–1758.
[63] I. Segal, N. Nisan, The communication requirements of efficient allocations and supporting prices, Journal of Economic Theory 129 (1) (2006) 192–224.
[64] M.D. Smith, J. Bailey, E. Brynjolfsson, Understanding digital markets: Review and assessment, Understanding the Digital Economy (2000) 99–136.
[65] K. Steiglitz, Snipers, Shills, and Sharks: eBay and Human Behavior, Princeton University Press, 2007.
[66] M. Stryszowska, Late and multiple bidding in competing second price Internet auctions, Working Paper, Fondazione Eni Enrico Mattei (2004).
[67] W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, Journal of Finance 16 (1961) 8–37.
[68] J.T.-Y. Wang, Is last minute bidding bad? Working Paper, University of California, Los Angeles, 2003.
[69] M.P. Wellman, P.R. Wurman, Market-aware agents for a multiagent world, Robotics and Autonomous Systems 24 (1998) 115–125.
[70] M. Yokoo, Y. Sakurai, S. Matsubara, The effect of false-name bids in combinatorial auctions: New Fraud in Internet Auctions, Games and Economic Behavior 46 (1) (2004) 174–188.
[71] R. Zeithammer, An equilibrium model of a dynamic auction marketplace, Working Paper, University of Chicago, 2005.

Please cite this article in press as: A.I. Juda, D.C. Parkes, An options-based solution to the sequential auction problem, Artificial Intelligence (2009), doi:10.1016/j.artint.2009.01.002