Online Ascending Auctions for Gradually Expiring Items

Ron Lavi and Noam Nisan *

School of Computer Science and Engineering,
The Hebrew University of Jerusalem, Israel

Abstract

In this paper we consider online auction mechanisms for the allocation of M items that are identical to each other except for the fact that the items have different expiration times, and each item must be allocated before it expires. A computational application is the allocation of time slots in a scheduling problem, and an economic application is the allocation of transportation tickets.

If we ignore incentive issues, then 2-approximation online allocation algorithms are known. We, however, are interested in situations where players act "selfishly" and may mis-report values, deadlines, or arrival times. We first design two auction-like mechanisms and prove that a 3-approximation allocation is obtained for a wide class of "semi-myopic" selfish behaviors of the players. The rest of the paper aims at providing a game-theoretic rational justification for acting in such a semi-myopic way.

We first show that the usual notions of incentive compatibility in dominant strategies can not be used in this case, since any such equilibria can not obtain better than an M-approximation. Instead we suggest a new notion of "Set-Nash" equilibria, where we can not pin-point a single best-response strategy, but rather only a set of possible best-response strategies. These strategies are all semi-myopic and thus our mechanisms will perform well on any of them. We believe that this notion is of independent interest.

*E-mails:{tron, noam}@cs.huji.ac.il. Supported by grants from the Israeli Ministry of Science and the Israeli Academy of Sciences.

1 Introduction

In recent years we have seen a growing body of work that analyzes distributed computer systems under the the assumption that participants will behave "rationally" in the game-theoretic sense of optimizing their own utility, rather than behaving obediently or maliciously as was commonly considered in computer science [26, 1, 11, 12, 9, 23]. The term "algorithmic mechanism design", coined in [26], refers to *designing* distributed systems that will function well under these assumptions.

The most convincing notion used to justify why rational players will follow the specification is that of playing a dominant strategy: no matter what others do, the player will benefit most by playing according to the specification. Unfortunately, for many settings, solutions with dominant strategies are rare, and a need to find other suitable equilibrium notions rises. This seems to be the case in the setting of online auctions with worst-case guarantees. Here, the collective experience of the CS community [21, 8, 3, 5, 7, 13, 2, 19, 27, 16] seems to indicate that strong assumptions about the players are necessary to obtain truthfulness, and the resulting auctions are sometimes somewhat artificial (e.g. without real competition among the different bidders). We aim to replace existing equilibrium concepts with weaker ones, while seeking auctions with realistic flavor in which *many* "reasonable" behaviors lead to approximate social optimum.

The problem we study is the online allocation of M items that are all identical except that they "expire" at different times: the first item expires at time 1, the second at time 2, and so on. Players arrive over time, and items must be allocated at or before their expiration time. Each player j desires any single item between his arrival time, r_j, and his deadline, d_j. Player j has a value v_j for receiving the item. All information r_j, d_j, v_j is private to player j, and players act rationally to maximize their utility: the value v_j, if they are allocated an item, minus any payment that they must pay. Our goal is to design a mechanism that maximizes the "social welfare", i.e. to allocate the items so that the sum of values of players that receive an item is maximized.

This model seems applicable to many scenarios in which items are sequentially allocated as time progresses, where both items and players have a finite "life-time". In a computational setting, this model is equivalent to online scheduling of unit length jobs with deadlines. Focusing on the algorithmic question only, and ignoring incentive issues, it is known that the offline problem can be solved exactly in polynomial time and that the online problem has a simple greedy algorithm that achieves a 2-approximation [18, 4], but no online algorithm can achieve a competitive ratio better than the "golden ratio" [15]. Of course, this algorithm requires each player to reveal his true type (value, arrival time, and deadline) and ignores players' strategic considerations.

Our first attempt was to design a truthful (incentive compatible) mechanism for this problem, in the usual sense of dominant strategies. Indeed a truthful polynomial-time mechanism for the *offline* problem is obtained using the known offline polynomial time algorithm and VCG payments. However, this can not be applied to the online problem, since online algorithms can only achieve an approximation and the difficulty of obtaining truthful mechanisms from approximation algorithms is well known [25, 20]. Indeed we first prove a strong impossibility result:

Theorem: *Any truthful online mechanism cannot obtain a better than M-approximation.*

One could approach this difficulty in several directions. First, we could attempt to replace truthfulness with a Bayesian-Nash equilibrium. However, the behavior of players in such mechanisms critically depend on the details of their assumptions regarding the a-priori distribution of other players' types. Such distributional assumptions are commonly considered problematic in computer science[1], and are also considered less desirable than "detail free" mechanisms in the economics literature [30]. A second approach is to insist on truthfulness, to assume a-priori knowledge of the range $[v_{min}, v_{max}]$ of possible player values (where v_{min} is strictly positive), and to allow randomization. In this case, the market clearing algorithm of [8] can be used to obtain a truthful randomized $2(\log(v_{max} - v_{min}) + 1)$-approximation. This demonstrates well the limitations of truthfulness, as a deterministic 2-approximation without any assumptions on players exists once truthfulness is dropped.

Instead, our approach is to continue with a worst-case analysis over players' types, but to relax the required notion of equilibria: instead of specifying a single tuple of strategies (the equilibrium point) that provides a good approximation ratio, we will specify a large set of strategies with the property that the mechanism will perform well on *any* of them. Our strategic analysis will not be able to pinpoint exactly which of these strategies will be rationally chosen, but rather only that one of them will be – this is enough to guarantee good performance.

At this point we embark with the algorithmic analysis of two (variants of) classic ascending auctions, under a wide class of possible player behaviors. We begin by providing intuitive reasoning why players will conform to one of these behaviors, and describe exact strategic justifications (and a new notion of equilibria) below.

The first auction we consider is a natural adaptation of the iterative auction of Demange, Gale, and Sotomayor [10] (similar offline scheduling auctions were also considered by [28]). *The Online Iterative Auction* constantly maintains a current price p_t and a current winner win_t for every item t. At each time t, each player (in his turn) may place his name as the temporary winner of some item t' (bid on t'), deleting the previous temporary winner, and increasing the price by some fixed small δ (a player can be a temporary winner only for one item). When none of the players wishes to bid, the time t phase ends: item t is sold to player win_t for a price of $p_t - \delta$. At time $t+1$ the prices and temporary winners from time t are kept, and the auction continues similarly.

In the offline setting, where all players arrive at time 1, [10] show that if all players behave *myopically*, i.e. always bid on the item with the lowest price among those that interest them, then the auction will reach the optimal allocation. Moreover, such myopic behavior is indeed the player's best interest [14]. But what will players choose, facing this auction in the online setting? This clearly depends on the beliefs about the future: if a player fears that new bidders will arrive in the future, then he may bid aggressively for earlier items, offering a higher price for them but reducing his risk of future competition. To incorporate such considerations, we call a player *semi-myopic* if he always bids on some item with price lower than his value (not necessarily the item with the lowest price, as the myopic behavior require). Thus there exist many semi-myopic behaviors, that represent different beliefs.

Theorem: *If all players are semi-myopic then the Online Iterative Auction achieves a 3-approximation of the welfare.*

The second auction we consider is *The Sequential Japanese Auction*: Item t is sold at time t as follows. The auctioneer gradually raises a price, starting from 0. As the price ascends, each participating player should decide whether to drop out or to stay (once a player drops out, he cannot re-join). The participants are allowed to observe how many drop-outs occur, and incorporate this, among other factors, into their drop-out decision. The price halts exactly when all players, besides one, have dropped. The remaining player wins, and pays the price that was reached[2].

[1]Indeed the reader is challenged to come up with a distribution on players' types in this model (value, arrival time, and deadline), whose theoretical analysis would seem to shed light on real world problems.

[2]Actually, our pricing scheme is a bit more delicate and described in the paper body.

Considering the offline case first, we observe that a myopic behavior again leads to an optimal allocation, were myopic here has the following somewhat surprising meaning: a myopic player drops exactly when either the price reaches his value or when there are $d-t$ remaining players (where d is the player's deadline, and t the current item number). The logic for dropping when $d-t$ players remain is that at this point the player is assured that there are enough items before his deadline to be allocated to all bidders who are willing to pay the current price. In the offline setting, if all players are myopic, then the optimal allocation is reached. In the online setting, however, players may bid more aggressively for the current item, not dropping when the number of players goes down to $d-t$, for fear of new players arriving and competing for later items. In this setting we call a player *semi-myopic* if, at every time t, he drops no later than when the price reaches his value, and no earlier than when only $d-t$ other players remain in the auction.

Theorem: *If players are semi-myopic then the Sequential Japanese Auction obtains a* 3*-approximation.*

But why should the players play as we expect? We now turn to give a more accurate game-theoretic analysis of the players' behaviors. As truthful auctions does not exist, we instead seek an equilibrium notion that will capture the idea advocated above, i.e. that the best we can do is recommend on a set of strategies, and not on a specific, single strategy.

Our setting is that of games with incomplete information, without any a-priori distribution of types. Player i's utility function is $u_i(t_i, s_1 \ldots s_n)$, where $t_i \in T_i$ denotes his privately known type and $s_j \in S_j$ is the action performed by player j. Each player's strategy $s_i : T_i \to S_i$ denotes which action he plays as a function of his type. Our notion of equilibria captures the situation where we do not describe a single strategy $s_i(t_i)$ for each type as the equilibrium point, but rather a set $R_i(t_i) \subseteq S_i$ of "recommended strategies" to choose from. We denote $R_i(*) = \cup_{t_i} R_i(t_i)$.

Definition: *The set functions $\{R_i(\cdot)\}$ are in* **Set-Nash equilibrium** *if for any player i, and any action combination of the other players $s_{-i} \in R_{-i}(*)$, player i has a best response to s_{-i} in $R_i(t_i)$.*

Thus the definition requires that a best response to any tuple of recommended strategies may be found within the recommended strategies. This becomes equivalent to regular Nash equilibrium when $|R_i(t_i)| = 1$ for all i and t_i. This definition only considers best responses to pure strategies, and stronger variants of this definition may also require a best response for mixed strategies of the other players (or even for correlated mixed strategies of the other players.)

Although Set-Nash equilibrium is a weak notion, certainly weaker than regular ex-post-Nash equilibrium, it seems to us that it carries some weight, especially in computerized environments, in which appropriate protocols and software programs that act "as recommended" are available, and so a deviation would seem to require some effort. In such cases players can be realistically expected to act as recommended unless they have clear incentives to deviate. Such a clear incentive would seem to be absent when the recommended strategies are in Set-Nash equilibrium (but see the discussion in the paper body.)

Completing this line of thought, we show that both our online ascending auctions have Set-Nash equilibria that are all semi-myopic. We leave the description of the appropriate sets of recommended strategies to the body of the paper. The main point we arrive at is that players do not have a clear incentive to deviate outside of these sets of recommended strategies; and when they do stay inside the set of recommended strategies, the mechanism obtains a 3-approximation.

Main Theorem: *The Online Iterative Auction and the Sequential Japanese Auction both have a Set-Nash equilibrium which results in a* 3*-approximation of the welfare.*

As a by-product of the proof, we also get an interesting equivalence between the two auctions. We show a one-to-one mapping, such that any recommended strategy combination of the online iterative auction has a recommended strategy combination of the Sequential Japanese Auction that leads to an identical allocation (and vice-versa).

The rest of the paper is organized as follows. The model and basic definitions are given in section 2. In section 3 we describe the two online ascending auctions, and show their algorithmic properties by characterizing a family of 3-approximations. Section 4 returns to the strategic setting, showing that no truthful auction can achieve an approximation better than M. In section 5 we define the new notion of Set-Nash equilibrium, and analyze our auctions accordingly. Appendix A describes some useful observations about the offline allocation problem, used as basic building blocks in our proofs. The full proofs appear in appendices B,C,D, and F.

2 Model and Basic Definitions

Items: We wish to sell M identical items with different expiration times. W.l.o.g. we assume that the first item expires at time 1, the second at time 2, and so on. Each item must be sold (and received by the buyer) at or before its expiration time.

Players: The potential buyers of the items (players/bidders) arrive over time. Player i arrives to the market at time $r(i)$, and stays in the market for some fixed period of time, until his deadline $d(i)$. We assume w.l.o.g. that the arrival and departure times are integers[3]. Each player desires only one item (unit demand), that expires no earlier than his arrival time. He must receive it at or before his departure time[4]. Player i obtains a value of $v(i)$ from receiving such an item, otherwise his value is 0. We assume w.l.o.g. that different players have different values[5].

We assume the standard game-theoretic setting: Player i privately obtains his variables $r(i)$, $d(i)$, and $v(i)$, at time $r(i)$. He acts selfishly in order to maximize his own utility: his obtained value minus his price. I.e., a player may arrive at or after his true arrival time, and declare or act as if he has any value, and any deadline.

Our goal: is to maximize the social welfare: the sum of (true) values of players that receive an item.

Basic definitions: Player i is *active* at time t if $r(i) \leq t \leq d(i)$, and i did not win any item before time t. Let A_t be the set of all active players at time t. An *allocation* is a mapping of items to players such that, if player i receives item t, then $r(i) \leq t \leq d(i)$. Let X_t be an allocation of items $t, ..., M$. $X_t[d]$ denotes the player that receives item d according to X_t, and $X_t[d_1, d_2] = \cup_{d=d_1}^{d_2} X_t[d]$, the set of players that receive items d_1 through d_2. By a slight abuse of notation we also use X_t as the set of players $X_t[t, M]$. The *value* of X_t is $v(X_t) = \sum_{d=t}^{M} v(X_t[d])$, i.e. the welfare obtained by X_t. A set S of players is *independent* with respect to items $t, ..., M$ if there exists an allocation of (part of) the items $t, ..., M$ such that every player in S receives an item.

The offline allocation problem: The offline problem, in which all players arrive at time 1, is a matroid: a set of *players* is independent if there exists an allocation of (part of the items) to these players. This was shown in [17] for the unit-demand scheduling problem, which is equivalent to ours. This matroid structure is used extensively in our proofs. See Appendix A for more details.

[3]As actions in a non-integral time point can be deferred to the next integral point with no affect.

[4]Our auctions also fit the more severe restriction that player i cannot get an item $t > d(i)$. E.g., player i cannot attend Saturday's show if he is leaving on Friday, even if he receives the ticket before Friday.

[5]I.e. fix some arbitrary order over players, and set $v(i) \succ v(j)$ iff $v(i) > v(j)$ or $v(i) = v(j)$ and $i \succ j$.

3 Two Online Ascending Auctions

We first describe online adaptations of two well-known ascending auctions. These have the property that players do not have to choose specific actions for the auction to perform well: a 3-approximation is obtained for a large, reasonable family of behaviors that we term "semi-myopic". Under any such player behaviors, each of our auctions belongs to a general family of snapshot planning algorithms, that we characterize. We then show that any snapshot planning algorithm obtains a 3-approximation. Full proofs are given in appendix B.

In this section, we focus on the algorithmic side. Therefore we give only intuitive justifications for the player behaviors that we assume. For the same reason, we also omit few technicalities about prices and tie-breaking rules from the definitions. All these are detailed in section 5, where we analyze the strategic properties of our auctions and apply game-theoretic arguments.

3.1 The Online Iterative Auction

We consider an online adaptation of the iterative auction of Demange, Gale, and Sotomayor [10]:

Definition 1 (The Online Iterative Auction (intuitive version)) *The Online Iterative Auction constantly maintains a current price p_t and a current winner win_t for every item t. These are initialized to zero at $t = 0$, and updated according to players' actions at each time t, as follows:*

- *Each player, in his turn, may place his name as the temporary winner of some item t', causing the previous winner to be deleted, and the price to increase by some fixed small δ. A player cannot perform this action, and must relinquish his turn, if he is already a temporary winner.*
- *When none of the players that are not temporary winners wishes to place their names somewhere, the time t phase ends: item t is sold to the player win_t for a price of $p_t - \delta$.*
- *At time $t + 1$ the prices and temporary winners from time t are kept. If additional players arrive then the auction continues according to the above rules.*

Before analyzing the online auction, it is useful to take a glimpse at the offline case, in which all players arrive at time 1. This is actually a special case of the unit-demand model studied by [10], [14]:

Definition 2 ([10]) *Player i has a* **myopic strategy** *in the iterative auction if, in his turn, he always places his name on the item $t \leq d(i)$ with the minimal price, unless the minimal price $\geq v(i)$, in which case he does not place his name at all.*

Lemma 1 ([10], [14]) *If all players are myopic and arrive at time 1 then the online iterative auction obtains the optimal allocation, and each winner pays the VCG price.*

In the online setting, a player might not be completely myopic, depending on his beliefs about the future. For example, he may bid aggressively for the current item, not placing his name on future items at all. This is reasonable if he anticipates tight competition from players that will arrive later on. Viewing this behavior as one extreme, and the completely myopic behavior as the other, it seems that any combination of the two cannot be "ruled-out". On the other hand, a player might choose not to participate at all for some time units – if, for example, there are M high valued players that desire any item 1 through M, but they all do not participate up to time M, then the resulting welfare will be low. As it turns out, this is the only type of behavior we need to exclude:

Definition 3 *Player i is* **semi-myopic** *if, in his turn, i places his name on some item t with $p(t) \leq v(i)$ and $r(i) \leq t \leq d(i)$. If there is no such item, i stops participating.*

Theorem 1 *If all players are semi-myopic then the online iterative auction achieves almost a 3-approximation: $v(OPT) \leq 3 \cdot v(ON) + 2 \cdot M \cdot \delta$, where OPT, ON are the optimal, online allocations.*

The proof shows that, under any semi-myopic behavior, the online iterative auction is a snapshot planning algorithm (see section 3.3 below), hence obtains the desired approximation (by lemma 4).

3.2 The Sequential Japanese Auction

A different possibility is to sell item t at time t using a simple one item ascending auction:

Definition 4 (A Japanese Auction) *The (classic, one item) Japanese auction operates as follows: An auctioneer gradually raises a price, starting from 0. Each participating player should decide whether to drop out or to stay (once a player drops out, he cannot join again), as the price ascends. The price stops increasing exactly when all players, besides one, have dropped out. The winner is the player that did not drop out, and he pays the price that was reached.*

A natural adaptation of this to the online case is:

Definition 5 (The Sequential Japanese Auction (intuitive version)) *The Sequential Japanese Auction sells each item t at time t, separately, using a Japanese auction with one modification: the participants are allowed to observe how many drop-outs occur as the price ascends (and to incorporate this into their drop-out decision).* [6]

As before, it is useful to first consider this auction for the offline case, in which a rather surprising notion of myopic behavior leads to the optimal allocation:

Definition 6 *Player i is* **myopic** *in the Sequential Japanese Auction if, in the auction of any time t, (for $r(i) \leq t \leq d(i)$), he drops exactly when either the price reaches $v(i)$, or when there are exactly $d(i) - t$ other players that did not drop yet.*

Lemma 2 *If all players are myopic and arrive at time 1 then the Sequential Japanese Auction obtains the optimal allocation.*

It is not hard to verify that this lemma is actually a special case of theorem 6 from the online strategic setting (specifically, it follows from claim 17 in the appendix). We note that the winners pay the VCG price (in the offline case) only when using the modified prices of section 5.5.

In the online setting, players might not play myopically, and may insist on closer items (i.e. stay longer in the auction) if they anticipate much competition in the future. In the extreme, when every player remains in the auction until the price reaches his true value, we actually simulate the simple greedy algorithm, which is a 2-approximation. As before, any behavior between the two extremes can cause only a minor performance degradation. All we wish is that players will not drop out "too soon". Indeed, dropping out early in the auction also have disadvantages, as future auctions might be much more competitive, due to new arriving players.

[6] Prices are also modified. The time-t-winner pays the highest price among all time-t'-auctions in which he tied the time-t'-winner. Defining "a tie" is delicate, and requires both players to drop simultaneously. See section 5.5.

Definition 7 *Player i's strategy is* **semi-myopic** *(for the Sequential Japanese Auction) if, at every time t, he drops no later than when the price reaches his value, $v(i)$, and no earlier than when only $d(i) - t$ other players remain in the auction.*

Theorem 2 *If all players play semi-myopic strategies then the Sequential Japanese Auction achieves a 3-approximation.*

In a similar manner to the iterative auction above, this theorem is proved by showing that, under any semi-myopic behavior, the Sequential Japanese Auction is a snapshot-planning algorithm.

3.3 Snapshot-Planning Algorithms

We characterize a family of snapshot-planning algorithms, that include both the Online Iterative Auction and the Sequential Japanese Auction, when players are semi-myopic. The main point is that any snapshot-planning algorithm obtains a 3-approximation of the welfare.

Specifically, the *snapshot plan* of time t, S_t, is the allocation with maximal value among all allocations of items $t, ..., M$ among the active players, A_t[7]. Define

$$f_t = \{\ j \in S_t \mid S_t \setminus j \text{ is independent w.r.t items } t+1, ..., M\ \}, \tag{1}$$

i.e. the set f_t contains all players that can receive item t, when one plans to allocate items $t, ..., M$ to the players of S_t. Now define the critical value at time t, v_t^*, as:

$$v_t^* = \begin{cases} \min_{j \in f_t}\{v(j)\} & S_t \neq f_t \\ 0 & \text{otherwise} \end{cases}$$

Definition 8 (A snapshot-planning algorithm) *An algorithm is called a snapshot-planning algorithm if every item t is sold at time t to some player j with $v(j) \geq v_t^*$.* [8]

Lemma 3 *The Online Iterative Auction with semi-myopic players and the Sequential Japanese Auction with semi-myopic players are both snapshot-planning algorithms.*

The family of snapshot-planning algorithms can be viewed as the entire range between the following two extremes: the first is the greedy algorithm that always chooses the player with maximal value[9], and the second is the algorithm that always chooses the player that determined v_t^*. These two extremes are 2-approximations (both were studied in the context of online scheduling [18, 6]). The entire family has only a slightly larger approximation ratio:

Lemma 4 *Any snapshot-planning algorithm is a 3-approximation of the welfare, and this is tight.*

4 The Impossibility of Truthful Approximations

We now return to more accurate game-theoretic arguments. Since our goal is to find approximately optimal allocations with respect to the *true* variables of the players, we would prefer to design truthful algorithms: an algorithm is truthful if there exist price functions such that, regardless of

[7] There exists one such allocation, by the matroid structure, and since different players have different values.

[8] A worst-case approximation cannot sell item t before time t, as a player with high value only for t may appear.

[9] Interestingly, this is a special case of the greedy algorithm of [22] for combinatorial auctions with sub-modular valuations. They study the offline case, but it is easy to verify that their algorithm actually works online.

how the other players act, player i will maximize his utility by declaring his true variables to the algorithm[10]. More formally, let T_i be the domain of all valid player i types/bids $(r(i), v(i), d(i))$, and let $T_{-i} = \times_{j \neq i} T_j$. Consider the allocation constructed by the algorithm upon receiving the type $b_i \in T_i$ from player i and $b_{-i} \in T_{-i}$ from the other players, and let $v(i, b)$ be the value that player i obtains from this allocation, i.e. $v(i)$ if i receives one of his desired items, and 0 otherwise.

Definition 9 (Truthfulness) *An algorithm is truthful if there exist price functions $p_i : T_1 \times \cdots \times T_n \to \Re$ such that, for any i, any $b_{-i} \in T_{-i}$, any true type $b_i \in T_i$, and any $\tilde{b}_i \neq b_i$[11]:*

$$v(i, b_i, b_{-i}) - p_i(b_i, b_{-i}) \geq v(i, \tilde{b}_i, b_{-i}) - p_i(\tilde{b}_i, b_{-i}).$$

Such a property is highly desirable, as it guarantees that each player will be motivated to reveal his true type to the algorithm, by an argument similar to the traditional worst-case arguments of Computer Science. Indeed, many recent examples show truthful algorithms for various CS models (some are cited in our introduction). However, for our model, no such algorithm performs well:

Theorem 3 *Any truthful deterministic algorithm for our online allocation problem cannot obtain an approximation ratio better than M.*

Remark 1: Although the proof utilizes an extreme scenario with players with very large values, the worst case ratio occurs in common, simple scenarios, as the proof demonstrates. I.e., since the algorithm defends itself against such extremes, it must make wrong decisions even in simple cases.

Remark 2: A simple truthful deterministic M-approximation exists: let $p_i(t)$ be the highest bid ever seen, except for i's own bid. Sell item t to a player i with $v(i) > p_i(t)$, for a price of $p_i(t)$.

5 Will Players Act As Expected?

Our main motivation at this point is to justify the assumption that players will behave "as expected". We desire a rational justification, i.e. one that shows that expected strategies are, in some sense, utility maximizers for the players. The settings that we are interested in are ones in which "recommended" strategies are indeed to be intuitively expected, and deviating from them would seem to require some effort. In such cases, even rather weak notions of rational justification carry some weight. Such settings include, in particular, situations where computer protocols are announced and appropriate software that acts "as expected" is available. From the onset, we should note that our notions are intended for cases where the existing standard notions of games with incomplete information do not apply: ex-post Nash equilibria do not exist, and no reasonable common prior can be assumed (i.e. we seek "worst-case" notions as in computer science rather than Bayesian notions common in economics).

5.1 Set-Nash Equilibria

Our setting is the standard one of games with incomplete information, but without any distributional assumptions. There are n players, where each player i has a privately known type (input) $t_i \in T_i$. No probability distribution is assumed on $T = T_1 \times ... \times T_n$. A game is defined in which each player i has a strategy space S_i. The outcome of the game is given by the n utility functions $u_i : T_i \times S \to \Re$ where $u_i(t_i, s_i, s_{-i})$ denotes i's payoff when his type is t_i, he plays strategy s_i and

[10]By the revelation principle [24], it is w.l.o.g. to consider only algorithms that receive, as input, the players' types.

[11]We actually restrict the possible $\tilde{b}_i$'s such that $\tilde{r}_i \geq r_i$.

the others play the strategy tuple s_{-i}. The basic assumption is that, given that the other players play s_{-i}, player of type t_i will choose a strategy $s_i \in \text{argmax}\{u_i(t_i, s_i, s_{-i})\}$.

In our setting, a set of recommended strategies is defined for each player. Specifically, a function $R_i : T_i \to 2^{S_i}$ is given, where $R_i(t_i) \subseteq S_i$ is the set of strategies that player i may be expected to follow. We denote also $R_i(*) = \cup_{t_i \in T_i} R_i(t_i)$. The motivating scenario is where it is known that if all players i play strategies $s_i \in R_i(t_i)$, then the outcome is "good" in some sense. E.g., in our case, the obtained social welfare approximates the optimal one. We would like to capture the notion that the sets R_i are in equilibrium. In other words, formalize when can it be said that given that other players $j \neq i$ all play strategies in $R_j(t_j)$, then player i also rationally plays some strategy in $R_i(t_i)$.

Definition 10 *The set functions $\{R_i(\cdot)\}$ are in* **Set-Nash equilibrium** *(in pure strategies) if for every i, every $t_i \in T_i$, every $s_{-i} \in R_{-i}(*)$, and every $s_i \in S_i$ there exists $r_i \in R_i(t_i)$ such that $u_i(t_i, r_i, s_{-i}) \geq u_i(t_i, s_i, s_{-i})$. I.e. if all others play some recommended strategies (not necessarily according to their true types) then there exists a best response strategy for i that is one of i's recommended strategies (according to his true type).*

This definition can be strengthened in several levels, maintaining the spirit of this "set equilibrium" notion. Such strengthenings answer the weak points of the original definition. First, one may require that $R_i(t_i)$ will contain best responses also for mixed recommended strategies of the other players, or even against any distribution over $s_{-i} \in R_{-i}(*)$ (this represents correlated player behaviors). A second level is then to require that any strategy outside of $R_i(t_i)$ will be dominated by some strategy from $R_i(t_i)$ (given that the others play strategies in $R_{-i}(*)$). All of these definitions behave the same on the two extremes: when each R_i is a singleton set (i.e. $|R_i(t_i)| = 1$) then they are equivalent to ex-post Nash equilibrium[12], and when $R_i(t_i) = S_i$ then they are all trivially satisfied.

We require the existence of a pure recommended strategy $r_i \in R_i(t_i)$. One can in principle relax the definition to allow r_i to be a mixed strategy, i.e. a probability distribution on $R_i(t_i)$. It is easy to verify that this does not change the definition (the best mixed strategy is always a pure one). It also does not change the first level of strengthening, while for the second level, this will weaken the definition to become equivalent to the first level (using von-Neuman's max-min principle, in the sense of Yao showing equivalence between distributional complexity and probabilistic one).

This definition suffer from the same caveats of regular Nash-equilibria, in particular noting that inequalities are not strict. Thus for example one can have any of these equilibria in strictly dominated strategies. More refined notions may require that strategies in $R_i(t_i)$ are undominated, or even that all undominated best-responses are in $R_i(t_i)$.

Another refinement is to show that the best response is in $R_i(t_i)$ even when other player's strategies reside in a wider class than $R_{-i}(*)$ (this may be interesting also when i assumes only partial rationality of the other players). One may formally define the wider set of acceptable strategies $A_i \subseteq S_i$, where $R_i(*) \subseteq A_i$, and replace the quantification of $s_{-i} \in R_{-i}(*)$ in the definition with $s_{-i} \in A_{-i}$.

In this paper we use the basic definition. In addition, all our Set-Nash strategies are undominated, and one can show that they are best response to a set of acceptable strategies wider than $R_{-i}(*)$. Indeed, we feel that an interesting problem we leave open is to find Set-Nash equilibria that contain best responses for mixed recommended strategies as well.

[12] I.e. for all i, all $t_i \in T_i$, all $s_i \in S_i$, and all $s_{-i} \in S_{-i}$, $u_i(t_i, R_i(t_i), s_{-i}) \geq u_i(t_i, s_i, s_{-i})$.

5.2 Snapshot Planning Mechanisms

We now devise a sealed bid auction with our two basic building blocks: it has a Set-Nash equilibrium, and, for these equilibrium strategies, the auction is a snapshot planning algorithm. We will later use this to show that both our online ascending auctions also have such a structure. The players' types and their utility functions are as described in section 2.

Definition 11 *We define the* **snapshot planning mechanism** *as follows:*

Strategy space: *Each player declares, as he arrives, his value, his deadline, and a tentative deadline between his arrival time and his deadline. The variable $d(i,t)$ holds i's tentative deadline if t is not larger than his tentative deadline, otherwise $d(i,t)$ equals his final deadline.*

Winner determination at time t: *Let A_t, S_t, and f_t be the natural parallels of the notions in definition 8, where the deadline of each player in A_t is $d(i,t)$. The mechanism allocates item t to some player in f_t (this choice may depend on the contents and structure of A_t, S_t, and f_t).*

Prices: *For each player i, the mechanism maintains a tentative price for each time t, $p_t(i)$, as follows: If $i \notin S_t$ then $p_t(i) = 0$. For any $i \in S_t$, let*

$$c_t(i) = max\{v(j) \mid j \in A_t \setminus S_t,\ S_t \setminus i \cup j \text{ is independent w.r.t items } t, ..., M\}. \quad (2)$$

For any $i \in f_t$, the mechanism sets $p_t(i) = c_t(i))$. For any $i \in S_t \setminus f_t$, the mechanism may set any price $p_t(i) \in [0, c_t(i)]$. The winner i of time t pays $max_{r(i) \le t' \le t} p_{t'}(i)$.

The recommended strategies: *In a recommended strategy, i declares his true value and deadline at time $r(i)$, and may declare any tentative deadline.*

Lemma 5 *For any player i, and any $s_{-i} \in R_{-i}(*)$, i has a best response to s_{-i} in $R_i(t_i)$.*

Lemma 6 *When all players play recommended strategies according to their true types then the snapshot planning mechanism is a snapshot planning algorithm.*

And as an immediate result, we get:

Theorem 4 *The recommended strategies of the snapshot planning mechanism form a Set-Nash equilibrium, which guarantees a 3-approximation of the welfare.*

In appendix D.1 we show that, unfortunately, the recommended strategies of the snapshot planning mechanism do not form a Set-Nash equilibrium for mixed strategies.

5.3 Ignorable Extensions of Games

We now aim to show that each of our ascending auctions has "inside" it a snapshot planning mechanism. But first, let us describe this type of a building block more generally.

When one actually attempts to implement a game as a software protocol, it often turns out that the set of strategies that is available to players has grown: the protocol that allows a player to play any strategy in S_i turns out to enable also other strategies, informally ones that are "locally" like a valid strategy $s_i \in S_i$, but that do not correspond to any single valid strategy. These new strategies may open up new strategic behaviors. We will specify the requirements from such implementations needed to maintain Set-Nash equilibria.

Formally, given a game with incomplete information $G = (T, S, u)$ (where T, S, u are the players' type space, the players' strategies, and the players' utility functions, as described in section 5.1

above) we say that $\bar{G} = (T, \bar{S}, \bar{u})$ is an extension of G if $S_i \subseteq \bar{S}_i$ for all i and $\bar{u_i}(t_i, s) = u_i(t_i, s)$ for all $t_i \in T_i$ and $s \in S$ (i.e. $\bar{u}$ when restricted to S is identical to u).

Clearly a strategy that was best response in G need not be a best response in $\bar{G}$ since the new strategies $\bar{S_i} \setminus S_i$ may be better. "Ignorable" extensions of G will not allow such better strategies:

Definition 12 *We say that $\bar{G}$ is an* **ignorable extension** *if for all i, all $t_i \in T_i$, all $s_{-i} \in S_{-i}$ and all $\bar{s_i} \in \bar{S_i}$ there exists $s_i \in S_i$ such that $u_i(t_i, s_i, s_{-i}) \geq u_i(t_i, \bar{s_i}, s_{-i})$. I.e. if others play an original strategy then I have an original strategy which is a best response.*

Proposition 1 *If $\{R_i(\cdot)\}$ are a Set-Nash equilibrium of G and $\bar{G}$ is an ignorable extension of G then $\{R_i(\cdot)\}$ are a Set-Nash equilibrium of $\bar{G}$.*

5.4 The Online Iterative Auction

We now show that our Online Iterative Auction is an ignorable extension of the snapshot-planning mechanism, thus having a Set-Nash equilibrium which approximates the welfare, according to theorem 4. For this, we need to refine our intuitive definition:

Definition 13 (The Online Iterative Auction) *We apply the following modifications to Def. 1:*

Prices: *The auction maintains a tentative price $p_t(i)$ for each player i at time t, as follows: if i is a tentative winner at the end of the iterations of time t then $p_t(i)$ equals to the tentative price of i's item, otherwise $p_t(i) = 0$. The winner i of time t pays $max_{r(i) \leq t' \leq t}\{p_{t'}(i)\}$.*

Recommended strategies: *i's strategy is recommended if i chooses a tentative deadline $d \leq d(i)$, plays myopically (as in Def 2) with value $v(i)$ and deadline d in all times $r(i) \leq t \leq d$, and plays myopically with value $v(i)$ and deadline $d(i)$ in all times $t > d$.*

It is not hard to verify that these recommended strategies are semi-myopic.

Theorem 5 *The Online Iterative Auction is an ignorable extension of a snapshot planning mechanism.*

Corollary 1 *The recommended strategies of the Online Iterative Auction form a Set-Nash equilibrium which obtains a 3-approximation of the welfare.*

5.5 The Sequential Japanese Auction

To show that our Sequential Japanese Auction is an ignorable extension of a snapshot planning mechanism, we need to modify payments similarly to the modification of the Online Iterative Auction. For this, we need to handle simultaneous "drop" announcements more carefully: At any price level p, several players may want to drop. Furthermore, this may be an on-going process, as after one player drops, another one now wants to drop as well. We need to determine more accurately the order among them. This information is used in order to determine f_t (interestingly, we are not able to compute S_t entirely, only f_t, which is enough).

Definition 14 (The Sequential Japanese Auction) *The basic auction structure remains the same as in Def 5. Two additional points should be handled:*

Simultaneous "drop" announcements: *Define $D(p,n)$ as the set of players (among those who did not drop yet), that wish to drop when the price level is p and the number of remaining players is n. At every price level p, the auction solicits drop announcements by repeatedly accepting only one drop announcement out of $D(p,n)$, and decreasing n by 1.*[13] *When $D(p,n) = \emptyset$, the price increases. The winner is, as before, the last remaining player.*

Prices: *Prices $p_t(i)$ for every player i at every time t are maintained as follows: Let k be the number of non-drop-outs just before the price ended its time-t ascend, at a level of p^*. Let $D(p^*,k), D(p^*,k-1), ..., D(p^*,1)$ be the order of drop-outs at this level. Define the critical number $x^* = \min\{0 < x < k \ : \ |D(p^*,x+1)| = 1\}$, and $D^* = \cup_{x \leq x^*} D(p^*,x)$. For any player i, if $i \in D^*$ set $p_t(i) = p^*$, otherwise $p_t(i) = 0$. The winner i of time t pays $max_{r(i) \leq t' \leq t}\{p_{t'}(i)\}$.*

Recommended strategies: *i's strategy is recommended if he arrives at $r(i)$, choose a tentative deadline $d \leq d(i)$, plays myopically with parameters $v(i), d$ until time d, and plays myopically with parameters $v(i), d(i)$ thereafter.*

Again, these recommended strategies are semi-myopic.

Theorem 6 *The Sequential Japanese Auction is an ignorable extension of a snapshot planning mechanism.*

Corollary 2 *The recommended strategies of the Sequential Japanese Auction form a Set-Nash equilibrium which achieves a 3-approximation of the welfare.*

Acknowledgements

We wish to warmly thank Yair Bartal and Ahuva Mu'alem for their help in the early stages of this work. We also thank Moshe Babaioff, Liad Blumrosen, Rica Gonen, Daniel Lehmann, and Motty Perry for many helpful discussions and comments.

References

[1] Aaron Archer and Eva Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42st Annual Symposium on Foundations of Computer Science(FOCS'01)*, 2001.

[2] Baruch Awerbuch, Yossi Azar, and Adam Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proc. of the 35th ACM Symposium on Theory of Computing (STOC'03)*, 2003.

[3] Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *Proc. of the 13th Symposium on Discrete Algorithms (SODA'02)*, 2002.

[4] Yair Bartal, Francis Y. L. Chin, Marek Chrobak, Stanley P. Y. Fung, Ron Lavi, Jiri Sgall, and Tomas Tichy. Online competitive algorithms for maximizing weighted throughput of unit jobs. In *Proc. of the 21st Symposium on Theoretical Aspects of Computer Science (STACS'04)*, 2004.

[13] E.g. if $D(p,n) = X$ then some $i \in X$ is chosen to be dropped, $X \setminus i \subseteq D(p,n-1)$ and $i \notin D(p,n-1)$.

[5] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi-unit combinatorial auctions. In *Proc. of the ninth Conference on Theoretical Aspects of Rationality and Knowledge (TARK'03)*, 2003.

[6] Yair Bartal and Ron Lavi. Analysis of the β-opt algorithm, 2002. Unpublished manuscript.

[7] Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *Proc. of the 14th Symposium on Discrete Algorithms (SODA'03)*, 2003.

[8] Avrim Blum, Toumas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. In *Proc. of the 13th Symposium on Discrete Algorithms (SODA'02)*, 2002.

[9] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communication. In *Proc. of the 43rd Annual Symposium on Foundations of Computer Science (FOCS'02)*, 2002.

[10] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.

[11] Joan Feigenbaum, Christos Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1), 2001.

[12] Amos Fiat, Andrew Goldberg, Jason Hartline, and Anna Karlin. Competitive generalized auctions. In *Proc. of the 34th ACM Symposium on Theory of Computing (STOC'02)*, 2002.

[13] Eric Friedman and David Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *Proc. of the ACM Conference on Electronic Commerce (EC'03)*, 2003.

[14] Faruk Gul and Ennio Stacchetti. The english auction with differentiated commodities. *Journal of Economic Theory*, 92:66–95, 2000.

[15] Bruce Hajek. On the competitiveness of on-line scheduling of unit-length packets with hard deadlines in slotted time. In *Proc. of the 2001 Conference on Information Sciences and Systems*, 2001.

[16] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David Parkes. Adaptive limited-supply online auctions. In *Proc. of the ACM Conference on Electronic Commerce (EC'04)*, 2004.

[17] Ellis Horowitz and Sartaj Sahni. *Fundamentals of Computer Algorithms.* Computer Science press, Rockville, Maryland, 1978.

[18] Alexander Kesselman, Zvi Lotker, Yishay Mansour, Boaz Patt-Shamir, Baruch Schieber, and Maxim Sviridenko. Buffer overflow management in qos switches. In *Proc. of the 33rd ACM Symposium on Theory of Computing (STOC'01)*, 2001.

[19] Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *Proc. of the 44rd Annual Symposium on Foundations of Computer Science (FOCS'03)*, 2003.

[20] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *Proc. of the 44rd Annual Symposium on Foundations of Computer Science (FOCS'03)*, 2003.

[21] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science*, 310:159–180, 2004.

[22] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proc. of the ACM Conference on Electronic Commerce (EC'01)*, 2001.

[23] Daniel Lehmann, Liadan O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[24] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory.* Oxford university press, 1995.

[25] Noam Nisan and Amir Ronen. Computationally feasible vcg mechanisms. In *Proc. of the ACM Conference on Electronic Commerce (EC'00)*, 2000.

[26] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

[27] Ryan Porter. Mechanism design for online real-time scheduling. In *Proc. of the ACM Conference on Electronic Commerce (EC'04)*, 2004.

[28] Michael P. Wellman, William E. Walsh, Peter R. Wurman, and Jeffrey K. MacKie-Mason. Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 35:271–303, 2001.

[29] Dominic J. A. Welsh. *Matroid theory.* London : Academic Press, 1976.

[30] Robert Wilson. Game-theoretic analyses of trading processes. In Truman Bewley, editor, *Advances in Economic Theory: Fifth World Congress*, pages 33–70. Cambridge University Press, 1987.

A Useful Properties of Offline Allocations and Matroids

Definition 15 (A Matroid) *A Matroid is a finite set S and a collection $I \subseteq 2^S$ of independent sets, such that:*

1. $\emptyset \in I$
2. *If $X \in I$ and $Y \subseteq X$ then $Y \in I$.*
3. *If $X, Y \in I$ and $|X| = |Y| + 1$ then there exists $j \in X \setminus Y$ such that $Y \cup j \in I$. If $X \subseteq S$ but $X \notin I$ then it is a dependent set. A base of a matroid is a maximal independent set, and a cycle is a minimal dependent set.*

Claim 1 *The offline allocation of M items among a set A of players is a matroid, where S is the set of players, and a subset X of players is independent if there exists an allocation of (part of) the items to all the players in X.*

proof: The first two conditions of the matroid are trivially satisfied. Let us verify the third one. Let X, Y are be two independent sets with $|X| > |Y|$. We first claim that there exists allocations for X, Y such that, for any $j \in X \cap Y$, j receives the same item in both allocations. To see this, start from arbitrary two allocations, and choose some $j \in X \cap Y$. Suppose j receives items t_1, t_2 in the allocations $1, 2$, where $t_1 < t_2$. Suppose that player j' receives item t_2 in allocation 1. Then we can swap players j and j' in allocation 1, so that j will receive item t_2 (this is valid as we know he can receive this item) and j' will receive item t_1 (this is valid as $t_1 < t_2$). Notice that we have strictly decreased the number of players in $X \cap Y$ that receive different items, and so repeating this implies the result. Now, choose some item t which is being allocated for X but not allocated to any player of Y. Suppose that t is allocated to j in the allocation of X. By our assumption, $j \notin Y$, and so $Y \cup j$ is independent: use the previous allocation of Y, and allocate item t (that beforehand was not allocated) to j. ∎

The following claim lists some useful matroid properties. For extensive discussion and proofs, see e.g. the textbook [29].

Claim 2 *Let $M = (S, I)$ be a matroid. Then:*

1. *If $X, Y \in I$ and $|X| < |Y|$ then there exists $Z \subseteq X \setminus Y$ such that $|X \cup Z| = |Y|$ and $X \cup Z \in I$.*

2. *If B_1, B_2 are bases then $|B_1| = |B_2|$.*

3. *If B_1, B_2 are bases, then, for any $j \in B_1 \setminus B_2$ there exists $j' \in B_2 \setminus B_1$ such that $B_1 \setminus j \cup j' \in I$ and $B_2 \setminus j' \cup j \in I$.*

The following claims are slight alterations of classical properties:

Claim 3 *Let $X, Y \in I$, and $X \nsubseteq Y$. Then, for any $j \in Y \setminus X$ such that $X \cup j \notin I$ there exists $j' \in X \setminus Y$ such that $X \setminus j' \cup j \in I$ and $Y \setminus j \cup j' \in I$.*

proof: If $|X| = |Y|$ then we can assume w.l.o.g. that both are bases (as $I' = \{\ Z \in I \mid |Z| \leq |X|\ \}$ are also the independent sets of a matroid), and the claim immediately follows.

If $|X| > |Y|$ then assume, as before, that X is a base. There exists $Z \subseteq X \setminus Y$ such that $B = Y \cup Z$ is a base. Since $j \in Y \setminus X$ then $j \in B \setminus X$ and so there exists $j' \in X \setminus B$ such that $X \setminus j' \cup j \in I$ and $B \setminus j \cup j' \in I$. Since $Y \subseteq B$ and $j' \in Y \setminus X$ as well, the claim follows.

If $|X| < |Y|$ then assume that Y is a base, take some $Z \subseteq Y \setminus X$ such that $B = X \cup Z$ is a base, and notice that $j \notin Z$ as $X \cup j \notin I$. Thus we can essentially repeat the above logic: j is also in $Y \setminus B$ so there exists $j' \in B \setminus Y$ such that $B \setminus j' \cup j \in I$ and $Y \setminus j \cup j' \in I$. Since $B \setminus Y = X \setminus Y$, and $X \subset B$, then the claim follows. ∎

Claim 4 *Let B be a base of the matroid, and $Y \in I$ such that $|B \setminus Y| = 1$. Then $|Y \setminus B| \leq 1$.*

proof: $|B| \geq |Y| = |B \cap Y| + |Y \setminus B| = |B| - |B \setminus Y| + |Y \setminus B| = |B| - 1 + |Y \setminus B|$. Therefore $|Y \setminus B| \leq 1$, as claimed. ∎

A.1 Some Useful Properties of Offline Allocations

For the following discussion, it will be convenient to assume the following ϵ-assumption: There exists many small valued players in A_t that desire any one of the items $t, ..., M$.

Definition 16 (A prefix) *A subset $X \subseteq S_t$ is called a prefix if it is a prefix of any allocation $S_t[1, M]$ of S_t.*

Claim 5 *$X \subseteq S_t$ is a prefix if and only if for all $j \in X$, $d(j) \leq t + |X| - 1$.*

proof: Suppose first that X is a prefix, and, by contradiction, that there exists some $j \in X$ with $d(j) > t + |X| - 1$. Let $S_t[t, M]$ be some allocation of S_t. Since $j \in X$ and X is a prefix then j is allocated some item $\leq t + |X| - 1$. Suppose player j' is allocated item $d(j)$. Then we can switch between j and j' and have an allocation in which X is not a prefix, a contradiction. In the other direction, if $X \subseteq S_t$ and $d(j) \leq t + |X| - 1$ for any $j \in X$ then, in any allocation, $j \in S_t[t, t + |X| - 1]$. Therefore $X \subseteq S_t[t, t + |X| - 1]$, and since $|S_t[t, t + |X| - 1]| = |X|$ then it follows that $S_t[t, t + |X| - 1] = X$, i.e. it is a prefix. ∎

Definition 17 *For any $t \leq d \leq M$, we build the set of players $P_t(d)$ using the following process (fix any allocation of S_t):*

1. *Let $x_0 = d$.*
2. *For $i > 0$, define inductively $x_i = max\{\ d(j) \mid j \in S_t[t, x_{i-1}]\}$.*
3. *Let k be some index such that $x_{k+1} = x_k$, and fix $P_t(d) = S_t[t, x_k]$.*

Claim 6 *$P_t(d)$ is the prefix with minimal length among all prefixes with length $\geq d - t + 1$.*

proof: First notice that, from the ϵ-assumption it immediately follows that $|P_t(d)| = x_k - t + 1$. Also notice that, by our construction, any $j \in P_t(d)$ has $d(j) \leq x_k = t + |P_t(d)| - 1$. Therefore, by claim 5, $P_t(d)$ is a prefix. Suppose by contradiction that there exists a prefix P' with $d \leq t + |P'| - 1 < x_k$. Choose index i such that $x_i \leq t + |P'| - 1 < x_{i+1}$. But then, by the construction process of $P_t(d)$, we must have a player in P' with deadline at least x_{i+1}, contradicting claim 5. ∎

Claim 7 *$j \in P_t(d)$ if and only if there exists an allocation of S_t in which $j \in S_t[t, d]$.*

proof: If $j \in S_t[t, d]$ then by definition $j \in P_t(d)$. Let us verify the other direction. Fix any allocation of S_t, and compute $P_t(d)$ by that allocation. Assume $j = S_t[d']$ for some $d' > d$ (otherwise the claim immediately follows). Let j_i be the player that determined x_i. Then we have $j_1 \in S_t[t, d]$. Consider the following allocation replacements: allocate item x_1 to player j_1 (this is his deadline, so this is valid), j_2 will get item x_2, ... , j_k will get item x_k. Finally, allocate j's item to j_{k+1} (that received x_k), and allocate j_1's item to j. Therefore we have an allocation in which j receives some item $\leq d$, as claimed. ∎

Claim 8 *$f_t = P_t(t) = P_t(|first_t| + t - 1)$.*

proof: If $j \in first_t$ then there exists an allocation of S_t such that $j = S_t[t]$. Since $P_t(t)$ is a prefix of $S_t[1, M]$ then $j \in P_t(t)$. On the other hand, claim 7 tells us that for any $j \in P_t(t)$ there exists an allocation such that $j = S_t[t]$, and therefore $j \in f_t$. We conclude that $f_t = P_t(t)$. From claim 6 we now get also that $P_t(t) = P_t(|first_t| + t - 1)$, as $P_t(t)$ is a prefix with length $|first_t|$. ∎

Claim 9 *For any t, d with $t < d$, $min_{j \in P_{t+1}(d)}\{v(j)\} \geq min_{j \in P_t(d)}\{v(j)\}$.*

proof: We will actually show that $min_{j \in P_{t+1}(d)}\{v(j)\} \geq min_{j \in P_t(d) \setminus ON[t]}\{v(j)\}$. Let $x = |P_t(d)| + t - 1$, the last item allocated to a player in $P_t(d)$. By the above claims, for any $j \in P_t(d)$, $d_j \leq x$, and $x \geq d$. Let j be the player with minimal value in $P_{t+1}(d)$, and assume by contradiction that $v(j) < min_{j \in P_t(d) \setminus ON[t]}\{v(j)\}$. Therefore $j \notin P_t(d) \setminus ON[t]$. Consider some allocation of S_{t+1} such that j receives item $\leq d$. Now consider $P_t(d) \setminus ON[t]$ and $S_{t+1}[t + 1, x]$. These are two bases of the matroid over items $t + 1, ..., x$. Since $j \in S_{t+1}[t + 1, x] \setminus (P_t(d) \setminus ON[t])$ then there exists $j' \in P_t(d) \setminus ON[t] \setminus S_{t+1}[t+1, x]$ such that $S_{t+1}[t+1, x] \setminus j \cup j'$ is independent (w.r.t items $t+1, ..., x$). As $d_{j'} \leq x$, it follows that $j' \notin S_{t+1}$, and therefore $S_{t+1} \setminus j \cup j'$ is independent as well. As $j' \in A_{t+1}$, and by the maximality of S_{t+1}, we must have $v(j) > v(j') \geq min_{j \in P_t(d)}\{v(j)\}$, a contradiction. ∎

Claim 10 *f_t is independent w.r.t items $t+1,...,M$ if and only if S_t is independent w.r.t items $t+1,...,M$.*

proof: Since $f_t \subseteq S_t$ then the right to left direction is immediate. Let us verify the other direction, i.e. that if f_t is independent w.r.t items $t+1,...,M$ then so is S_t. Let $\tilde{A}_t, \tilde{f}_t, \tilde{S}_t$ be the variables after adding many ϵ players, as in the ϵ-assumption. By the optimality of S_t it follows that $S_t \subseteq \tilde{S}_t$ (when ϵ is small enough). As $\tilde{f}_t$ is a prefix, it cannot be independent w.r.t items $t+1,...,M$. Thus there exists $j \in \tilde{f}_t \setminus f_t$. By definition, $\tilde{S}_t \setminus j$ is independent w.r.t items $t+1,...,M$, and therefore $S_t \setminus j$ is independent w.r.t items $t+1,...,M$. If $j \in S_t$ this will therefore imply $j \in f_t$, a contradiction. Thus j is an ϵ player, and $S_t \subseteq \tilde{S}_t \setminus j$. Since $\tilde{S}_t \setminus j$ is independent w.r.t items $t+1,...,M$, then this implies that so is S_t, as needed. ∎

Sub-claim 1 *Let $A'_t = A_t \cup j'$. Let S_t, S'_t and f_t, f'_t be derived from A_t, A'_t, respectively. Then:*

1. *If $j' \in f'_t$ then $f_t \setminus S'_t \neq \emptyset$.*
2. *$f_t \neq f'_t$ if and only if $j' \in f'_t$.*

proof: From the prefix properties detailed in sub-section A.1 it immediately follows that, if $f_t \subseteq S'_t$ then $f'_t = f_t$, and thus the first claim follows. This also implies the right to left direction of the second claim. We are left to show that, if $f_t \neq f'_t$ then $j' \in f'_t$. By the maximality of S_t, S'_t it follows that either $S_t = S'_t$, or $S'_t = S_t \setminus j \cup j'$ for some $j \in S_t \setminus S'_t$. Since $f_t \neq f'_t$, the latter alternative must hold. If $j \notin f_t$ then $f_t \subseteq S'_t$, implying that $f'_t = f_t$, a contradiction. Thus $j \in f_t$. Therefore there exists an allocation with $j = S_t[t]$. Since $S'_t = S_t \setminus j \cup j'$ then there exists an allocation with $j' = S'_t[t]$ (simply use the previous allocation, changing only the player who receives item t from j to j'). By definition, this implies that $j' \in f'_t$. ∎

B Proofs of Section 3

B.1 Snapshot Planning Algorithms

Lemma 4 *Any snapshot-planning algorithm is a 3-approximation of the welfare, and this approximation factor is tight.*

proof: The lemma follows immediately from claim 11, as any snapshot planning algorithm has $v(ON[t]) \geq v^*_t$, and therefore $v(OPT) = v(OPT \setminus ON) + V(ON) \leq 2\sum_{t=1}^{M} v^*_t + v(ON) \leq 2 \cdot v(ON) + v(ON) = 3 \cdot v(ON)$. This ratio is tight, as demonstrated by example 1. ∎

Claim 11 *Any online allocation algorithm that produces an allocation ON have $v(OPT \setminus ON) \leq 2\sum_{t=1}^{M} v^*_t$, where OPT is the optimal allocation.*

proof: We first prove two useful claims:

Sub-claim 2 *Let A, B be sets of players, where $A \subset B$. Let S_A, S_B be the allocation with optimal value for A, B, respectively (both are over the same set of items). Then if $j \in A$ but $j \notin S_A$ then $j \notin S_B$*

proof: Assume by contradiction that there exists $j \in S_B \cap A$ but $j \notin S_A$. Notice that S_A and S_B are both independent sets of the matroid over players in B. Notice also that, by the contradiction assumption, $S_A \not\subseteq S_B$, otherwise also $S_A \cup j \subseteq S_B$, implying that $S_A \cup j$ is independent, with players only from A, contradicting the maximality of S_A. Therefore, since $j \in S_B \setminus S_A$, there exists $j' \in S_A \setminus S_B$ such that $S_A \setminus j' \cup j$ and also $S_B \setminus j \cup j'$ are both independent. From the maximality of S_A and since $j \in A$, the first condition implies that $v(j') > v(j)$. But then we obtain a contradiction to the maximality of S_B. ∎

Sub-claim 3 *Let S be the allocation with maximal value over the set of players A and the set of items $t, ..., M$. Assume that S is not independent w.r.t items $t+1, ..., M$. Let $j \in S$ be the player with minimal value such that $S \setminus j$ is independent w.r.t items $t+1, ..., M$. Then $S \setminus j$ has maximal value among all independent sets w.r.t items $t+1, ..., M$ and players in A.*

proof: Denote $S' = S \setminus j$. Suppose by contradiction that the maximal allocation X over items $t+1, ..., M$ has $v(X) > v(S')$. If $j \notin X$ then this contradicts the maximality of S, as $X \cup j$ is independent w.r.t items $t, ..., M$. Otherwise $j \in X \setminus S'$. $S' \not\subseteq X$, since otherwise $S = S' \cup j \cup X$ contradicting the fact that S is not independent w.r.t items $t+1, ..., M$. Hence there exists $j' \in S' \setminus X$ such that $X \setminus j \cup j'$ and $S' \setminus j' \cup j$ are independent w.r.t items $t+1, ..., M$. Therefore $S \setminus j'$ is independent w.r.t items $t+1, ..., M$, and from the choice of j it follows that $v(j) < v(j')$, contradicting the maximality of X. ∎

We now prove claim 11. Fix some scenario, and let OPT and ON be the optimal and online allocations for this scenario. We describe $f : OPT \setminus ON \to \{1, ..., M\}$ such that f is 2 to 1 and $v(j) \leq v^*_{f(j)}$ for any $j \in OPT \setminus ON$. From this, claim 11 immediately follows. The function f is defined as follows. Let X_t be the optimal allocation of items $t+1, ..., M$ among players in $OPT[1,t] \setminus ON$. For any $j \in OPT \setminus ON$ (say $j = OPT[t']$), let $t^*_j = min\{t \geq t' \mid j \notin X_t\}$. Then we fix $f(j) = t^*_j$.

Sub-claim 4 *For any $j \in OPT \setminus ON$, $v^*_{f(j)} \geq v(j)$.*

proof: Let $t = f(j)$. First notice that $j \in A_t$: $j \notin ON$, $r(j) \leq t$ as $j \in OPT[1,t]$, and $d(j) \geq t$ since either $j \in X_{t-1}$ or $j = OPT[t]$. Let $m_t \in S_t$ be the player who determined v^*_t, (if $v^*_t = 0$ then set $m_t = null$, so $S_t \setminus m_t = S_t$). We first show that, by subclaim 2, $j \notin S_t \setminus m_t$: define A as $OPT[1,t] \setminus ON$ minus all players with deadline $< t$, and $B = A_t$. Clearly $A \subseteq B$. By definition, X_t is optimal for A (over items $t+1, ..., M$). $S_t \setminus m_t$ is optimal for B (over items $t+1, ..., M$): if $m_t = null$ this follows from the optimality of S_t, and if $m_t \neq null$ this follows from sub-claim 3. Therefore, since $j \notin X_t$ then $j \notin S_t \setminus m_t$. If $j \neq m_t$ then $j \notin S_t$, and since $j \in A_t$ it follows from the optimality of S_t that $v(j) \leq v(m_t)$. If $j = m_t$ then this trivially holds. Therefore $v(j) \leq v(m_t) = v^*_{f(j)}$, and the claim follows. ∎

Sub-claim 5 *f is 2 to 1.*

proof: Fix any time t. We need to show that f maps at most two players to t. Let $j_1 \in X_{t-1}$ be the player with minimal value such that $X_{t-1} \setminus j_1$ is an allocation of items $t+1, ..., M$, and denote $Y = X_{t-1} \setminus j_1$ (if X_{t-1} itself is independent w.r.t items $t+1, ..., M$ then set $Y = X_{t-1}$). If $X_t \subseteq Y$ then by the optimality of X_t it follows that $X_t = Y$ and the claim follows: by definition, f maps only j_1 and $OPT[t]$ to t. Otherwise, $X_t \setminus Y \neq \emptyset$. We first show that $X_t \setminus Y = \{OPT[t]\}$. This is implied by sub-claim 2: set $A = OPT[1, t-1] \setminus ON$, and $B = OPT[1,t] \setminus ON$. Since Y is optimal for A (by sub-claim 3) and X_t is optimal for B (by definition) it follows that, if $j \in OPT[1, t-1]$

but $j \notin Y$ then $j \notin X_t$, i.e. that $X_t \setminus Y = \{OPT[t]\}$. To conclude, we observe that X_t is a base in the matroid over items $t+1, ..., M$ and players $OPT[1,t] \setminus ON$, and that Y is an independent set of that matroid. Therefore $|Y \setminus X_t| \leq |X_t \setminus Y| = 1$, and thus $|X_{t-1} \setminus X_t| \leq 2$. Since $OPT[t] \in X_t$ then, by definition, the players mapped to t are exactly those in $|X_{t-1} \setminus X_t|$, and the claim follows. ∎

The following example shows that the 3-approximation factor analysis is tight:

Example 1 *Consider the following scenario for three items. At time 1 arrive two players, j_1 has value ϵ and deadline 1 and j_2 has value 1 and deadline 2. It is easy to verify that $v_1^* = 0$, and so the online algorithm allocates item 1 to j_1. At time 2 arrive two additional players, j_3 has deadline 2 and j_4 has deadline 3, and both have a value of 1. Therefore $v_2^* = 1$ and the online algorithm chooses j_4. At time 3 no new players arrive, so item 3 remains unallocated by the online algorithm. Therefore its welfare is $1 + \epsilon$. The optimal welfare is, however, 3, as needed.*

B.2 The Two Online Ascending Auctions are 3-Approximations

In this section, for the Online Iterative Auction, we will assume that the price increment of the iterative auction is infinitesimally small. This assumption is not necessary, it will just make notations and proofs simpler.

Theorem 1 *If all players are semi-myopic then the online iterative auction achieves a 3-approximation of the welfare: $v(OPT) \leq 3 \cdot v(ON) + 2 \cdot M \cdot \delta$, where OPT, ON are the optimal and the online allocations, respectively.*

proof: We use here the notations of section 3.3. We show that for each time t, $v(win_t) \geq v_t^* - \delta$, hence the claim follows from claim 11: $v(OPT) = v(OPT \setminus ON) + v(ON) \leq 2 \sum_{t=1^M} v_t^* + v(ON) \leq 2(v(ON) + M \cdot \delta) + v(ON)$.

If $v_t^* = 0$ then, trivially, $v(win_t) \geq v_t^* - \delta$. Thus assume that $v_t^* > 0$. Let Y_t be the allocation of items to the temporary winning players at the end of time t iterations. According to claim 10 in section A.1, f_t is independent w.r.t items $t+1, ..., M$ if and only if $v_t^* = 0$. Therefore f_t is not independent, so there exists some player $j \in f_t$ such that $j \notin Y_t[t+1, M]$. Since $j \in f_t$ then $v(j) \geq v_t^*$. if $j = Y_t[t]$ $(= win_t)$ then we are done. Otherwise, j is not a temporary winner at the end of time t iterations. Since j is semi-myopic, this implies that $v_t^* \leq v(j) < p(t)$. Let $i = win_t$. Since i is also semi-myopic then $v(i) \geq p(t) - \delta$. Therefore $v(win_t) \geq v_t^* - \delta$, as needed. ∎

Theorem 2 *If all players play semi-myopic strategies then the Sequential Japanese Auction achieves a 3-approximation.*

proof: We again use the notations of section 3.3. We will show that the winner has value at least v_t^*, and the claim will follow from lemma 4. Let $j \in f_t$ be the first player in f_t that dropped. If he dropped because the price reached v_j then the winner has value at least v_j, which is at least v_t^*. Otherwise there were at most $d(j) - t + 1$ players that did not drop yet, including j. By claim 5[14], $d(j) - t + 1 \leq |f_t|$. Since no player in f_t dropped yet, it follows that every player that did not drop yet belongs to f_t, hence the winner belongs to f_t. ∎

[14] We can assume that there are no ϵ players in f_t, otherwise $v_t^* = 0$ and the claim trivially holds.

C Proof of Impossibility Theorem

Theorem 3 *Any truthful deterministic algorithm for our online allocation problem cannot obtain an approximation ratio better than M.*

proof: Assume w.l.o.g. that a player that does not win any item pays 0. This implies that i's price must not be higher than his value.

Claim 12 *Fix some algorithm with some fixed approximation ratio. Then, for any player i with $r(i) = 1$ there exists a price function $p_i : T_{-i} \to \Re$ such that, for any combination of players that arrive at time 1, b_{-i}:*

- *If $v(i) > p_i(b_{-i})$ then i wins item 1 and pays $p_i(b_{-i})$ (regardless of his deadline).*
- *If $v(i) < p_i(b_{-i})$ then i does not win any item.*

proof: Fix any combination of players that arrive at time 1, b_{-i}. Suppose first that i has deadline equals to 1. For this case, the player becomes one parameter, and by truthfulness there exist a price function according to the claim [1][15].

We now show that this function p_i satisfies the conditions of the claim, regardless of i's deadline. Fix any deadline $d(i)$ of i. If $v(i) > p_i(b_{-i})$ then i must win some item until his deadline, otherwise he can declare $\tilde{d}_i = 1$ and have strictly better utility. But then, if i does not win item 1, the adversary will produce players with higher and higher values, forcing the algorithm not to allocate any item to i in order to maintain the approximation ratio, thus contradicting truthfulness. Therefore i will receive item 1. He will pay $p_i(b_{-i})$ as otherwise, if he pays a higher price, he will declare $\tilde{d}_i = 1$ and will reduce his price, and if he pays less, then if i will have deadline equals 1 he will declare $d(i)$ instead, thus still winning item 1 but paying less. Therefore the function p_i satisfies the first condition.

Suppose now that $v(i) < p_i(b_{-i})$, and suppose there exists a scenario in which i wins one of his desired items. His price must be at most $v(i) < p_i(b_{-i})$. But then, if i had some value larger than $p_i(b_{-i})$ he would have been better off declaring $v(i)$ instead, by this still winning but paying less. Therefore i cannot win any item at all, and the claim follows. ∎

We can now quickly finish the proof of theorem. Fix any price functions $p_i : T_{-i} \to \Re$. For any $\epsilon > 0$ we will show that there exist player types $b_1, ..., b_M$ such that, for all i, $r(i) = 1$, $d(i) = M$, $1 \leq v(i) \leq 1 + \epsilon$, and $v(i) \neq p_i(b_{-i}$. By the above claim, it follows that the algorithm can obtain welfare of at most $1 + \epsilon$, while the optimal allocation is at least M, and the theorem follows. To verify that such types exist, fix $L > M$ real values in $[1, 1 + \epsilon]$. Choose values $v(i)$ uniformly at random from these L values. Then, for any given i, $Pr(v(i) = p_i(v(-i))) \leq 1/L$, as the values were drawn i.i.d. Thus, $Pr(\exists i,\ v(i) = p_i(v(-i))) \leq M/L < 1$, hence there exist a choice of values with $v(i) \neq p_i(v(-i))$ for all i. ∎

D Full Proofs: Snapshot Planning Mechanisms

[15] The argument essentially states that, if i wins for some $v(i)$ then he ins with any higher value, and pays the same. Therefore there exists a threshold value $p_i^* = p_i^*(b_{-i})$, such that i wins and pays p_i^* if $v(i) > p_i^*$, and looses otherwise.

Lemma 6 *When all players play recommended strategies according to their true types then the snapshot planning mechanism is a snapshot planning algorithm.*

proof: Fix any time t. We need to show that the mechanism chooses a player with value at least v_t^*. Let f_t^{true} be the "true" one, i.e. the relevant set computed with the true player deadlines, and let S_t, f_t be the actual sets computed by the mechanism according to the declared tentative deadlines. If $f_t^{true} \subseteq S_t$ then, by the prefix construction process described in section A.1, and since tentative deadlines are not larger than true ones, $f_t \subseteq f_t^{true}$, and the claim follows. Otherwise there is some $j \in f_t^{true} \setminus S_t$, and so for every $i \in f_t$, $v(i) \geq v(j) \geq v_t^*$ (recall that the declared values are the true ones), as claimed. ∎

Lemma 5 *For any player i, and any $s_{-i} \in R_{-i}(*)$, i has a best response to s_{-i} in $R_i(t_i)$.*

proof: Let σ be the scenario in which all players besides i play s_{-i}, and i does not show up at all. Let

$$t^* = argmin_{r_i \leq t \leq d_i}\{v_t^*(\sigma)\}. \tag{3}$$

Notice that player i can win and pay exactly $v_{t^*}^*$ by arriving at time t^*, declaring any value larger than $v_{t^*}^*$, and a deadline equals to t^*.

Claim 13 *t^* and $v_{t^*}^*$ does not depend on the choice of the winner $i \in f_t$ of time $t \in [r(i), d(i)]$ (where the winners prior to time $r(i)$ are fixed).*

proof: By contradiction, assume that there exist two different scenarios, σ_1, σ_2, that differ only in the choice of the winners (notice that the f_t's themselves might become different during the scenario run due to a previous choice of different winners). Let $v^*(\sigma_i) = min_{r(i) \leq t \leq d(i)}\{v_t^*(\sigma_i)\}$, and let t_i^* be the minimal time in which $v^*(\sigma_i)$ is obtained.

We first assume w.l.o.g. that $v_{t_2^*}^*(\sigma_1) > v_{t_2^*}^*(\sigma_2) = v^*(\sigma_2)$. Let us justify this. If $v^*(\sigma_1) \neq v^*(\sigma_2)$ then w.l.o.g. $v^*(\sigma_1) > v^*(\sigma_2)$ and therefore also $v_{t_2^*}^*(\sigma_1) > v^*(\sigma_2)$. If $v^*(\sigma_1) = v^*(\sigma_2)$ then, by the contradiction assumption, $t_1^* \neq t_2^*$, so w.l.o.g. $t_2^* < t_1^*$. Therefore $v_{t_2^*}^*(\sigma_1) > v^*(\sigma_1) = v^*(\sigma_2)$, as needed. Notice also that from this it follows that $t_2^* > r(i)$, as $A_{r(i)}(\sigma_1) = A_{r(i)}(\sigma_2)$.

Since $v_{t_2^*}^*(\sigma_1) > v_{t_2^*}^*(\sigma_2)$ then $f_{t_2^*}(\sigma_1) \neq f_{t_2^*}(\sigma_2)$, and therefore, by the prefix properties of section A.1, $f_{t_2^*}(\sigma_1) \not\subseteq S_{t_2^*}(\sigma_2)$. Fix some $j \in f_{t_2^*}(\sigma_1) \setminus S_{t_2^*}(\sigma_2)$. Since $v(j) > v_{t_2^*}^*(\sigma_1)$ it follows that $j \notin A_{t_2^*}(\sigma_2)$. This implies that, in σ_2, j is the winner of some time $t' < t_2^*$, i.e. $j \in f_{t'}(\sigma_2) = P_{t'}(t', \sigma_2)$. As $d(j) \geq t_2^*$ then $P_{t'}(t', \sigma_2) = P_{t'}(t_2^*, \sigma_2)$. By claim 9 of section A.1, it therefore follows that $v_{t'}^*(\sigma_2) \leq v_{t_2^*}^*(\sigma_2)$, contradicting the choice of t_2^*. ∎

Claim 14 *i's price in any strategy s_i is at least $v_{t^*}^*$ (where the other players play s_{-i}).*

proof: Recall that σ denotes the scenario in which i does not show up at all. Let σ' be the scenario in which i plays some strategy s_i and the others play s_{-i}. Denote by t_0 the minimal t with $i \in f_t(\sigma')$. We claim that there exists a scenario σ'', that differs from σ only in the choice of winners in f_t, such that $A_{t_0}(\sigma') = A_{t_0}(\sigma'') \cup i$. This follows by the an inductive argument: At time $t < t_0$, $A_t(\sigma') = A_t(\sigma'') \cup i$. Since $i \notin f_t(\sigma')$ then, by sub-claim 1, $f_t(\sigma') = f_t(\sigma'')$. Choose the winner in σ'' to be the winner of σ'. Therefore $A_{t+1}(\sigma') = A_{t+1}(\sigma'') \cup i$, and the inductive claim follows.

Now, at time t_0, since $i \in f_{t_0}(\sigma')$ then, by sub-claim 1, there exists some $j \in A_{t_0}(\sigma') \setminus S_{t_0}(\sigma')$ such that $j \in f_{t_0}(\sigma)$. Therefore i's price is at least $v(j) \geq v_{t_0}^*(\sigma'') \geq v_{t^*}^*(\sigma)$ (where the last inequality follows by claim 13, and the lemma follows. ∎

Claim 15 *The (recommended) strategy of arriving at time $r(i)$, declaring the true value and deadline and declaring a tentative deadline equals to t^* is a best response of i against s_{-i}.*

proof: If $v(i) \leq v^*_{t^*}$ then i cannot possibly gain positive utility, as claim 14 shows, and indeed any recommended strategy will not allocate any item to i.

If $v(i) > v^*_{t^*}$ then, if player i arrives at time t^* and declares tentative deadline t^* he will win item t^* for a price of $v^*_{t^*}$. Let σ be the scenario in which i does not show up at all and σ' be the scenario in which i arrives at $r(i)$ and declares tentative deadline t^*. We claim by induction that, for any $t < t^*$, the winners of σ and σ' are identical, and that i's tentative price is at most $v^*_{t^*}$. Therefore i will win item t^* for a price of $v^*_{t^*}$, and the claim follows. For any $t < t^*$, we have by sub-claim 9 and the construction of t^* that $min_{j \in P_t(t^*,\sigma)}\{v(j)\} \leq min_{j \in P_{t^*}(t^*,\sigma)}\{v(j)\} = v^*_{t^*}(\sigma) < v^*_t(\sigma)$. By the maximality of $S_t(\sigma')$ it follows that, in σ', i replaces the minimal player in $P_t(t^*,\sigma)$, therefore $f_t(\sigma) \subseteq S_t(\sigma')$, and so $f_t(\sigma) = f_t(\sigma')$. By claim 13 we can assume w.l.o.g. that the winner has not changed in the transition from σ to σ'. i's price at time t is (at most, as the mechanism has some freedom in setting this) $min_{j \in P_t(t^*,\sigma)}\{v(j)\} \leq v^*_{t^*}$, and therefore i's final price was not affected as well. ∎

D.1 Bad Examples

We would like to show, by an example, that the recommended strategies of the snapshot planning mechanism do not form a Set-Nash equilibrium in the distribution-free sense. We will start with a basic problematic scenario, and then add to it a second scenario, together obtaining the counter example.

The basic problematic scenario demonstrates that a player might be tempted to arrive later, or to declare a deadline higher than his true one, although this is not his best response:

Example 2 *Consider the following scenario, where (v, d) denotes a player with value v and deadline d):*

- *At time 1 arrive players $(\epsilon, 1), (x_1, 4), (x_2, 4), (x_3, 4), (x_4, 4)$.*
- *At time 2 arrive players $(y_1, 2), (y_2, 3)$.*
- *At time 3 arrive players $(z_1, 5), (z_2, 5)$.*
- *At time 4 arrives a (very large) player $(z_3, 4)$.*

where the values satisfy: $\epsilon < x_2, x_3 < y_1 < x_1 < z_1, z_2 < y_2 < x_4 < z_3$.

If all players declare their true value and tentative deadline equals to their true deadline, a snapshot planning algorithm can choose the winners (first to last) x_1, y_1, y_2, z_3, z_1. So player x_4 looses. However, if he delays his arrival to time 2, or, equivalently, declares a deadline of 5, the winners will be $\epsilon, y_2, x_4, z_1, z_3$, so x_4 will win, with price x_1. Notice, however, that this is not his best response. His best response, to arrive at time 1 and declare tentative deadline 1, is still of-course recommended.

Example 3 *Let scenario 1 be the scenario of example 2, where we consider the decisions faced by x_4, and scenario 2 be as follows:*

- *At time 1 arrive player $(x, 1)$ and our player $(x_4, 4)$.*
- *At time 2 arrive player $(x, 2)$.*

- *At time* 3 *arrive player* $(x, 3)$.

(where $x = x_4 - \epsilon$*). The best response of* x_4 *to scenario 1 is to arrive at time* 1 *and declare deadline* 1*. The best response to scenario 2 is to arrive at time* 1 *and declare a deadline of* 4 *(thus winning item* 0 *with price* 0*). Now suppose that player* x_4 *knows/estimates that both scenarios have probability half. Then, a quick calculation shows that if* x_4 *plays some recommended strategy (and thus arrives at time* 1*) with tentative deadline lower than* 4*, then with probability half (for scenario 2) he will win of the items* 1 *to* 3 *with a resulting utility (i.e. value minus price) of* ϵ*. If his tentative deadline will be* 4 *then with probability half (for scenario* 1*) he will lose. Therefore, any recommended strategy has resulting utility at most* $(x_4 + \epsilon)/2$*. However, if* x_4 *will arrive at time* 2 *and will declare deadline* 4*, a non-recommended strategy, his resulting utility will be half times* $x_4 - 0$ *(for scenario* 2*) plus half times* $x_4 - x_1$*, better than* $(x_4 + \epsilon)/2$ *for small enough* ϵ*.*

E Full Proofs: The Online Iterative Auction has a Set-Nash Equilibrium

Theorem 5 *The Online Iterative Auction is an ignorable extension of a snapshot planning mechanism.*

proof: We first prove that the iterative auction is an extension of a snapshot planning mechanism. We will then show that this extension is ignorable.

Claim 16 *If all players* i *play strategies in* $R_i(*)$ *then the iterative auction is a snapshot-planning mechanism.*

proof: We need to map every recommended strategy of the iterative auction to a strategy of the snapshot-planning mechanism, such that the result of the iterative auction (winners plus payments) will match the criteria of a snapshot planning mechanism. This is done as follows. At time t, map every players that plays myopically with (v, d) to a type (v, d), and denote this set of types as A_t. Let S_t be the optimal allocation of items $t, ..., M$ to the players of A_t. All we need to show is that the iterative auction selects a winner from f_t and sets correct payments. In what follows, we use the notion of a prefix and the claims of section A.1. Let $Y[t, ..., M]$ and $p_t[t, ..., M]$ be the tentative allocation and prices of the iterative auction with the myopic strategies, at the end of time t. For any $d \geq t$, let $P_Y(d)$ be the appropriate prefix of Y, according to definition 17. Define $l(d) = min\{d' \geq t \mid P_Y(d') = P_Y(d)\}$, and

$$c_t(d) = max\{v(j) \mid j \in A_t \setminus Y \text{ and } d(j) \geq l(d)\ \}.$$

(notice that, by abuse of notation, we have defined both $c_t(d)$ for an item $d \in \{t, ..., M\}$, and $c_t(i)$ for a player i. Those are two differently defined terms, although we will see below that they are equal, for $d = Y[i]$).

Sub-claim 6 $p_t(d) \geq c_t(d)$.

proof: Fix any $j \in A_t \setminus Y$ with $d(j) \geq l(d)$. If $d(j) \geq d$ then j has positive value for receiving d. Since j is myopic, it therefore follows that $p_t(d) \geq v(j)$. If $d(j) < d$, then since $P_Y(l(d)) = P_Y(d)$, By the construction of $P_Y(l(d))$, since it is equal to $P_Y(d)$, then there exist players $i_1, ..., i_k$ and items $t_1, ..., t_k$ such that, for any index $x \in \{1, ..., k\}$, $i_x = Y[t_x]$, $d(i_x) \geq t_{x+1}$, $t_1 \leq l(d)$, and $t_k = d$. Since $d(i_x) \geq t_{x+1}$ it follows that $p_t(t_x) \leq p_t(t_{x+1})$, otherwise i_x would have placed his name on item t_{x+1}. Therefore $p_t(d) = p_t(t_k) \geq p_t(t_1)$. Since $t_1 \leq l(d) \leq d(j)$ it follows that $p_t(t_1) \geq v(j)$, and the claim follows. ∎

Sub-claim 7 *If $p_t(d) > p_{t-1}(d)$ then $p_t(d) \leq c_t(d)$.*

proof: Suppose by contradiction that d is the maximal one with $p_t(d) > c_t(d) + \epsilon$, for some small $\epsilon > 0$. Thus, at some point in the iterative process of time t, the price of item d was $c_t(d) + \epsilon/2$, and then some player, j, placed his name on item d, further increasing its price. Let $X[t, ..., M]$ be the tentative allocation at this point, just before j's action. Let us examine the identity of this player j. Notice that any item $d' < d$ has price at most $c_t(d)$, as any player that placed his name on d could have placed his name on d'. We first claim that $Y[l(d), ..., d] \subseteq X[l(d), ..., d]$. Otherwise, fix some $i \in X[l(d), ..., d] \setminus Y[l(d), ..., d]$. If $i \in Y[d+1, M]$ then i placed his name on an item with price strictly larger than $c_t(d) \geq c_t(d+1) \geq p_t(d+1)$ which is larger or equal to the current price of item $d+1$, a contradiction to the myopic behavior of i. If $i \in Y[t, l(d)-1]$ then, by the prefix properties, $d(i) < l(d)$, a contradiction. And if $i \in A_t \setminus Y$ with $d(i) \geq l(d)$ then $v(i) \leq c_t(d)$ by definition, therefore i placed his name on an item with price higher than his value, again a contradiction. Therefore $Y[l(d), ..., d] \subseteq X[l(d), ..., d]$. Now, j places his name on item d. But $j \notin Y[l(d), ..., d]$ as these players are already tentative winners, and $j \notin A_t \setminus Y[l(d), ..., d]$, by repeating exactly the same arguments from above, thus reaching a contradiction. ∎

Sub-claim 8 *$Y = S_t$, and, for any $d \geq t$ and $i = Y[d]$, $c_t(d) = c_t(i)$ (as defined in eq. 2).*

proof: We first show that, for any $j \in A_t \setminus Y$, $Y \setminus i \cup j$ is independent w.r.t items $t, ..., M$ if and only if $d(j) \geq l(d)$. Since $Y[t, ..., l(d)-1]$ is a prefix, any allocation X that contains it cannot allocate an item $\leq l(d)-1$ to player $j \notin Y[t, ..., l(d)-1]$. Therefore $d(j) \leq d$. In the other direction, if $d(j) \leq d$ then we can simply allocate d to player j instead of to i, thus having an allocation to $Y \setminus i \cup j$. Otherwise, $l(d) \leq d(j) \leq d$, and we can use the exact same chain argument of subclaim 6 to obtain an allocation, when replacing i with j.

From this and subclaim 6 we have that, for any $i \in Y$ and $j \in A_t \setminus Y$ such that $Y \setminus i \cup j$ is independent w.r.t items $t, ..., M$, $v(i) \geq p_t(d) \geq c_t(d) \geq v(j)$. This property immediately implies, by the matroid basic properties, that Y is the optimal allocation. By using the above claim again we now get that $c_t(d) = c_t(i)$. ∎

From this last claim it follows that the winner of time t belongs to f_t, as $f_t \subseteq S_t = Y$, and therefore all first $|f_t|$ items of S_t must be sold to the players of f_t. It remains to show that the prices charged by the auction match the criteria of the snapshot planning mechanism.

Sub-claim 9 *In the Online Iterative Auction, the winner i of time t pays $max_{r(i) \leq t' \leq t}\{c_{t'}(i)\}$.*

Let $p_t(i)$ be i's tentative price at time t. Let t' be such that $i = Y[t']$. By the above sub-claims, $p_t(i) = p_t(t') \geq c_t(t') = c_t(i)$. We additionally show that either $p_t(i) = p_{t-1}(i)$ or $p_t(i) = c_t(i)$, and the claim will follow. Assume $p_t(i) \neq p_{t-1}(i)$. Therefore i must have placed his name on item t' during the iterative process of time t. Thus $p_t(t') > p_{t-1}(t')$, and, by the above sub-claims, it follows that $p_t(i) = p_t(t') = c_t(t') = c_t(i)$. ∎

This concludes the proof of claim 16. ∎

We now continue with the proof of the theorem. By claim 16 it follows that the set $R(*)$ of the Online Iterative Auction forms a snapshot planning mechanism, and so the Online Iterative

Auction is an extension of the snapshot planning mechanism. Fix any player i and and combination of recommended strategies of the other players, $s_{-i} \in R_{-i}(*)$. We need to show that i has best response to s_{-i} in $R_i(*)$. Since all players beside i are myopic with tentative deadline and then with final deadline, we can map them to types (v, d) as in claim 16. Let σ be this scenario, where i does not show up at all, and define t^* as in equation 3 of the proof of lemma 5 of section 5.2. Now suppose i plays some strategy $\bar{s}_i$, and denote this scenario by σ'. Let $Y_t(\sigma), Y_t(\sigma')$ be the tentative winners at time t in scenarios σ, σ', respectively. Let t_0 be the first time t such that $f_t(\sigma) \not\subseteq Y_t(\sigma')$. Therefore, for every $t < t_0$, σ' chooses a winner from $f_t(\sigma)$, and by sub-claim 13 we can assume w.l.o.g. that σ and σ' choose the same winner. Therefore $A_{t_0}(\sigma') = A_{t_0}(\sigma) \cup i$. Now suppose that $i = Y_{t_0}(\sigma')[d]$ for some d (if $i \notin Y_{t_0}(\sigma')$ then $d = M + 1$).

Sub-claim 10 *Then $Y_{t_0}(\sigma')[t_0, ..., d-1]$ is independent with respect to items $t_0 + 1, ..., d$.*

proof: By contradiction, let $f \subseteq Y_{t_0}(\sigma')[t_0, ..., d-1]$ be its minimal prefix. Fix any $j \in f_{t_0}(\sigma) \setminus Y_{t_0}(\sigma')$. Since $j \in A_{t_0} \setminus Y_{t_0}(\sigma')$ we have that the tentative price of item t_0 in σ' is at least $v(j)$. By a chain argument as in sub-claim 7 if follows that every $j' \in f$ has value at least than $v(j)$. Since $j \in f_{t_0}(\sigma)$ it then follows that $j' \in Y_{t_0}(\sigma)$. Thus $f \subseteq Y_{t_0}(\sigma)$. Therefore $f_{t_0}(\sigma) = f \subseteq Y_{t_0}(\sigma')$, a contradiction. ∎

From this, and by using again the chain argument of sub-claim 7 we get that d's price is at least $v^*_{t_0}(\sigma) \geq v^*_{t^*}(\sigma)$. As i can win and pay $v^*_{t^*}(\sigma)$ by a strategy in $R_i(*)$ (e.g. arriving at time t^* and bidding only on item t^*), the claim follows. ∎

This concludes the proof of the Theorem. ∎

F Full Proofs: The Sequential Japanese Auction has a Set-Nash Equilibrium

Theorem 6 *The Sequential Japanese Auction is an ignorable extension of the snapshot planning mechanism.*

proof: The theorem will follow by the following claims.

Claim 17 *If all players i play strategies in $R_i(*)$ then the Sequential Japanese Auction forms a snapshot-planning mechanism.*

proof: Let p^* be the last price reached by the auction of time t, and suppose there are k players that did not drop out just before p^* was reached.

Sub-claim 11 *Fix any $j \in A_t \setminus S_t$. As long as j does not drop, then every $i \in P_t(d(j))$ does not drop.*

proof: By contradiction, let $i \in P_t(d(j))$ be the first to drop, say at price p. Since j did not drop, $v(j) \geq p$. By the maximality of S_t, $v(i) > v(j)$. Thus i did not drop because of the price. But the number of non-dropped players is at least $|P_t(d(j))| + 1 > d(i)$. Therefore i could not have dropped at this point, a contradiction. ∎

Sub-claim 12 $p^* = max\{\ v(j) \mid j \in A_t \setminus S_t\ \}$.

proof: Let j^* be the player with maximal value among those in $A_t \setminus S_t$. By the previous sub-claim, j^* will drop because the price will reach his value, as $|P_t(d(j^*))| \geq d(j^*)$. Thus $p^* \geq v(j^*)$. Suppose by contradiction that $p > v(j^*)$, and choose some p in between. Thus, when the price reaches p, all the non-drop-outs belong to S_t. Consider the one that receives, according to S_t, the latest item. The number of non-drop-outs is smaller than his deadline, so he will drop. The one that receives the item before last will next drop, by the same argument, and so on. Therefore the price will not increase beyond p, a contradiction. ∎

Sub-claim 13 *For any $i \in f_t$, $p^* = c_t(i)$.*

proof: For any $j \in A_t \setminus f_t$, $S_t \setminus i \cup j$ is independent: choose an allocation in which i receives item t, and then remove i and allocate t to j. Therefore the claim follows from the previous sub-claim, and from the definition of $c_t(i)$. ∎

Sub-claim 14 *For any l', $|D(p^*, l') \cup \cdots \cup D(p^*, 1)| = l'$.*

proof: Since $D(p^*, l'), ..., D(p^*, 1)$ includes only players that did not actually drop before phase (p^*, l'), and there are exactly l' of those, then $l' \geq |D(p^*, l') \cup \cdots \cup D(p^*, 1)|$. On the other hand, every player among the l' players that did not drop yet will drop in some phase $D(p^*, l'), ..., D(p^*, 1)$, so $l' \geq |D(p^*, l') \cup \cdots \cup D(p^*, 1)|$. ∎

Sub-claim 15 *If $|D(p^*, l' + 1)| = 1$ then $D(p^*, l') \cup \cdots \cup D(p^*, 1)$ is a prefix.*

proof: Since $|D(p^*, l' + 1)| = 1$ then any $j \in D(p^*, l') \cup \cdots \cup D(p^*, 1)$ has deadline $d(j) < t + (l' + 1) - 1$, i.e. $d(j) \leq t + l' - 1$. Since $|D(p^*, l') \cup \cdots \cup D(p^*, 1)| = l'$ it follows from claim 5 that $D(p^*, l') \cup \cdots \cup D(p^*, 1)$ is a prefix. ∎

Sub-claim 16 *Let x^* be the critical number of drop-outs, and $D^* = \cup_{x \leq x^*} D(p^*, x)$, as in def. 14. Then $D^* = f_t$.*

proof: Let $l = |f_t|$. Notice that, for any $l' > l$, $f_t \cap D(p^*, l') = \emptyset$: If $i \in f_t$ then $v(i) > p^*$ and $d(i) \leq |f_t| + t - 1 < l' + t$, so i will not drop. This, in turn, implies that a player in f_t will drop in one the of phases $(p^*, l), ..., (p^*, 1)$, so $f_t \subseteq D(p^*, l) \cup \cdots \cup D(p^*, 1)$. Since $|D(p^*, l) \cup \cdots \cup D(p^*, 1)| = l$, we conclude that $f_t = D(p^*, l) \cup \cdots \cup D(p^*, 1)$. It is left to show that $x^* = l$. As $D(p^*, l) \subseteq f_t$ and $f_t \cap D(p^*, l+1) = \emptyset$ then $D(p^*, l+1) \cap D(p^*, l) = \emptyset$. This implies that $|D(p^*, l+1)| = 1$, so $x^* \leq l$. But if $x^* < l$ then $D(p^*, x^*) \cup \cdots \cup D(p^*, 1) \subsetneq f_t$ is a prefix, contradicting the minimality of f_t (by claims 6, 8). Therefore $x^* = l$ and $D^* = f_t$. ∎

From all the above, the proof of the claim immediately follows: First, the winner belongs to $D^* = f_t$. Second, all time t prices for players not in f_t equal 0, and for players in f_t, time t prices equal $p^* = c_t(i)$, i.e. as required by the price rules of the snapshot-planning mechanism. ∎

We now continue with the proof of the theorem. Using the above claim, it only remains to show that, fixing some player i and some strategies $s_{-i} \in R_{-i}(*)$ of the other players, i has a best response in $R_i(*)$. Consider some strategy s_i of i. Let t_0 be the first time in which i enters D^*. We first notice that, in every time prior to t, i can wave participation without affecting the winner: If the price when i participates reached a level p^*, then clearly, when i does not participate the price cannot rise above p^*. By definition, a player in D^* will not drop before there will be at most $|D^*| - 1$ other non-drop-outs (as the price does not reach his value). Therefore the last non-drop-outs will be exactly all players in D^*, and so the winner will be the same.

Now suppose the price level at time t, in which i entered D^*, is p^*. Therefore i's price will be at least p^*. We claim that, by arriving at time t and playing the fixed confidence strategy $(p^*, 1)$, i can win item t for a price p^*. Since this strategy is in $R_i(*)$, the claim will follow. To see this, observe that $|D(p^*, x)| > 1$ for any $1 < x < x^*$ (since $D(p^*, x) \cap D(p^*, x-1) \neq \emptyset$). Therefore, even if i will not be willing to drop out until being the last non drop out, all others will drop out at price p^*, and so i will win t and will pay p^*. ∎