

An Options-Based Method to Solve the Composability Problem in Sequential Auctions

Adam I. Juda and David C. Parkes

Division of Engineering and Applied Sciences,
Harvard University,
33 Oxford Street, Cambridge MA 02138
{juda, parkes}@eecs.harvard.edu

Abstract. Current auctions often expose bidding agents to two difficult, yet common, problems. First, bidding agents often have the opportunity to participate in successive auctions selling the same good, with no dominant bidding strategy in any single auction. Second, bidding agents often need to acquire a bundle of goods by bidding in multiple auctions, again with no dominant bidding strategy in any single auction. This paper introduces an options-based infrastructure that respects the autonomy of individual sellers but still enables bidders to utilize a dominant, truthful strategy across multiple auctions.

1 Introduction

Many authors (e.g., [8]) have written about a future in which commerce is mediated by automated trading agents. Yet, we believe that one leading place of resistance is in the lack of optimal bidding strategies for any but the simplest of market designs. Although it is popular to appeal to computational mechanism design [5], and try to design *truthful* auctions to address this problem, it is nevertheless clear that a single truthful mechanism cannot exist for all transactions in which an agent has an interest. Somewhere, at some point, there must be boundaries between mechanisms [20].

This work proposes a new *options-based* market infrastructure, that can enable simple yet optimal bidding strategies, while retaining the seller autonomy that is the defining feature of the most successful of today's electronic markets. Although eBay acts as an intermediary of sorts, eBay does *not* gather up the goods for sale by multiple sellers and run them within a single coordinated event. Rather, eBay is an open environment in which each seller chooses: a) when to bring a good to market; and b) the kind of auction to use. Buyers then pick-and-choose across auctions, before submitting bids and making purchases.

One problem encountered in environments like eBay is the *exposure problem*. People would like to acquire multiple items, but may end up only holding a subset of those items at the end. For example, imagine Alice would like to buy both Peanut Butter (*PB*) and Jelly (*J*), but has to participate in two different auctions in order to acquire both items. Alice may bid enough to win the *PB*, but may subsequently not win the *J*. In doing so, she is left exposed, having spent money in acquiring the *PB*, but unable to derive any benefit without also acquiring the other item she desired. Another problem occurs when *multiple copies* of an item are offered for sale sequentially. For example,

Alice may want PB , and could potentially bid in either a 1 o'clock or 3 o'clock auction. While Alice would prefer to participate in whichever auction will have the lower winning price, she can not determine beforehand which auction that may be, and could end up winning in the wrong auction. In general, we refer to problems including the *exposure problem* and the *multiple copies* problem as the “composability problem,” because they relate to issues with composing strategies across the boundaries of multiple mechanisms [14].

Combinatorial auctions (CAs) [18] can be used to solve the exposure problem by auctioning many different goods simultaneously, allowing bidders to submit desired bundles as their bids (e.g., Alice may submit that she wants both PB “AND” J). However, there are a number of reasons to doubt that CAs can extend to address fully the problem of composability. First, the computational requirements for deriving solutions to combinatorial auctions is worst-case exponential in the number of bids (unless $P = NP$). In an eBay environment, the presence of millions of different goods (with an exponentially larger number of possible *bundles* of goods) makes the design of a single mechanism impractical. Second, CAs do not resolve *temporal issues*, assuming instead that all agents are present at the same time. Third, CAs assume an unrealistic market scope, with one market-maker able to control and bring together all participants into a single market.

Retail stores have customers that face similar strategic problems as the composability problem, and they have devised different policies to alleviate the problems that their customers face. Return policies alleviate the *exposure problem* by allowing customers to return goods at the purchase price. Price matching alleviates the *multiple copies* problem by allowing agents to receive from sellers after product purchase the difference between the price paid for a good and a lower price found elsewhere else for the same good. These two retail policies provide the basis for the scheme proposed in this paper. In particular, we propose an *options-based infrastructure* to address the composability problem.

To participate in the options scheme, a seller must agree to sell an *option* for her good, which will ultimately lead to either a sale of the good, or (if the option is not exercised) going back to the market and offering another option on the same good. The process by which a seller sells options and leaves/returns to market can be made easier via a proxy, whereby a seller may tell a proxy her patience and good for sale, after which the proxy could then sell an option for that good at auction, observe the status of the sold option, and resell another option if the initial option is not exercised and the seller’s patience has yet to expire.

Buyers can collect portfolios of options before deciding which to exercise. We provide buyers with *mandatory* proxy agents, that carefully control the number of outstanding options that each buyer can hold, and yet still follow one of the dominant bidding strategies that an agent could follow if there was no proxy in the system. A buyer reports her value and patience to a proxy agent upon arrival, and then agrees to let her proxy: a) bid in auctions to acquire options; b) exercise options to maximize reported utility once the buyer’s patience is expired. The proxy agents are essential to prevent a buyer from obtaining options on which they have no intention of exercising. The options-based protocol makes truthful and immediate revelation to a proxy a dominant strategy for buyers, whatever the future auction dynamics.

A benefit for sellers as demonstrated through simulation is that the options-based protocol maintains a market even with buyers with complementary values for goods. In comparison, we show that a traditional market can fail, because it quickly becomes difficult for buyers to participate in the market without becoming exposed to partial bundles and losing money. Thus, the options-based scheme has appeal to both buyers and sellers.

1.1 Related Work

The composability problem was previously observed in Wellman & Wurman [20], in the context of a discussion about the boundaries that must inevitably exist between mechanisms. This theme was continued by Parkes [14] and Ng et al. [13] in the context of *strategyproof computing*, in which the goal is to promote the deployment of strategyproof mechanisms within an open and shared infrastructure. The problem has often been identified in the context of simultaneous ascending price auctions, where it is termed the *exposure problem* [4].

Previous work to address the problem has considered two different directions. First, one can change the mechanism to define an expressive bidding language and incentive-compatibility. This is the approach taken in the work on combinatorial auctions (see Rothkopf et al. [18]). Second, one can attempt to provide agents with smarter bidding strategies. This is the approach taken in the work of Boutilier et al. [2], Byde et al. [3], Anthony & Jennings [1], and Reeves et al. [7]. Unfortunately, it seems hard to design artificial agents with equilibrium bidding strategies, even for a simple simultaneous ascending price auction.

Iwasaki et al. [10] have considered *options* in the context of a single, monolithic, auction design to help bidding agents with marginal-increasing values avoid exposure in a multi-unit homogeneous item auction problem. Sandholm & Lesser [19] have considered options in the form of *leveled commitment contracts* for facilitating multi-way recontracting in a completely decentralized market place. Rothkopf & Engelbrecht-Wiggans [17] discuss the advantages associated with the use of options.

Recent work of Porter et al. [16] has considered auctions with uncertainty in an agent's ability to successfully complete a task. As in our work, there is uncertainty for a seller, in their setting due to whether or not a task will be performed. The chief difference is that bidders in their model *know* their fault probabilities, and the authors can design a mechanism around the revelation of this information.

Finally, a recent direction taken in computational mechanism design is that of *online* mechanisms [15] and *online* auctions [11, 9], in which agents can dynamically arrive and depart across time. We leverage a price-based characterization in Hajiaghayi et al. [9] to provide a dominant-strategy equilibrium for buyers within our options-based protocol, creating a decentralized, truthful, option-based implementation of an online combinatorial auction.

2 The Composability Problem

To illustrate the composability problem, consider the following simple example in which a bidder does not have a dominant strategy equilibrium even though the individual auctions in the world are strategyproof.

Example 1. Alice values acquiring PB before Wednesday for \$10. Bob will hold a Vickrey auction for PB on Monday, and then again on Tuesday. In which auction(s) should Alice participate? How much should Alice bid in each auction?

Clearly, Alice no longer has the simple strategy that an individual Vickrey auction offers. By participating in the Monday auction, Alice might win and so would forgo bidding in the Tuesday auction. However, the Tuesday auction may yield a better price. But if Alice decides to skip the Monday auction and only participate in the Tuesday auction, she may not win the Tuesday auction when she could have won the Monday auction, or may have been better off participating in the Monday auction and securing a lower price than the Tuesday auction. Alice has no dominant bidding strategy.

Example 2. Alice values PB and J at \$10. Bob holds a Vickrey auction for PB on Monday. Charlie holds a Vickrey auction for J on Tuesday. How much should Alice bid in each auction?

Again, Alice has no dominant bidding strategy. If Alice wins the PB for \$ y , it is in Alice's best interest to bid \$10 in the second auction, as winning the J derives a value of \$10 for Alice while losing the auction for J garners Alice no value. However, if Alice loses the second auction, she will have a net surplus of $-\$y$, having spent money on PB but receiving no benefit from it. Even if Alice wins the J , it is possible that Alice will have to spend as much as $\$(10 + y)$ to obtain the set, and so again may lose as much as \$ y in surplus.

2.1 The Model

Consider a world with goods, \mathcal{G} , buyers \mathcal{B} and sellers \mathcal{S} . Buyer $i \in \mathcal{B}$ has a value $v_i(L) \geq 0$ for each subset of goods $L \subseteq \mathcal{G}$. Let \mathbb{V} denote the domain of agent valuations. Let $T = \{1, 2, \dots\}$ denote time periods. Each buyer has an arrival time, $a_i \in T$, and departure time $d_i \in T$. (Sometimes we refer to the *patience* of a buyer, which is simply $d_i - a_i$.) Let $D_i \subseteq \mathcal{G}$ denote the goods of interest to agent i , defined as

$$D_i = \{k : \exists L \subset \mathcal{G} \text{ s.t. } v_i(L \cup k) > v_i(L)\} \quad (1)$$

These are the goods that, when taken with some combination of other goods, have non-zero value to the agent. Buyers are indifferent between receiving a bundle of goods at any time before and up to her departure time. Each seller, $j \in \mathcal{S}$, sells a *single* item $k_j \in \mathcal{G}$.¹ All individual auctions in our model will therefore be for a single item. Seller j has an arrival time, $a_j \in T$, and a departure time $d_j \in T$, when they will leave the system if no agent has yet obtained the right to purchase their good.

Agents arrive or leave at the end of each period, but multiple auctions can be sequenced within a period. Let A^t denote this sequence of auctions in period t , each one associated with a single seller. From buyer i 's perspective, let A_i denote the sequence of auctions that occur during the time in which buyer i is in the market. For example, if buyer i arrives in period 1 and departs in period 3, then $A_i = (A^1, A^2, A^3)$.

¹ This is the main limitation in our current model, but a good first-order approximation to eBay style markets.

We require that each individual auction is strategyproof (SP). Following Parkes [14] we term this *local strategyproofness*, to emphasize that it does not imply that a buyer has a dominant strategy when bidding across multiple such auctions. The utility to a buyer, given goods L and price p , is defined as $u_i(L, p) = v_i(L) - p$. An individual auction can be defined in terms of an *allocation rule*, $x(b)$ and a *payment rule* $p(b)$, given bids $b = (b_1, \dots, b_N)$. Bid b_i reports a value $b_i(L)$ for each bundle of goods L in the auction, and need not be truthful. In particular, agent i receives $x_i(b) \subseteq \mathcal{G}$ goods and makes payment $p_i(b) \in \mathbb{R}$.

Definition 1 (locally strategyproof). *An auction A , defined as (x, p) , is (locally) strategyproof when its bidding language is expressive given valuations \mathbb{V} , and given the goods available in the auction, and when truthful bidding is a dominant-strategy for an agent that can only bid in this one auction.*

Formally, truthful bidding, $b_i = v_i$, is a *dominant* bidding strategy when $v_i(x_i(v_i, b_{-i})) - p_i(v_i, b_{-i}) \geq v_i(x_i(\hat{v}_i, b_{-i})) - p_i(\hat{v}_i, b_{-i})$, for all bids b_{-i} from other agents, and for all $\hat{v}_i \neq v_i$. For instance, a single-item Vickrey auction is locally-SP for all agents that will bid only in that auction.

When facing a sequence of auctions, a bidding strategy, b_i for buyer i defines the bid that the agent will make in each auction, and can be contingent on: i) her own value; ii) her beliefs about other agents; iii) the outcomes and feedback from earlier auctions.

Definition 2 (The composability problem). *The composability problem exists for an agent facing a sequence of auctions A_i , when each auction in A_i is locally-SP, but the agent does not have a dominant bidding strategy across the sequence of auctions.*

In fact, the composability problem exists more often than not! In what follows, we assume that all goods in an agent's valuation function are available in A_i , and that each auction is locally-SP. All proofs are omitted in the interest of space. First, consider a single-minded buyer, and let (B_i, w_i) denote that the buyer demands bundle B_i at price w_i . Refer to auction $j \in A_i$ as an *interesting auction* when the good $k_j \in B_i$.

Proposition 1. *The composability problem exists for a single-minded buyer whenever there are two or more interesting auctions.*

Proof omitted for space. The effect of multiple auctions is that the agent must anticipate the level of competition, and prices, in future auctions when deciding how to bid. For instance, an agent that wants a single item and faces a sequence of Vickrey auctions does not have a dominant bidding strategy, but would prefer to bid in the auction with the lowest second price.

Next, we can consider an agent that demands multiple disjoint bundles $B_{im} = \{B_{i1}, B_{i2}, \dots, B_{iM}\}$, with an additive value w_{im} for each bundle.

Proposition 2. *An agent with additive values across bundles faces the composability problem whenever there is at least one bundle B_{im} for which the composability problem exists.*

Proof omitted for space. We can also consider an agent with a general valuation that cannot be expressed as an additive value across disjoint bundles, which precludes a single-minded agent.

Proposition 3. *An agent i with a general valuation faces the composability problem whenever it faces two or more interesting auctions.*

Proof omitted for space. Here, there must exist bundles that are either substitutes or complements. If the bundles are substitutes, an agent faces the problem of determining which bundle to pursue (analogous to the problem an agent faces when the same single item is sold at multiple auctions). If the bundles are complements, an agent can face the exposure problem (analogous to when a single bundle contains multiple items).

3 The Opportunity to Use Options

An option is a right to acquire a good at a certain price, called the *exercise price*. For instance, Alice may obtain from Bob the right to buy PB from him at an exercise price of \$3. What makes options unique is that the *right* to purchase a good at an exercise price does not imply the *obligation* to purchase a good at an exercise price. Therefore, when Alice obtains an option from Bob, Bob is not guaranteed that Alice will actually exercise the option at the exercise price and obtain the good. This flexibility makes options useful in addressing the composability problem. Buyers can put together a collection of options on goods, and then decide whether to exercise each option.

Options are typically sold, obtained at a price called the *option price*. However, options obtained at a non-zero option price can not generally support a simple dominant bidding strategy, as an agent must compute the expected value of an option [6] to justify the cost. This computation requires a model of the future, which in our setting requires a model of the bidding strategies and the values of other bidders. This is the very reasoning that we are trying to avoid by introducing options! Instead, we consider *costless options*, where the option price is zero. This will require some care.

The basic problem arises because agents are always (weakly) better off with an option than without an option, whatever its exercise price, because an agent can always choose for free not to exercise an option won. Therefore, an agent would be interested in obtaining a costless option at *any* exercise price (including infinity), subsequently choosing to exercise the option only if doing so would result in a gain of surplus. However, multiple bidders pursuing options with no intention of exercising them could cause market efficiency to unravel. We address this issue through mandatory *proxy agents*, which intermediate between buyers and the market.

4 Auctions for Options

In our scheme, sellers run an auction for costless options on goods, and buyers bid through mandatory proxy agents. These proxy agents are critical to addressing the potential for an inefficient allocation of options through hoarding. Proxy agents, coupled with auctions for options, make it a buyer's dominant strategy to truthfully reveal her valuation and patience. Proxies follow a dominant bidding strategy for a buyer (by bidding at a value high enough that no higher bid could make the agent *strictly* better off), but restrict a buyer from pursuing options on which it is *indifferent*, such as a second option for a good when only one instance of the good is desired, or an option with an exercise price that could never be exercised for positive surplus.

We now define the two main elements of our market:

Seller Auctions. Each seller j sells a costless option in a Vickrey auction. The option is issued to the highest bidder, with an exercise price equal to the second-highest bid, and is set to expire at the end of the buyer’s patience. Sellers also agree, by joining the market, to allow the proxy representing a winning buyer to *adjust downwards* the exercise price if the proxy discovers that it could have achieved a better price by waiting to bid in a later auction for the same good (i.e., sellers agree to price match their competitors). Sellers can run additional auctions if the options are returned.

Proxy Agents. Each buyer i must submit to a proxy an expressive language bid, \hat{v}_i , reporting values for desired bundles. Each buyer also must submit a departure period, \hat{d}_i , to the proxy. The proxy computes the *maximal value* for each item k desired by i as $v_i^{\max}(k) = \max_{L \supseteq k} \hat{v}_i(L)$.² The proxy bids this price in any auction for item k , until it holds an option on this item. At that point, the proxy tracks future auctions on that item, determines what the world would look like if it had delayed its entry into the market until that later auction, and will reduce the exercise price of the option it has if it discovers that it could have secured a lower price by waiting to bid in that later auction.³ The proxy determines this information by asking each future auction to report the identities (can be pseudonymous) of the winner and second-highest winner, together with their bid prices. The identities are necessary because they are used by the agent when creating its view of the world had it decided to delay its entry.⁴ Finally, at the end of the buyer’s reported patience \hat{d}_i , the proxy exercises options to maximize reported value \hat{v}_i , solving $\max_L \hat{v}_i(L) - \sum_{k \in L} p_{\text{opt}}(k)$, where $p_{\text{opt}}(k)$ is the option’s exercise price, and $p_{\text{opt}}(k) = \infty$ if the proxy does not hold an option. All other options are returned. No options are exercised when no combination of options are priced below a buyer’s reported value.

The proxy agent forces a link between the valuation function \hat{v}_i used to acquire options and the valuation \hat{v}_i used to exercise options. Without this, agents could indifferently acquire options with exercise prices too high to ever be exercised. The proxy

² While sellers of the same item type k' may not have different reserve prices for their goods (due to potential conflicts in being able to price match), sellers may agree (or be required) to have a universal reserve price for each item type, $rp(k')$. In such a scenario, bidding agents can incorporate this information into their bids for multi-item bundles because it provides a tighter lower-bound on the price; specifically, $v_i^{\max}(k) = \max_{L \supseteq k} \left(\hat{v}_i(L) - \sum_{k' \in L, k' \neq k} rp(k') \right)$.

³ However, the proxy does *not* at any point acquire a second option for the good. Rather, it retains the single option it has been holding, but reduces its exercise price to the later price.

⁴ In particular, the proxy maintains a *candidate* agent, cand_k^i , for each item k on which it holds an option. Agent cand_k^i is the agent still present in the market that is currently not allocated an option for k , but would have been by now had the proxy delayed its entry. There may be no such candidate agent if the displaced winner leaves the market without winning in a subsequent auction, at which point the state of the market looking forward is unaffected by i winning its option. Initially, cand_k^i is set to the highest outside bidder in the auction in which proxy i wins an option for k . In subsequent auctions for k : if cand_k^i wins, the exercise price for the option held by i is adjusted to the minimal of its current value and the second-highest price in this new auction, and the second-highest bidder in this new auction becoming the new candidate; else, the proxy’s price is adjusted to the minimal of its current price and the highest bid price in this new auction.

agent also ensures that no buyer can hold more than one option on each good, and can hold options on no goods outside its demand set. Without this, agents could indifferently obtain options that they have no intention of exercising. These two properties help to provide a well-functioning market.

Example 3. Alice desires PB and J for \$30. Bob desires PB for \$5. Charlie desires J for \$10. All agents have a patience of 2 days. On day one, a J auction is held, where Alice’s proxy bids \$30 and Charlie’s bids \$10. Alice wins an option to purchase the J for \$10. On day two, a PB auction is held, where Alice’s proxy bids \$30 and Bob’s bids \$5. Alice wins an option to purchase the PB for \$5. At the end of the second day, Alice’s proxy holds an option to buy PB for \$5 and an option to buy J for \$10, and so exercises both options, spending a total of \$15 to acquire her entire desired bundle.

Example 4. Alice desires PB for \$20. Bob desires PB for \$10. Charlie desires PB for \$5. All agents have a patience of 2 days. On day one, an auction is held for Peanut Butter where each agents’ proxy bids their value, and Alice’s proxy wins an option to buy PB for \$10, and Alice’s proxy notes that Alice prevented Bob from winning the option it now holds. On day two, another auction for PB is held where only Bob and Charlie’s proxies bid. Bob’s proxy wins an option to buy PB for \$5. Alice’s proxy looks at its notes and observes that had Alice delayed her entry until now, Bob would not be bidding. Therefore, Alice would have won today’s auction at Charlie’s bid price of \$5, and Alice’s proxy adjusts Alice’s option price down to \$5. At the end of the second day, both Alice and Bob hold options to buy PB for \$5, and so both proxies exercise their options, each spending \$5.

Example 5. Alice desires a red hat for \$20 “XOR” a blue hat for \$10 (i.e., Alice’s value is \$20 for a red hat, \$10 for a blue hat, and \$20 for both). Bob desires a red hat for \$15. Charlie desires a blue hat for \$30. On day one, a red hat auction is held where Alice’s proxy bids \$20 and Bob’s proxy bids \$15, resulting in Alice winning an option for the red hat with an exercise price of \$15. On day two, a blue hat auction is held where Alice’s proxy bids \$10 and Charlie’s proxy bids \$30, resulting in Charlie winning an option for the blue hat with an exercise price of \$10. At the end of day two, Alice’s proxy exercises her red hat option and Charlie’s proxy exercises his blue hat option.

4.1 Truthful Bidding to Proxy Agent

What remains to be shown is that it is a dominant strategy for the buyer to truthfully reveal her value and arrival and departure time to her proxy agent. The proof builds on the price-based characterization of time-strategyproof auctions in Hajiaghayi et al. [9].

Lemma 1. *An online unit-demand auction is time and value strategyproof when:*

- *it constructs a price function for the agent over time that is independent of the type reported by the agent,*
- *it allocates the good to the agent at the minimal price during period $[a_i, d_i]$ in the market, in a period no earlier than this minimal price, and only when the agent’s value is higher than this minimal price.*

Theorem 1. *It is a dominant strategy for a buyer to truthfully reveal her valuation function and patience to her proxy agent in the options-based market.*

Proof. (sketch) The options scheme constructs an agent-independent price schedule, $p_i^t(k)$, for each item k (in period t), defined as the highest bid received among those agents not holding an option at time t , not including i herself and not including any agents that would have already won options had i never entered the system. The proxy agent holds an option for item k at time t at the *minimal* price from \hat{a}_i to t , whenever this price is less than the maximal value the agent could have for the item (given possible future prices). Overstating the value on a bundle can lead to the proxy holding an option on some item, k , at some price greater than it would ever want to pay. Understating the value can lead to the proxy missing a useful option on some item, k , and will not reduce the price otherwise. Strategies that misstate arrival and departure are not useful because reporting a later arrival or earlier departure can forfeit opportunities, while reporting a later departure, $\hat{d}_i > d_i$, is not useful because the agent will not receive its goods until after d_i . Thus, there is no useful manipulation of the final options on individual items, and finally the proxy makes a purchasing decision by looking at prices on options and exercising those that maximize reported utility. ■

5 Experiments

Up to this point we have focused solely on buyers. However, sellers can also benefit because the options-based scheme fixes the market failure that exists when buyers have complex values but face a sequence of auctions. The experimental results in this section demonstrate that there are many scenarios in which the average buyer surplus in a market without options is *negative*. In such a world, buyers would not enter the market to begin with (such a decision is not individually-rational) and there would be no market formation.

We simulate a simple market to better understand the economic effect of the options scheme, for both buyers and sellers. We construct values for buyers according to a *quadratic method* [12], which is parameterized with (M, γ) . Each buyer receives value on M bundles, and each bundle contains γ distinct goods. The value of one bundle, B_{im} for buyer i , is determined by first choosing γ (distinct) items uniformly at random and assigning each item k a value $v_k \sim U[0, 1]$. The value on the bundle is defined as $w_{im} = \sum_{k \in B_{im}} v_k + \sum_{k' \in B_{im}, k' \neq k} v_k v_{k'}$. Each seller sells a single item, chosen uniformly at random from the set of all goods. We choose to model an *identical* reserve price for all goods and for all sellers, which is set to $rp(k) = 0.5$ for each good k unless stated otherwise.⁵

Buyer patience is set to 50, while seller patience is set to 100. We vary the buyer entry-rate and seller entry-rate to model different levels of supply and demand. We compare the options-based market with a market in which there is a sequence of Vickrey auctions for traditional goods. To model sell-side auctions in each round we choose to run f^t auctions in each period t . Upon arrival, sellers wait in a queue for their auction to be scheduled, on a first-come first-served basis. The rate, f^t , is adjusted to keep the

⁵ Even though all sellers use the same reserve price, we assume that the buyers do not know this, and do *not* pursue the alternate bidding strategy introduced in Footnote 2.

wait time, defined as the time that a seller needs to wait to have her auction scheduled, below 5 periods.⁶ In the options world, a seller returns to the end of the queue if there are no bidders in her auction, or if the winner returns her option. In the non-options world, a seller only returns to the queue if she fails to sell her item.

Finally, we need a model of buyer strategies. In the options world, we assume that each buyer reports her true value to her proxy immediately upon arrival. In the non-options world, we need to adopt a bidding strategy for buyers, to provide a meaningful comparison with the performance of the options world. Ideally we would adopt an equilibrium bidding strategy, but this analysis is not available for such a complex game. Instead, we adopt a “sunk-aware” bidding strategy, following the ideas in Reeves et al. [7]. At any point in time an agent has purchased goods L^t , for price $\sum_{k \in L^t} p(k)$, where $p(k)$ was the price paid in the Vickrey auction for good k . Consider some item $k' \notin L^t$. The agent estimates her value for k' as

$$\hat{v}_i^t(k') = \max_{L|k' \in L} \left[v_i(L) - \alpha \sum_{k \in L^t} p(k) - \sum_{k \in L, k \notin L^t, k \neq k'} \hat{p}_\beta^t(k) \right] \quad (2)$$

where $\hat{p}_\beta^t(k)$ is defined as the average price for item k in the last $\beta > 1$ auctions for k (we assume the agent has access to this information).⁷ In the event of an auction for item k' , the agent then bids this estimated value, $\hat{v}_i^t(k')$.

Parameter $\alpha \in [0, 1]$ determines how sunk-cost aware the bidder will act. With $\alpha = 0$ the agent ignores the sunk cost and continues to bid aggressively. With $\alpha = 1$ the agent considers the cost of items already purchased, as though it is deciding over again whether to buy those items and the new items. Thus, for higher α the agent is more conservative, and bids less aggressively. We model buyers in the non-options world as leaving the market as soon as they have purchased a bundle for which they have positive value, or at time d_i (whichever occurs first). The sunk-cost parameter α was selected so as to maximize buyer performance in the market for each (M, γ) and each level of buyer entry-rate. *Buyer surplus* is measured as the average *value-normalized* surplus upon exit from the market, with value-normalized surplus for a buyer that purchases bundle B_{im} at price p defined as $(v_{im} - p)/v_{im}$, for true value v_{im} . Note that this can go negative, when a buyer pays more than her value for a bundle. We normalize in this way to remove dependencies on absolute values of goods in our empiric analysis. In the non-options world, when a buyer can fail to put together a complete bundle, we substitute $-p/V$, where $V = \max_L v_i(L)$, i.e. the value of the *most-preferred* bundle. *Seller surplus* is measured as the ratio of total revenue generated by all sellers divided by the total value of goods allocated to buyers. *Losers* are the percentage of buyers that leave the market with negative surplus. In some markets without options buyers can lose money, even when bidding conservatively. In this case, we also calculate the *adjusted-seller* surplus, by factoring out any revenue that sellers were achieving from buyers that were losing money (which we would not expect in a sustainable equilibrium).⁸

⁶ In practice, we set $f^t = \lceil \frac{N^t}{5} \rceil$, where N^t is the number of sellers in the queue in period t .

⁷ β was set to 25 in experimentation.

⁸ This adjustment is deemed conservative, with a more realistic revenue for sellers in these scenarios expected to be lower. If the negative surplus bidders had never entered the market,

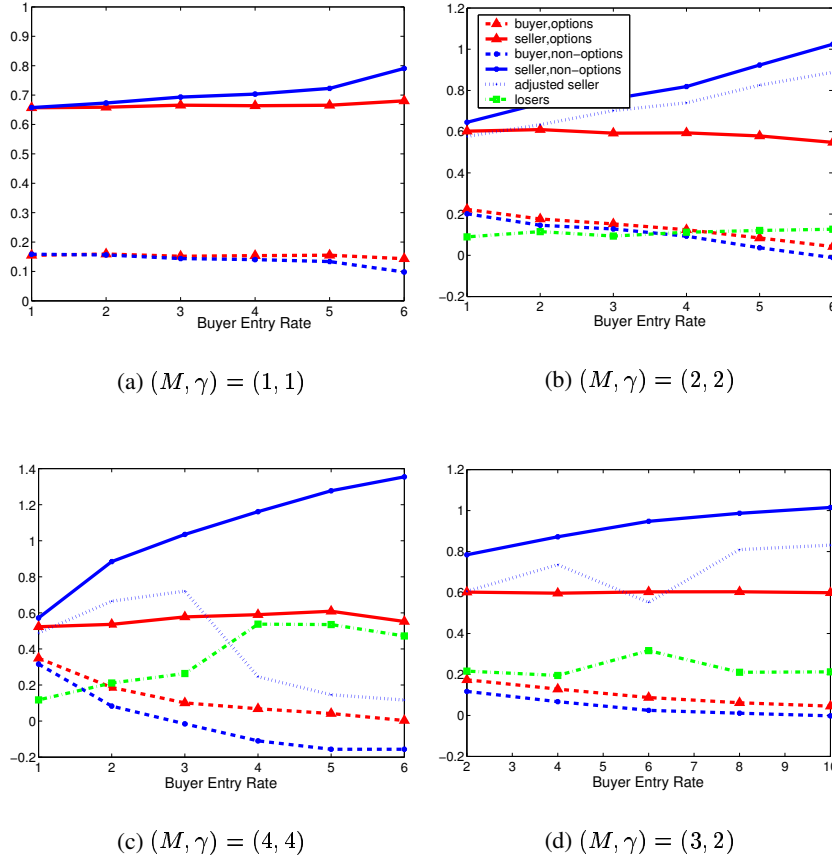


Fig. 1. Buyer and seller surplus vs. buyer entry rate, for different valuations parameterized with (M, γ) . We also plot the percentage of buyers that are *losers*. Subplots (a)–(c) hold the *seller entry rate* fixed, whereas the *seller entry rate* in (d) is scaled with *buyer entry rate* at a 2:1 ratio.

Figure 1 illustrates buyer and seller surplus against an increasing buyer entry-rate, with each subplot dedicated to a different structure (M, γ) for buyer valuations. We consider values of $(M, \gamma) \in \{(1, 1), (2, 2), (4, 4)\}$ in subplots (a), (b) and (c) respectively, with the seller entry-rate set to 3, 6 and 12 in each scenario (this increases supply in line with increased buyer demand as the number of items M demanded in each bundle increases). Figure 1 (a) demonstrates that the options and non-options world produce very similar results when the population has the most simple of valuations, with seller surplus increasing and buyer surplus decreasing as demand increases. Similar results were experimentally confirmed for all $(M, \gamma) = \{(2, 1), (3, 1), (4, 1), (5, 1)\}$.

In Figure 1 (b) and (c) the non-options world “breaks” when demand gets too high because the composability problem becomes more challenging. The average buyer in-

not only would their income be lost to sellers, but also the prices paid by those people who remained would be lower due to the decreased competition.

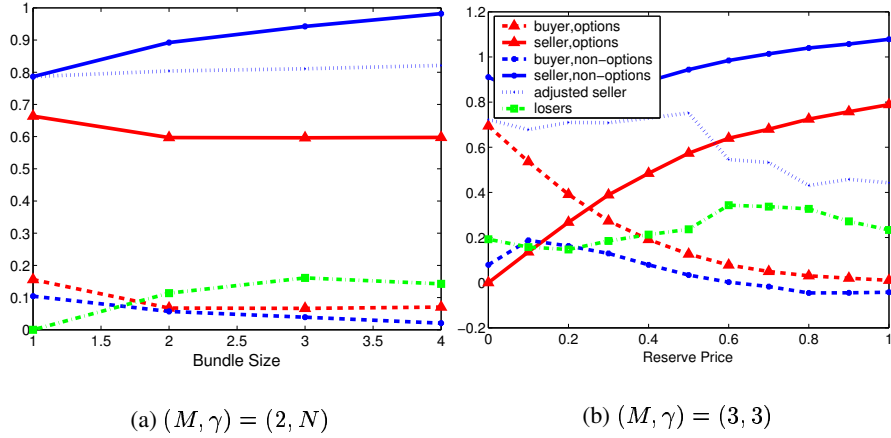


Fig. 2. Buyer and seller surplus, for different valuations parameterized with (M, γ) . Subplot (a) shows the effects of increasing the number of items in the bundles in a valuation. Subplot (b) shows the effects of varying the reserve price.

curs negative surplus and one would reasonably expect that buyers would not enter this market in the first place. On the other hand, buyer surplus in the options world remains positive, indicating that buyers would continue to enter the market. This suggests the existence of scenarios in which introducing options can create new markets. Further evidence for market breakdown in the non-options world can be found by considering the *Losers* rates. Figure 1 (c) shows that the *Losers* rates is near 0.5 when there is high demand, indicating that nearly half of the agents who are entering the market are losing surplus upon exit.

In Figure 1 (d), we consider $(M, \gamma) = (3, 2)$, and scale the seller entry-rate continuously as we scale the buyer entry-rate, keeping the seller entry at twice the buyer entry. Whereas scaling the world degrades the buyer surplus in the non-options world to the point of being negative, seller surplus in the options world is steady and accompanied by positive buyer surplus.

Figure 2 (a) fixes $M = 2$ and increases the number of items in each bundle. We scale the buyer entry-rate from 16 to 8 to 5 to 4, while the seller entry-rate is fixed at 16, to keep supply and demand in the same proportion. The exposure problem is present in this scenario, and indeed we see a significant drop in buyer surplus and a rise in the percentage of buyers with negative surplus in the non-options world. Figure 2 (b) illustrates the effect of changing the reserve price on all items. Buyer entry-rate is 3, seller entry-rate is 9, and buyers demand 3 bundles, each with 3 items. A higher reserve price in the non-options world drives buyer surplus negative, and results in market failure. On the other hand, sellers in the options world are able to raise reserve prices to increase their surplus, and buyers can still manage to obtain positive surplus from the market even when reserve prices are set very high. Noteworthy, although the seller surplus is increasing in Figure 1, one should appreciate that this is a *relative* metric. In fact, we observe that the *total* seller surplus tends to decrease in the options

world as buyer entry-rate grows (even though the normalized surplus remains quite flat). Less buyers complete their desired bundles and less buyers eventually exercise their options. Of course, we believe this is preferable to the complete market failure in the non-options world. However, while markets continue to form with options, trades are more infrequent as buy-side competition increases— which suggests that it would be interesting to explore additional methods, such as a prequalification stage, the use of stronger reserve prices, or the “throttling” of buyer entry-rate and pooling into separate markets. We reserve these topics for future study.

6 Conclusion

We introduced an options-based auction protocol to address the composability problem that exists when buyers with complex values must bid in sequences of simple auctions. Our approach combines costless options with proxy agents, which acquire, maintain, and exercise options on the agent’s behalf and best interest. Simple trading agents have dominant bidding strategies in our options-based market, even though the markets remain fundamentally disintermediated. We believe that options-based markets may provide an interesting new class of market designs for eBay-like electronic markets.

Future work should aim to better situate this work within the context of the theory of strategyproof *online* auctions. Future work may also address and resolve the strategic problems facing sellers in this work. While it is not a dominant strategy for sellers to try and keep prices on their goods artificially high (as doing so may prevent options from being exercised if the prices maintained are at a prohibitively high level), it is true that straightforward truthful behavior may not always be in the best interest of sellers in the current model. Furthermore, an investigation of the role of false-name behavior [21] should also yield interesting results. While buyers do not want to engage in false-name behavior as multiple buyers, and sellers do not wish to engage in false-name behavior as multiple sellers, we believe there are manipulations for buyers pretending to be sellers and for sellers pretending to be buyers.

Additionally, future work on the empirical aspects of this project should aim to utilize better benchmarks when analyzing the model, including the use of real data. In particular, there are three areas where real data could be particularly helpful to this model. First, we believe there is ample opportunity for further exploration as to modeling the arrival of sellers and timing of auctions in this setting, perhaps using data from eBay as a foundation. Second, the high rates of buyers that are losing surplus in our simulation of buyers in the non-options model when demand is high is cause to believe that agents may follow a different bidding strategy than the one assumed here. Real world data can be of great assistance in helping to empirically determine what those strategies might be. Third, real world data can also help in developing accurate valuation models for some set of niche goods in an existing market.

Acknowledgments

Useful comments and suggestions were received during seminar presentations of earlier versions of this work in DEAS, HBS, and AMEC VI as part of AAMAS ’04, as well as from Michael Rothkopf and Loizos Michael. This work is supported in part by NSF grant IIS-0238147.

References

1. P. Anthony and N. R. Jennings. Developing a bidding agent for multiple heterogeneous auctions. *ACM Trans. On Internet Technology*, 2003.
2. C. Boutilier, M. Goldszmidt, and B. Sabata. Sequential auctions for the allocation of resources with complementarities. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 527–534, 1999.
3. A. Byde, C. Preist, and N. R. Jennings. Decision procedures for multiple auctions. In *Proc. 1st Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS-02)*, 2002.
4. M. M. Bykowsky, R. J. Cull, and J. O. Ledyard. Mutually destructive bidding: The FCC auction design problem. *J. of Regulatory Economics*, 2000. .
5. R. K. Dash, N. R. Jennings, and D. C. Parkes. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems*, pages 40–47, November 2003. Special Issue on Agents and Markets.
6. A. K. Dixit and R. S. Pindyck. *Investment under Uncertainty*. Princeton University Press, 1994.
7. D.M.Reeves, M.P.Wellman, J. K.MacKie-Mason, and A.Osepayshvili. Exploring bidding strategies for market-based scheduling. *Decision Support Systems*, 2004. Special issue dedicated to the Fourth ACM Conference on Electronic Commerce (EC'03).
8. A. Greenwald and J. O. Kephart. Shopbots and pricebots. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 506–511, 1999.
9. M. T. Hajiaghayi, R. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. In *Proc. ACM Conf. on Electronic Commerce*, pages 71–80, 2004.
10. A. Iwasaki, M. Yokoo, and K. Terada. A robust open ascending-price multi-unit auction protocol against false-name bids. *Decision Support Systems*, 2004. Special issue dedicated to the Fourth ACM Conference on Electronic Commerce (EC'03).
11. R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC-00)*, 2000.
12. K. Leyton-Brown, M. Pearson, and Y. Shoham. Towards a universal test suite for combinatorial auctions. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC-00)*, 2000.
13. C. Ng, D. C. Parkes, and M. Seltzer. Strategyproof Computing: Systems infrastructures for self-interested parties. In *1st Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, June 2003.
14. D. C. Parkes. Five AI challenges in Strategyproof Computing. In *Proc. IJCAI'03 Workshop on Autonomic Computing*, Acapulco, Mexico, 2003.
15. D. C. Parkes and S. Singh. An MDP-based approach to Online Mechanism Design. In *Proc. 17th Annual Conf. on Neural Information Processing Systems (NIPS'03)*, 2003.
16. R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz. Mechanism design with execution uncertainty. In *Proc. 18th Conf. on Uncertainty in Artificial Intelligence (UAI'02)*, 2002.
17. M. H. Rothkopf and R. Engelbrecht-Wiggans. Innovative approaches to competitive mineral leasing. *Resources and Energy*, 14:233–248, 1992.
18. M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.
19. T. Sandholm and V. Lesser. Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35:212–270, 2001.
20. M. P. Wellman and P. R. Wurman. Market-aware agents for a multiagent world. *Robotics and Autonomous Systems*, 24(3–4):115–125, 1998. Earlier version in Proc. 8th Eur. Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW-97).
21. M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: New Fraud in Internet Auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.