



On Revenue-Optimal Dynamic Auctions for Bidders with Interdependent Values

Citation

Constantin, Florin and David C. Parkes. 2009. On revenue-optimal dynamic auctions for bidders with interdependent values. In Agent-mediated electronic commerce and trading agent design and analysis, ed. J. Collins, 1-15. Berlin, Heidelberg: Springer. Previously published in Lecture Notes in Business Information Processing 13: 1-15.

Published Version

doi:10.1007/978-3-540-88713-3_1

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:4039773>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

On Revenue-Optimal Dynamic Auctions for Bidders with Interdependent Values

Florin Constantin and David C. Parkes

School of Engineering and Applied Sciences, Harvard University
33 Oxford St, Cambridge, MA 02138 USA
{florin,parkes}@eecs.harvard.edu

Abstract. In a dynamic market, being able to update one’s value based on information available to *other* bidders currently in the market can be critical to having profitable transactions. This is the model of *interdependent values* (IDV): a bidder’s value can explicitly depend on the private information of other bidders. In this paper we present preliminary results about the revenue properties of *dynamic* auctions for IDV bidders. We adopt a computational approach to design single-item revenue-optimal dynamic auctions with known arrivals and departures but (private) signals that arrive online. In leveraging a characterization of truthful auctions, we present a mixed-integer programming formulation of the design problem. Although a discretization is imposed on bidder signals the solution is a mechanism applicable to continuous signals. The formulation size grows exponentially in the dependence of bidders’ values on other bidders’ signals. We highlight general properties of revenue-optimal dynamic auctions in a simple parameterized example and study the sensitivity of prices and revenue to model parameters.

1 Introduction

Suppose that every day a web content provider, such as the New York Times (NYT), auctions the right to have a banner ad appear next to the lead sports story the following day. The NYT has contracted with various advertisers, where the contract defines a fixed bidding interval wherein the advertiser can choose to bid for the right at the same time each day (the “arrival” time of the bid) and is guaranteed a response by some subsequent time (the “departure” time of the bid.) Only the bid of the advertiser, and indeed whether or not the advertiser will choose to bid, is uncertain. The NYT must complete the auction by 4am each morning, where upon the auction for the subsequent day commences.

It is reasonable to assume that advertisers will have a hard time known the value of having their own banner ad appear because this depends, in part, on the breaking news stories that day and on the associated user demographics. For this reason, an advertiser may naturally be interested in other bidders’ information regarding the value of the banner ad. The *interdependent values* (IDV) model from auction theory nicely captures a model in which each bidder has private information (his signal), but a value that depends on his own signal as well as the signals of other bidders.

Our model concerns interdependent values in a dynamic, known-interval setting. As in any domain where bidders report private information to the auction-

eer, a basic requirement is *incentive compatibility* (IC, also known as truthfulness): bidders must be best off when reporting truthfully.¹ We first consider a naive generalization of the optimal *static* and IC auction for interdependent bidders [3] and we point out that the solution obtained fails to satisfy IC constraints. The reason is that, in retrospect, a bidder i will regret reporting truthfully in the following scenario: i can misreport his signal changing the value or price of an earlier bidder h such that h is now precluded from winning and i is going to win (maybe because the future signals turn out to be “favorable” for i to misreport).

To determine the optimal, revenue-maximizing auction in our dynamic IDV environment we adopt a mixed-integer programming (MIP) formulation and follow the framework of automated mechanism design [4], which builds on the heritage of the Myerson [10] program. For practical scalability we require that the interdependencies between bidders is of bounded degree, that the designer is able to constrain the number of signals that must be propagated from earlier periods into defining the price of bidders in later periods, and that a coarse discretization of signals can be tolerated. The formulation is illustrated in a simple, three bidder scenario. Based on this formulation we can compare the revenue from this, IC formulation with the one obtained using the naive generalization and the one obtained using a clairvoyant approach.

Related work Branco [3] studies *static* revenue-optimal auctions and shows that under a certain regularity condition, an asymmetric critical signal-based auction is optimal (we review this auction in Subsection 3.1). The regularity condition is satisfied if bidders’ valuations are increasing and concave in their own signal and the signals’ distribution has a non-decreasing hazard rate.

Aoyagi [1] investigates optimal pricing schemes in dynamic settings with IDV bidders. Aoyagi is not concerned with incentive compatibility and adopts an indirect approach in which a bidder infers his value from the decisions of other bidders. If a previous bidder j accepted (respectively rejected), then a current bidder i ’s estimation of j ’s signal will increase (respectively decrease), leading to a corresponding change in i ’s *value*. Aoyagi shows that for any simultaneous selling scheme, there exists a sequential one with at least as high a revenue.

Hajiaghayi et al. [6, 7] provide competitive mechanisms for selling one or more goods in an dynamic environment, but they model bidders’ values as private. See Parkes [12] for a recent survey on mechanism design in dynamic environments. The strategy of finding optimal-revenue mechanisms through search is in the spirit of *automated mechanism design* [4]. However, rather than impose IC constraints directly our formulation amounts to an *informed search*, since we search only for critical signals that support a truthful allocation policy.

2 Preliminaries

We now describe our model of interdependent values in dynamic settings. Throughout we assume that there is a single indivisible good for sale and the auction

¹ The revelation principle tells us that it is without loss to focus on IC auctions, since any outcome that can be implemented in the equilibrium of some auction can also be achieved in the equilibrium of an IC auction.

will end as soon as the good is sold. Of course, this might be repeated every day (see example in the introduction).

Each bidder i is active in an *interval known* to the auctioneer, delimited by his *arrival* a_i and *departure* d_i . Bidder i 's private information about the item is encoded in his *signal* s_i . Each bidder's value can depend on the collection of signals reported in the market before his departure. Misreports are therefore significantly restricted in our model since a bidder can only misreport his signal. The auctioneer has uncertainty about s_i and models each s_i as an independent draw from a distribution on non-negative values with cumulative distribution function (cdf) F_i and probability density function (pdf), f_i . We consider only deterministic auction rules and assume for simplicity that all bidders have disjoint departures.²

Let us define now the semantics of interdependent values in environments where the set of signals is dynamic. By s_{-i} we will denote the signals of bidders that have arrived before i 's *departure*. Given this, then a bidder's value $v_i(s_i, s_{-i}) \geq 0$ is well defined in all periods no earlier than the last period $t \in [a_i, d_i]$ for which another bidder arrives. Valuation function v_i is known to the auctioneer but signals are private to bidders. In the standard online, private-value auctions model then $v_i(s_i, s_{-i}) = s_i$ and a bidder's valuation is known, with certainty, throughout his presence in the auction. The online environment also imposes an additional constraint: an bidder's value is zero for an allocation decision in a period t after his departure. The departure period can be conceptualized as the period by which an bidder must receive a decision.

Without loss of generality we can focus on auction protocols that sell to a bidder upon his departure. Not only does this ensure that an bidder's own value is known but this allows the auctioneer to gain maximal information about other demand in the market. We assume that v_i is differentiable with respect to s_i and: (1) *v-monotonicity*: $v_i(s_i^+, s_{-i}) \geq v_i(s_i, s_{-i}) \forall i, \forall s_{-i}, \forall s_i^+ \geq s_i$. That is, a higher private signal cannot result in a lower value for the item; (2) the *single crossing condition* (SCC): an infinitesimal change in bidder i 's private signal influences i 's value *more* than it influences the value of j if i 's value is equal to j 's and at least as high as the values of the other bidders. Any non-trivial incentive compatible auction in static IDV environments must satisfy (1) and (2) (see [9]).

Let bidder i 's true signal be s_i and his reported signal be s'_i . A dynamic auction defines an allocation rule $q_i(s'_i, s_{-i}) \in \{0, 1\}$ to indicate whether or not bidder i is allocated the item, and a payment rule $p_i(s'_i, s_{-i}) \geq 0$ to specify the payment made by bidder i (the time here is implicit: d_i). In a dynamic environment these must be online computable, i.e. they may not depend on signals not yet observed. Bidders are modeled with *quasilinear* utilities: the utility of bidder i with type s_i when reporting s'_i is $q_i(s'_i, s_{-i})v_i(s_i, s_{-i}) - p_i(s'_i, s_{-i})$. That is, a bidder i 's utility is his value minus price, where i 's value for the allocation is the value given all signals (aggregated by v_i) if i wins the item and 0 otherwise.

² If two bidders depart in the same period then they are effectively taking part in an one-shot IDV auction.

The auctioneer can take a good decision only if bidders' reports are accurate; therefore we require that bidders are motivated to truthfully report their signals:

Definition 1. *An auction is incentive compatible (IC) if, when other bidders report their true types, the ex post utility of any bidder is maximized if he reports his true type as well (i.e. truthful reporting is an ex post Nash equilibrium).*

This is *ex post* IC, meaning that a bidder's best strategy is to truthfully report his signal whatever the signals of other bidders. As is standard for interdependent value auctions, this is however weaker than the dominant-strategy equilibrium property that can be achieved in private-value settings. Intuitively, an bidder's best strategy critically relies on the auctioneer having the correct view of his value and this depends on the other bidders reporting truthful signals. We also require that the utility of any truthful bidder be non-negative when all bidders report truthfully. This is the familiar property of *individual rationality* (IR).

2.1 Incentive Compatibility Characterization

Consider a dynamic auction for IDV bidders that can only misreport their signal. Constantin et al. [5] establish three conditions that are necessary and sufficient for IC in dynamic, IDV auctions. Two of the conditions require that a bidder's price does not go down if he misstates his interval— they are trivially satisfied in the known-interval setting of this paper.³ The third condition, adapted to a no-interval-misreports domain, requires:

UNCONDITIONAL CRITICAL SIGNAL: Fix the signals of other bidders. For bidder i there is a signal, $c_i^u[s_{-i}]$, such that i is allocated if and only if $s_i \geq c_i^u[s_{-i}]$ (and is ∞ if i is not allocated for any s_i .) When allocated, the payment by i is $v_i(c_i^u[s_{-i}], s_{-i})$.

This implies that the allocation rule is monotonic in the bidder's signal. The existence of an unconditional critical signal generalizes the “critical-value” concept in private-value settings, where a bidder wins iff his value is higher than the critical value, which is also the price he pays [12].

In designing optimal, dynamic IDV auctions we find it easier to work with an equivalent characterization that is defined in terms of *conditional* critical signals, when coupled with additional inter-temporal constraints. This will lead to more natural multi-period optimization problems.

CONDITIONAL CRITICAL SIGNAL: Fix the signals of the other bidders. For bidder i there is a signal, $c_i[s_{-i}]$, such that i is allocated if and only if $s_i \geq c_i[s_{-i}]$ **and** there is an item available for allocation at i 's departure. When allocated, the payment by i is $v_i(c_i[s_{-i}], s_{-i})$.

It is quite easy to see that this property is not sufficient for IC. The reason is that it can be in a bidder's interest to influence whether or not the item is still

³ In Constantin et al. [5] it is also shown that if departures and signals are private then no “reasonable” IC auction is possible, but one can have a “reasonable” IC auction with private arrivals and private signals, but known departures.

available at its departure. Consider a scenario in which i loses (before departing) to a competitor h when reporting signal s_i , but in which i can misreport some signal s'_i causing h to lose (e.g. if her critical signal goes from below to above s_h when i 's signal changes from s_i to s'_i), and resulting in i now winning at a price less than his true value for the item. To address this we must combine conditional critical signals $c_i[s_{-i}]$ with additional inter-temporal constraints:

Theorem 1. *A dynamic auction in the known-interval, IDV model is IC if and only if it has conditional critical signals with the property that there are no signals $s_{<i}$, s_i , and s'_i such that: (a) $c_j[s_i, s_{-ij}] \leq s_j$ for some $j < i$; (b) $c_j[s'_i, s_{-ij}] > s_j$ for all $j < i$; and (c) $c_i[s_{-i}] \leq \min\{s_i, s'_i\}$.*

We refer to the constraints in Theorem 1 as the *inter-temporal IC constraints* (ITIC). Given this, there is never an instance for which some bidder i loses when reporting true signal s_i (a), could have prevented all earlier bidders from winning for some $s'_i \neq s_i$ (b), and wins for report s'_i and with a critical signal less than its true signal and thus a payment less than its true value (by v-monotonicity). Conditional critical signals that satisfy ITIC become *unconditional*: i wins if and only if his signal is at least $c_i[s_{-i}]$.

3 Special cases

We now present two simple dynamic IDV environments for which revenue optimal auctions can be easily constructed. We first review Branco's [3] solution for non-dynamic IDV environments and then provide a multi-step optimization formulation for the case of disjoint intervals.

3.1 Revenue-optimal static auctions

Definition 2. *If a bidder i 's valuation is $v_i(s_i, s_{-i})$ then i 's virtual valuation is*

$$\pi_i(s_i, s_{-i}) = v_i(s_i, s_{-i}) - \frac{\partial v_i}{\partial s_i}(s_i, s_{-i}) \frac{1 - F_i(s_i)}{f_i(s_i)} \quad (1)$$

For example, if $i = 3$, signals are distributed uniformly on $[0, 1]$ and $v_3(s_1, s_2, s_3) = s_3 + \frac{s_1}{4} + \frac{s_2}{4} + \frac{1}{4}$ then $\pi_3(s_1, s_2, s_3) = s_3 + \frac{s_1}{4} + \frac{s_2}{4} + \frac{1}{4} - 1 \frac{1-s_3}{1} = 2s_3 + \frac{s_1}{4} + \frac{s_2}{4} - \frac{3}{4}$. Branco [3] provides a revenue-optimal static auction for IDV bidders if i 's *virtual valuation* is increasing in s_i , for all i . This holds in particular if bidder valuations are increasing and concave in their own signal and the distribution of each bidder's signal has a non-decreasing hazard rate ($\frac{f_i(s_i)}{1-F_i(s_i)}$). Branco's auction generalizes Myerson's [10] revenue-optimal private-value optimal auction, and is thus based on the fact that in equilibrium, *virtual valuation = revenue*. This insight suggested the virtual valuation-based heuristic for the non-clairvoyant mechanism that we present in Subsec. 5.1. In Branco's solution, the bidder with the highest virtual valuation π_i wins, but only if π_i is non-negative. The winner pays his value computed at the lowest signal for which he still wins.

Branco's result extends to any dynamic setting in which all bidders' intervals have at least one point in common, making the auction a static one.

3.2 Disjoint intervals

In this subsection we analyze the case of *disjoint* intervals when the number of bidders n is known in advance and we show that the revenue-optimal auction can be obtained as a solution to a multi-period decision problem.

From our earlier characterization, we know that an IC auction in this environment must define a *critical signal schedule*, $(c_j)_{1 \leq j \leq n}$, where c_j denotes the critical signal for bidder j *conditioned on the item still being available*, and computed with knowledge of the signals $s_{<j}$ reported by earlier bidders but not with knowledge of the signal of bidder j himself (else it would not be IC). In this case the ITIC constraints are vacuously satisfied because no bidder can influence the critical signal faced by an earlier bidder.

We can compute an optimal schedule by adopting dynamic programming: bidder j 's critical signal should optimally balance the revenue from selling to him now (at a price of $v_j(c_j, s_{<j})$) and waiting. Let $c_j \in \operatorname{argmax}_c R_j(c, s_{<j})$, where $R_j(c, s_{<j}) = \mathbb{E}[v_j(c, s_{<j}) | s_j \geq c] + \mathbb{E}[R_{j+1}(c_{j+1}, s_{\leq j}) | s_j < c]$ and $R_{n+1}(\cdot) = 0$. $R_j(c, s_{<j})$ is the expected revenue from selling to bidder j at a price defined by critical signal c and selling to future bidders under the *optimal* critical signal schedule. For the last bidder, we get $c_n \in \operatorname{argmax}_c \{(1 - F(c)) \cdot v_j(c, s_{<n})\}$.

Recall that f_i and F_i define, respectively, the pdf and cdf on the signal distribution of bidder i . This probabilistic model is used in defining the base case for the last bidder, and also provides the distributional information with which the expected revenue is computed in earlier periods.

4 Working with a Specific Problem instance

For the remainder of the paper we will work with a set of four specific 3-bidder scenarios as shown in Fig. 1. In all scenarios bidders' arrivals and departures are known and fixed, with a slight variation for the scenarios in which bidder 3 is supposed to arrive later than time d_1 . In those scenarios, if bidder 3 has not arrived yet, the auctioneer only knows the (correct) probability of bidder 3 arriving $0 \leq p_{3A} \leq 1$. These simple scenarios capture, we think, the essence of interactions in dynamic environments for IDV bidders.

Each scenario is labeled $X_{12}X_{13}X_{23}$ where X_{ij} is 'T' or 'F' (shorthand for 'True' and 'False') and X_{ij} specifies whether s_j can be used in v_i , i.e. whether bidder j arrives before bidder i departs. For instance, in scenario TFT bidder 1 uses his signal and 2's while bidders 2 and 3 use the signal of both other bidders. Note that because of the bidder ordering, one cannot have scenarios FTF, FTT or TTF – bidders' arrival and departure order is the same: 1,2,3. We omitted scenario FFT as we expect it to be analogous to TFF.

For simplicity, we will model the signal of any bidder as being *uniformly* distributed on $[0, 1]$. In the continuous Branco formulation, the distribution-dependent part of i 's virtual valuation (see Eq. (1)) will then be $1 - s_i$.

4.1 MIP formulation

We present a mixed-integer programming (MIP) formulation for the TFT scenario. This formulation can be extended to any other dynamic scenario.

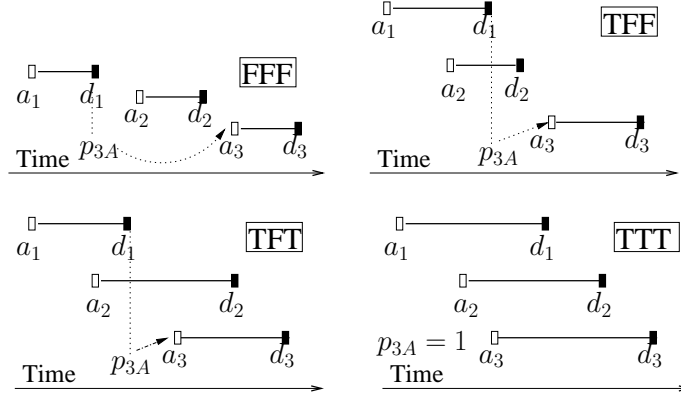


Fig. 1. Bidder configuration as known a priori (at a_1) in each of the four scenarios. At d_1 or d_2 , if 3's scheduled arrival a_3 is later, the probability that 3 will arrive at all is p_{3A} . If bidder 3 arrives, then his interval is *known* to be $[a_3, d_3]$.

We discretize the signal space of each bidder b such that s_b can take on values s_{b1}, \dots, s_{bm_b} . We will hereafter assume that bidders' valuations are weighted-linear in the signals available $v_i(s_i, s_{-i}) = s_i + \sum_{j \neq i} v_{ij} s_j$. In Section 5 we will use specific numeric values for the v_{ij} weights, i.e. for the weight that bidder i assigns to the signal of some bidder $j \neq i$. This discretized type space provides an approximation to a continuous one.

To use a mechanism defined on a discrete type space in a continuous one, compute the critical signal for a bidder $1 \leq i \leq 3$ as follows: let s_j and s_k be the other bidders' signals and \underline{s}_j and \underline{s}_k be the *highest* discrete signals *lower than* s_j and s_k respectively. Bidder i 's critical signal is then the discrete critical signal that was computed for $(\underline{s}_j, \underline{s}_k)$.

Decision variables As seen in Theorem 1, any IC auction can be defined by conditional critical signals when coupled with ITIC constraints. The MIP includes decision variables c_i , instantiated for each bidder i on the discretized signals of the other relevant bidders, and defining the conditional critical signals (just "critical signals" here after). A bidder gets the item if no earlier bidder won the item and the bidder's signal satisfies $c_i \geq c_i(s_{-i})$ upon its departure. Given critical signals, binary variables λ, μ and ν encode whether each bidder (1, 2 and 3 respectively) does not win the item:

$$\begin{aligned}
 \lambda_{ij} = 1 &\iff s_{1i} < c_1(s_{2j}) \\
 \mu_{ijk} = 1 \text{ iff } s_{2j} < c_2(s_{1i}, s_{3k}) &\iff \mu_{ij}^{-3} = 1 \text{ iff } s_{2j} < c_2^{-3}(s_{1i}) \\
 \nu_{ijk} = 1 &\iff s_{3k} < c_3(s_{1i}, s_{2j})
 \end{aligned} \tag{2}$$

Within the MIP, one can capture logic such as $\lambda_{ij} = 1 \iff s_{1i} < c_1(s_{2j})$ via linear constraint $-M\lambda_{ij} \leq s_{1i} - c_1(s_{2j}) < M(1 - \lambda_{ij})$ where a "big M " is adopted, and set to the smallest constant that can be proved to be larger than the maximal absolute value of $s_{1i} - c_1(s_{2j})$. We also impose monotonicity

constraints, $\lambda_{ij} \leq \lambda_{i-1,j}$, $\mu_{ijk} \leq \mu_{i,j-1,k}$, $\mu_{ij}^{-3} \leq \mu_{i,j-1}^{-3}$, and $\nu_{ijk} \leq \nu_{i,j,k-1}$ on the indicator variables.

These indicator variables are used both in the ITIC constraints and in the objective function. For example, λ_{ij} indicates whether bidder 1 does not win the item when its signal is s_{1i} and bidder 2's signal is s_{2j} . Note that in TFT bidder 1's critical signal depends on s_2 , bidder 3's on (s_1, s_2) and bidder 2's depends on whether or not bidder 3 arrives. Variables $c_2(s_{1i}, s_{3k})$ and μ_{ijk} capture the behavior of the auction in respect to bidder 2 when bidder 3 arrives while variables μ_{ij}^{-3} and $c_2^{-3}(s_{1i})$ are for the case without bidder 3.⁴

ITIC constraints In the TFT scenario the ITIC constraints are encoded as:

- If $p_{3A} < 1$, bidder 2 does not have a useful signal misreport when 3 does not arrive: he cannot report a signal $s_{2j'}$ instead of s_{2j} such that he loses with j , wins with j' and his critical signal is less than his true signal s_{2j} (note that bidder 2 can only lose to bidder 1). In critical signal notation,

$$\nexists s_{1i}, s_{2j}, s_{2j'} \text{ s.t. } \begin{cases} c_1(s_{2j}) \leq s_{1i} \text{ and } s_{1i} < c_1(s_{2j'}) \text{ and} \\ c_2^{-3}(s_{1i}) \leq s_{2j} \text{ and } c_2^{-3}(s_{1i}) \leq s_{2j'} \end{cases}$$

which, using the variables in Eq. (2), can be written as

$$\nexists i, j, j' \text{ such that } (\neg \lambda_{ij}) \wedge \lambda_{ij'} \wedge (\neg \mu_{ij}^{-3}) \wedge (\neg \mu_{ij'}^{-3})$$

or still: $\forall i, j, j' \quad (1 - \lambda_{ij}) + \lambda_{ij'} + (1 - \mu_{ij}^{-3}) + (1 - \mu_{ij'}^{-3}) \leq 3$, as a linear constraint.

- If $p_{3A} > 0$, bidder 2 does not have a useful signal misreport when 3 arrives: he cannot report a signal $s_{2j'}$ instead of s_{2j} such that he loses with j , wins with j' and his critical signal is less than his true signal s_{2j} (note that in this case bidder 2 can only lose to bidder 1). In critical signal notation,

$$\nexists s_{1i}, s_{2j}, s_{2j'}, s_{3k} \text{ s.t. } \begin{cases} c_1(s_{2j}) \leq s_{1i} \text{ and } s_{1i} < c_1(s_{2j'}) \text{ and} \\ c_2(s_{1i}, s_{3k}) \leq s_{2j} \text{ and } c_2(s_{1i}, s_{3k}) \leq s_{2j'} \end{cases}$$

which is easily expressed into linear constraints as above.

- If $p_{3A} > 0$, bidder 3 does not have a useful misreport when 3 arrives:

$$\nexists s_{1i}, s_{2j}, s_{3k}, s_{3k'} \text{ s.t. } \begin{cases} c_2(s_{1i}, s_{3k}) \leq s_{2j} \text{ and } s_{2j} < c_2(s_{1i}, s_{3k'}) \text{ and} \\ c_3(s_{1i}, s_{2j}) \leq s_{3k} \text{ and } c_3(s_{1i}, s_{2j}) \leq s_{3k'} \end{cases}$$

The TFT scenario misses one type of ITIC constraint that is present in scenarios where a bidder i can influence the ability of more than one earlier bidder to win the item. For example, if there is a bidder 4 that is visible to both 2 and 3: $a_3 < a_4 < d_2 < d_3 < d_4$ then the ITIC constraints for bidder 4 would require: he cannot report a signal $s_{4l'}$ instead of s_{4l} such that he loses

⁴ Note that $\sum_i \lambda_{ijk}$ may be different than 1 for some j and k as long as $[\min s_{1i}, \max s_{1i}]$ does not cover s_1 's domain: e.g. if $c_1(s_{2j}, s_{3k}) < \min s_{1i}$. A similar observation holds for $\sum_j \mu_{ijk}$ and $\sum_k \nu_{ijk}$.

with l , wins with l' and his critical signal is less than his true signal s_{4l} . Note that bidder 4 could lose to bidder 2 *or* bidder 3. In critical signal notation,

$\nexists s_{1i}, s_{2j}, s_{3k}, s_{4l}, s_{4l'}$ such that

$$\begin{cases} [c_2(s_{1i}, s_{3k}, s_{4l}) \leq s_{2j} \text{ or } c_3(s_{1i}, s_{2j}, s_{4l}) \leq s_{3k}] \text{ and} \\ \quad s_{2j} < c_2(s_{1i}, s_{3k}, s_{4l'}) \text{ and} \\ \quad s_{3k} < c_3(s_{1i}, s_{2j}, s_{4l'}) \text{ and} \\ \quad c_4(s_{1i}, s_{2j}, s_{3k}) \leq s_{4l} \text{ and } c_4(s_{1i}, s_{2j}, s_{3k}) \leq s_{4l'} \end{cases} \quad (3)$$

which can also be easily expressed in linear constraints.

Objective function The objective is to maximize the expected revenue given the probabilistic model on bidder signals and whether or not bidder 3 will arrive. Let $\mathbb{P}(s_{il})$ denote the discrete probability mass assigned to discrete signal level l for bidder i by evenly distributing pdf f_i . We write the objective as:

$$\sum_{k,j,i} \mathbb{P}[s_{3k}] \mathbb{P}[s_{2j}] \mathbb{P}[s_{1i}] (R_1(i, j) + (1 - p_{3A}) R_2^{-3}(i, j) + p_{3A} (R_2(i, j, k) + R_3(i, j, k)))$$

where

$$\begin{aligned} R_1(i, j) &= (1 - \lambda_{ij}) (c_1(s_{2j}) + v_{12}s_{2j} + v_{10}) \\ R_2^{-3}(i, j) &= \lambda_{ij} (1 - \mu_{ij}^{-3}) (c_2^{-3}(s_{1i}) + v_{21}s_{1i} + v_{20}) \\ R_2(i, j, k) &= \lambda_{ijk} (1 - \mu_{ijk}) (c_2(s_{1i}, s_{3k}) + v_{21}s_{1i} + v_{23}s_{3k} + v_{20}) \\ R_3(i, j, k) &= \lambda_{ijk} \mu_{ijk} (1 - \nu_{ijk}) (c_3(s_{1i}, s_{2j}) + v_{31}s_{1i} + v_{32}s_{2j} + v_{30}) \end{aligned}$$

Recall that v_{ij} is a constant, denoting the weight that the bidder assigns to the signal of some other bidder $j \neq i$. To linearize the objective, note that an objective term such as $R_1(i, j) = (1 - \lambda_{ij}) (c_1(s_{2j}) + v_{12}s_{2j} + v_{10})$ can be reduced to $R_1(i, j) \leq c_1(s_{2j}) + v_{12}s_{2j} + v_{10}$ and $R_1(i, j) \leq M(1 - \lambda_{ij})$ for a suitable big M constant. Similar tricks can be used for the other terms in the objective. The $R_h(\cdot)$ quantities measure the revenue obtained when signals take the specific values s_{1i}, s_{2j} and s_{3k} and bidder h wins the item. For instance $R_2(i, j, k)$ is only activated if bidder 1 has not won the item ($\lambda_{ij} = 0$ i.e. $s_{1i} < c_1(s_{2j})$) but bidder 2 wins the item ($\mu_{ij} = 1$ i.e. $s_{2j} \geq c_2(s_{1i}, s_{3k})$). The winning bidder pays his valuation at the critical signal and the signals of other bidders.

Formulation size In considering the size of the MIP in a general scenario we consider the following problem characteristics:

- D^+ ($\leq m$): the maximum number of earlier bidders' signals that can influence a bidder's value
- D^- : the maximum number of earlier bidders whose value can be influenced by any single bidder's signal
- S : the maximum size of the discrete signal space of any bidder
- n : the number of bidders
- m : the maximum number of other bidders' signals that can influence the critical signal to some bidder.

By generalizing Eq. (2), one can see that for each bidder i , each discrete signal s_i and each possible signal tuple of bidders that influence i 's value, there is a variable controlling whether s_i is smaller than the critical signal c_i computed at that particular tuple. Since in the tuple there can be at most D^+ bidders, *the number of variables is $O(nS^{D^++1})$.*

To analyze the number of constraints, one needs to extrapolate Eq. (3) since the bulk of constraints will be of this form. Each such ITIC constraint must ensure that no bidder i ($i = 4$ in Eq. (3)) can misreport his signal such that: all earlier bidders (2 and 3 in Eq. (3)) who should be winning when i is honest do not win anymore and furthermore, i wins. There are n bidders in total. Each bidder i can influence the value of at most D^- earlier bidders and each of those other bidder's ITIC constraint depends on its own signal, together with perhaps m other signals, where m is the maximal number of other signals that can influence its critical signal. Thus, *the number of constraints is $O(nD^-S^{m+1})$.* Note that there are no ITIC constraints when $D^- = 0$, which occurs either in the IDV but disjoint interval case or in the non-disjoint but private-values case. In both of these cases it is sufficient for IC to simply formulate the decision problem as one of setting critical signals.

Remarks: Signals are discrete and critical signals are only constrained by inequalities of the form $s_i < c_i$, from the ITIC constraints. Because critical-signals have non-negative weights in the objective function, if a critical signal c_i is in between two consecutive discrete points s_a and s_{a+1} then a solution as least as good as the current one can be obtained by setting $c_i = s_{a+1}$, without affecting any ITIC constraints. From this, we see that critical signals will only be defined at discrete points.

While it is reasonable to consider structured problems in which D^- and D^+ are small, the main bottleneck in encoding MIPs for large instances is in the dependence of the number of constraints on S^{m+1} . For practical formulations the designer will need to impose some limit to the number of earlier signals that can factor into setting the critical signal for the current bidder, or adopt an alternate formulation that restricts this dependence to some other derived statistic; e.g., the *maximal earlier value* of a departing bidder, or the *maximal signal* of an earlier bidder.

5 Instantiation

We now go on to instantiating particular numbers for the valuation model described above. This leads to our simulation study. Let $s_1, s_2, s_3 \sim U[0, 1]$ and assume the following valuations (whose choice we will motivate shortly):

$$\begin{aligned} v_3(s_1, s_2, s_3) &= s_3 + \frac{s_1}{4} + \frac{s_2}{4} + \frac{1}{4} \\ v_2(s_1, s_2, s_3) &= \begin{cases} s_2 + \frac{s_1}{4} + \frac{s_3}{4} + \frac{1}{4} & , \text{if } d_3 \leq a_2 \\ s_2 + \frac{s_1}{2} + \frac{1}{4} & , \text{otherwise} \end{cases} \\ v_1(s_1, s_2, s_3) &= \begin{cases} \frac{3}{2}s_1 + \frac{1}{4} & , \text{if 1 cannot see 2 or 3} \\ s_2 + \frac{s_1}{2} + \frac{1}{4} & , \text{if 1 can see 2, but not 3} \\ s_1 + \frac{s_2}{4} + \frac{s_3}{4} + \frac{1}{4} & , \text{if 1 can see both 2 and 3} \end{cases} \end{aligned}$$

The virtual valuations are

$$\begin{aligned}\pi_1(s_1, s_2, s_3) &= \begin{cases} 3s_1 - \frac{5}{4}, & \text{, if 1 cannot see 2 or 3} \\ 2s_1 + \frac{s_2}{2} - \frac{3}{4}, & \text{, if 1 can see 2, but not 3} \\ 2s_3 + \frac{s_1}{4} + \frac{s_2}{4} - \frac{3}{4}, & \text{, if 1 can see both 2 and 3} \end{cases} \\ \pi_2^{-3}(s_1, s_2, s_3) &= 2s_2 + \frac{s_1}{2} - \frac{3}{4} \\ \pi_2(s_1, s_2, s_3) &= 2s_2 + \frac{s_1}{4} + \frac{s_3}{4} - \frac{3}{4} \\ \pi_3(s_1, s_2, s_3) &= 2s_3 + \frac{s_1}{4} + \frac{s_2}{4} - \frac{3}{4}\end{aligned}$$

The valuations were chosen to be symmetric and such that:

- they depend linearly on the signals available,
- $\frac{\partial v_i}{\partial s_i} > \frac{\partial v_i}{\partial s_j} \forall i, j$, implying the v -monotonicity and SCC conditions
- $v_1 = v_2 = v_3$ and $\pi_1 = \pi_2 = \pi_3$ when $s_1 = s_2 = s_3$ in *all* scenarios, except $\pi_1 < \pi_2 = \pi_3$ in scenario FFF.

The final property ensures that there is no a priori bias between bidders and it is meant to facilitate the analysis of the interaction of interdependent values and uncertainty about the future. One cannot have, however, $v_1 = v_2 = v_3$ and $\pi_1 = \pi_2 = \pi_3$ when $s_1 = s_2 = s_3$ in *all* scenarios unless values are private. For the weights we have chosen, $\mathbb{E}[\pi_1(s_1)] = 0.25$ in scenario FFF, less than 0.5, the expected value of any bidder's virtual valuation in all other scenarios.

5.1 Experimental Results

We compare the performance of the MIP-based auction with two additional auctions as summarized in Table 1. The *continuous, clairvoyant* (ContCV) auction provides a best-case revenue. Here, we allow the auctioneer to observe the signals of all bidders from the start, regardless of the scenario. Bidder i 's critical signal is then computed as the least signal for which i has the highest non-negative virtual valuation π_i (recall that π_i depends on the scenario). In other words, ContCV implements Branco's auction as if the auctioneer could have perfect knowledge of all signals.

The *continuous, non-clairvoyant* (Cont NonCV) auction is a naive generalization of Branco's offline auction that is not IC in general. At each departure d_i , the item is sold to i iff i 's virtual valuation π_i is higher than the maximum of the expectations of the virtual valuations of bidders still to arrive and zero (as in Eq. (4)). If any bidder j has already departed, π_i is *not* compared with π_j . For example, in scenario TFT, this solution requires selling to:

$$\begin{aligned}\text{Bidder 1, if } \pi_1 &\geq (1 - p_{3A}) \max(\pi_2^{-3}, 0) + p_{3A} \mathbb{E}[\max(\pi_2, \pi_3, 0)] \\ \text{Bidder 2, if } \pi_2^{-3} &\geq 0 \text{ when 3 does not arrive and haven't sold to 1} \\ \text{Bidder 2, if } \pi_2 &\geq \max(\pi_3, 0) \text{ when 3 arrives and haven't sold to 1} \\ \text{Bidder 3, if } \pi_3 &\geq 0 \text{ when 3 arrives and haven't sold to 1 or 2}\end{aligned}\tag{4}$$

Auction	CV	Cont	IC		FFF	TFF	TFT	TTT
ContCV	✓	✓	✓	Cont CV	0.8875	0.9471	0.9474	0.9448
Cont NonCV	X	✓	X	Cont NonCV	0.8848	0.8551	0.85	0.9448
MIP policy	X	X	✓	MIP policy	0.8843	0.9015	0.9018	0.9175

Table 1. Auctions used for evaluation (left), as characterized by whether or not the auctioneer is clairvoyant (CV), whether or not a discretization is imposed when designing the auction (Cont == “continuous,” no discretization), and whether or not the auction is IC. The expected revenue is summarized by scenario (right) in an environment where bidder 3 always arrives.

For scenario TFT in particular however, this is *not IC*: one can verify that for $s_1 = 0.6, s_2 = 0.3$ and any $s_3 \leq 0.3$ bidder 2 is better off reporting $s'_2 = 0.5$.

For the *MIP-based policy*, recall that there are m_i discrete signals for each bidder i . We chose them to be $\frac{1}{1+m_i}$ through $\frac{m_i}{1+m_i}$, where $m_1 = 14, m_2 = 9, m_3 = 7$. These numbers were calibrated so as to balance solution quality in the time allotted with the granularity of the discretization. A higher emphasis was placed on bidders 1 and 2 since the critical signals vary the most for them. In the MIP formulation, 3’s critical signals are usually not used for high values of s_1 : if $s_3 \geq c_3$ then 1 wins the item since $s_1 \geq c_1$ as well (in other words, the unconditional critical signal of bidder 3 for high values of s_1 is ∞). As mentioned before, critical signals c_1 will only take values at discrete signals s_{1i} . Allowing bidder 1 the most discrete signals makes the solution more informative.

The MIP formulation was encoded using CPLEX and JOpt, a simplified Java Wrapper for mixed integer or linear programming [8]. It was run on each scenario and auction for three hours on a Pentium IV at 3GHZ, allowing 256MB of memory for CPLEX. For comparison, 256MB is not enough memory in the TTT scenario for $m_1 = m_2 = m_3 = 19$. Note that the CPLEX solution was stopped before reaching its tolerance level of 99.9%; any feasible solution to the MIP is however IC, because it must satisfy the ITIC constraints.

Throughout the experiments we fix $p_{3A} = 1$, meaning that bidder 3 always arrives and consider the four scenarios from before. All critical signal sets were evaluated using sampling of signals: 500000 independent uniform samples of $(s_1, s_2, s_3) \in [0, 1]^3$ were taken for the clairvoyant tests and 10000 for the non-clairvoyant, due to these taking significantly longer.

5.2 Empirical Results

The average revenue for the MIP formulation and the Clairvoyant and Non-Clairvoyant settings are shown in Table 1, right. Let us take a closer look at these numbers. We see that, in all 3 auctions, as the bidders intervals’ overlap increases, so does the revenue – we should expect this since the uncertainty in the model decreases with the amount of overlap.

The MIP solution generates a higher revenue than the non-clairvoyant one in scenarios TFF and TFT. Thus, it appears that tackling the online problem directly provides a better solution despite having to impose ITIC constraints and having the auction limited to be defined at discrete points. Scenarios TFF and TFT have produced close revenues for all auctions.

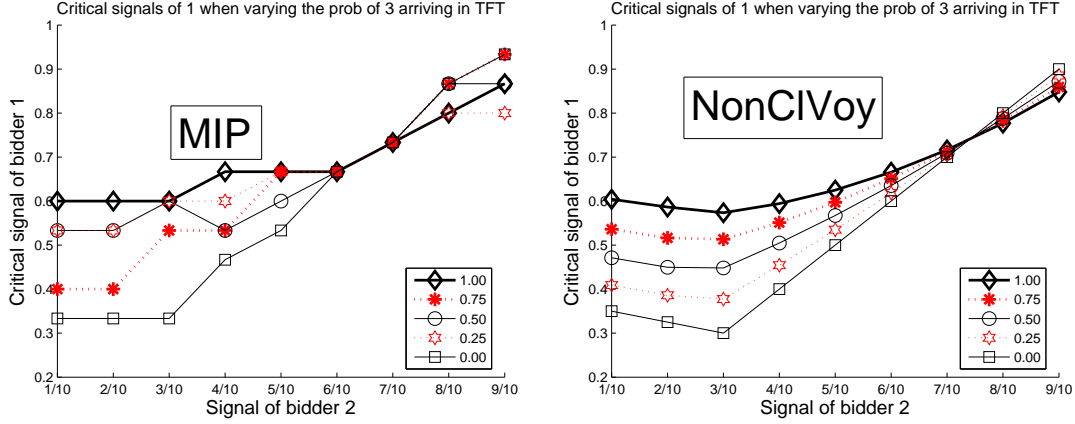


Fig. 2. Critical signals of bidder 1 as a function of bidder 2's signals in scenario TFT for several values of p_{3A} (shown in legend). The critical signals obtained from the MIP policy are on the left, while the ones obtained from the (Branco) ContNonCV heuristic are on the right.

We also experimented with varying p_{3A} in scenario TFT. Recall that in scenario TFT, bidder 1 can only observe bidder 2's signal, but bidders 2 and 3 observe every bidder's signal. At d_1 only the probability p_{3A} of 3 arriving is known – whether 3 arrives will be known before d_2 . This scenario offers the most interesting IC considerations (for bidder 2), at least without bidders fully overlapping.

As expected, as the probability p_{3A} of bidder 3 arriving goes down, so does the expected revenue of the optimal auction as approximated by our MIP formulation. The table below was obtained by sampling, in the same way as above, 100000 triples of signals:

p_{3A}	1.0	0.75	0.5	0.25	0.0
revenue	0.902	0.868	0.848	0.815	0.807

The critical signals for bidders 2 and 3 computed at time d_1 are basically identical for all values of p_{3A} . What differs are the critical signals of bidder 1, plotted in Fig. 2, left. The signals have an increasing, convex shape and appear to become flatter as p_{3A} increases.⁵ The similarity with the signals predicted by the Branco non-clairvoyant (on the right in the same Figure) is striking. Recall though that the Branco non-clairvoyant are *not IC*! We are currently investigating, through additional computation, whether the phenomenon we see here is actually *monotonicity ironing* in order to achieve IC (in the sense of Myerson [10]) or due to the discretizations imposed in the MIP methodology together with suboptimal solutions. We conjecture that we are in fact seeing ironing, in this case leveling of critical signals in regions of the signal space that may violate

⁵ The variance in this picture is probably due to CPLEX stopping in each instance on one of the many approximately optimal solutions.

IC (in our case, pairs of low s_1 and s_2 in Fig. 2). By way of comparison, we also ran our formulation on the TTT scenario, which is effectively a one-shot, non-dynamic scenario. The (approximate) optimal auction was almost identical to the optimal IC auction predicted by Branco’s result, as described in Subsec. 3.1.

6 Discussion and future work

In this paper we adopt a computational approach to the design of dynamic, interdependent-value auctions. We formulate the design problem as a mixed-integer program defined on discretized signals. The formulation has reasonable size if the amount of interdependence between bidders’ valuations is small, a coarse signal discretization can be tolerated, and the a bound is imposed on the maximal dependence on earlier signals of the price to any bidder. Our formulation leverages a “critical-signal based” characterization for truthful, online auctions and we formulate the design problem as one of setting appropriate critical signals, which in turn induce an optimal decision policy.

In future work we intend to use our optimization approach to better understand the relationship between Branco/Myerson-style virtual valuation-based methods and the IC solutions generated through optimization. Another relevant extension includes identifying graphical, interdependent value models that can provide structure to leverage computationally. It would also be nice to extend to allow for uncertainty both in the number of bidders, and also the time intervals that define bids.

References

1. M. Aoyagi, “Optimal Sales Schemes Against Interdependent Buyers”, Discussion paper No. 645, Osaka University, October 2005.
2. L. Blumrosen and N. Nisan, ”Auctions with Severely Bounded Communication”. In Proceedings of FOCS 2002, pp. 406-415.
3. F. Branco, “Multiple Unit Auctions of an Indivisible Good“, Economic Theory, vol. 8, pp. 77-101, 1996.
4. V. Conitzer and T. Sandholm, “Complexity of Mechanism Design”. In Proceedings of Uncertainty in Artificial Intelligence 2002, pp. 103-111.
5. F. Constantin, T. Ito and D. Parkes, “Online Auctions for Bidders with Interdependent Values”, Proceedings of AAMAS 2007.
6. M. Hajiaghayi, R. Kleinberg and D. Parkes, “Adaptive Limited-Supply Online Auctions”, Proceedings of ACM Electronic Commerce 2004, pages 71–80.
7. M. Hajiaghayi, R. Kleinberg, M. Mahdian and D. Parkes, “Online Auctions with Re-usable Goods”, Proceedings of ACM EC 2005, pages 165–174.
8. JOpt, a Simplified Java Wrapper for Mixed Integer and Linear Programming, <http://www.eecs.harvard.edu/econcs/jopt/>
9. V. Krishna, “Auction Theory”, Academic Press, 2002.
10. R. Myerson, “Optimal Auction Design”, Mathematics of Operations Research, VI (1981), 58-73.
11. M. Pai and R. Vohra, “Optimal Dynamic Auctions”, Technical Report, kellogg.northwestern.edu/faculty/vohra/ftp/vohra8.pdf, 2006.
12. D. Parkes, “Online Mechanisms”, in “Algorithmic Game Theory”, edited by N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani, Cambridge Univ. Press, 2007