Online Ascending Auctions for Gradually Expiring Items*

Ron Lavi†

Faculty of Industrial Engineering and Management
Technion – Israel Institute of Technology
ronlavi@ie.technion.ac.il

Noam Nisan

School of Computer Science and Engineering
The Hebrew University, Jerusalem, Israel
noam@cs.huji.ac.il

Abstract

We consider dynamic auction mechanisms for the allocation of multiple items. Items are identical, but have different expiration times, and each item must be allocated before it expires. Buyers are of dynamic nature, and arrive and depart over time. Our goal is to design mechanisms that maximize the social welfare. We begin by showing that dominant-strategy incentive-compatibility cannot be used in this case, since any such auction cannot obtain any constant fraction of the optimal social welfare. We then design two auctions that perform well under a wide class of "semi-myopic" strategies. For every combination of such strategies, the auction is associated with a different algorithm (or allocation rule), and so we have a family of "semi-myopic" algorithms. We show that any algorithm in this family obtains at least 1/3 of the optimal social welfare. We also provide some game-theoretic justification for acting in a semi-myopic way. We suggest a new notion of "Set-Nash" equilibrium, where we cannot pin-point a single best-response strategy, but rather only a set to which best-response strategies belong. We show that our auctions have a Set-Nash equilibrium which is all semi-myopic.

JEL Classification Numbers: C70, D44

Keywords: Multi-Item Ascending Auctions, Robust Mechanism Design, Dynamic Mechanism Design.

---

*An extended abstract of this work appeared in the proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA'05).

†Corresponding author. This author is funded by a Marie-Curie IOF grant from the European Commision for Research.

1 Introduction

Internet auctions exhibit many traditional aspects that are well-studied in auction theory, but are also different in some important aspects. One important difference is their extremely dynamic nature: buyers frequently enter and leave the electronic markets, and items are displayed with attached "expiration times" – a time limit by which the item must be sold. This differs from most classic models of auction theory, where buyers and items are usually static, present throughout the auction.

In recent years, many papers on "online mechanism design" (or "dynamic mechanism design") try to address this new aspect in auctions, for example, Lavi and Nisan (2004), Parkes and Singh (2003), Bergemann and Välimäki (2010), Athey and Segal (2013), and Babaioff, Immorlica and Kleinberg (2007), among many others (Section 1.1 below gives more details). Two standard solution concepts are widely used. Many papers insist on dominant-strategy incentive-compatibility ("truthfulness"), even at the cost of significantly limiting the models being studied. Other papers follow the standard methodology of Bayesian mechanism design, at the cost of making significant assumptions on the distributional knowledge and rationality assumptions that players have about the other players. The dynamic and cultural nature of the Internet raises doubts regarding the validity of such assumptions.

Here, we try a new trail. We design auctions for which *many* strategic choices lead to an approximately optimal allocation. Thus, instead of designing one equilibrium path, aiming to make this single outcome efficient, we design a *family* of outcomes, that correspond to a large set of possible strategic choices. We will predict that players will choose a strategy that belongs to this subset of reasonable strategies. From an algorithmic point of view, we design a family of algorithms, each one corresponds to a specific combination of players' strategies, and *all* of them incur only a bounded welfare loss. Thus, although players are not expected to follow a specific strategy, but only one out of a set of strategies, the outcome is still guaranteed to have bounded efficiency loss. We believe that this general idea of "set equilibria" offers a new way to bypass the inherent difficulties of robust mechanism design, in a way that suits the classic worst-case notions of algorithmic theory.

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, $d(j)$, and 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 "lifetime". Our main motivating example comes from online computational environments, where computational resources need to be allocated to users over time. In particular, our model is equivalent to online scheduling of unit length jobs with deadlines: jobs arrive over time, require one unit of computational resources, and have a deadline for completion. A job that is being completed by its deadline yields a value $v(j)$ for its owner. This is a common, classic computational scenario (Horowitz and Sahni, 1978).

Algorithmic theory offers many solutions for this model, but does not handle any strategic considerations the players/jobs/users might have. A first attempt should probably be to design a truthful mechanism for this problem, in which the dominant strategy of the players is to reveal their true types (arrival time, value, deadline). However, such a mechanism will have to compromise one of our goals, as we show the following impossibility: any truthful deterministic mechanism for our setting cannot always obtain a constant fraction of the optimal welfare.

One could approach this difficulty by adding more assumptions on the players and the environment, but we will try a different approach that builds on the strong connection of our model to assignment problems. In the static case, when all players are present from time 1, our model is a classic assignment model. Many auction formats have been proposed for this model, and it will be especially useful to consider in detail a variant of one of the iterative ascending auctions suggested by Demange, Gale and Sotomayor (1986). Adjusted to our model, this auction constantly maintains a current price p_t and a current winner win_t for every item t. Each player, in 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 for only one item). When none of the players wishes to bid, the auction terminates: each item t is sold to player win_t for a price of p_t. Demange et al. (1986) show that if all players are *myopic*, i.e. always bid on the item with the lowest price, the auction obtains maximal welfare.[1] Furthermore, Gul and Stacchetti (2000) have later shown that behaving myopically is an ex-post equilibrium in this case.

When the dynamic nature of buyers is taken into account, the above results no longer hold, and the situation changes significantly. Clearly, in order to guarantee a constant fraction of the social welfare, items must be sold over time, and the auction process cannot terminate at time 1 with a decision about all future times. As a result, myopic behavior will no longer be an equilibrium. In particular, the best response behavior of a specific player must depend on the player's beliefs about the future. Intuitively, if a player fears that new competitive bidders will arrive in the future, she may bid aggressively for earlier items, offering a higher price for them but reducing her risk of future competition, while if the beliefs suggest that most relevant bidders have already arrived, the bidder will tend to be more myopic.

We give a full analysis of this natural variant of the auction of Demange et al. (1986), adjusted to fit the dynamic setting. We show that even under a wide range of players' strategies, that correspond to different and contradicting beliefs, the resulting social welfare will still be *at least one third* of the optimal social welfare. This holds regardless of the number of items, the number of players, the range of player values, and even if an adversary sets players' types so as to intentionally "fail" the auction.

This type of "worst-case" analysis is common in computer science and algorithmic theory, where it is standard practice to avoid the assumption that the probability distribution over the input to the algorithm is known to the designer. In connection with mechanism design and auction theory, worst-case analysis leads more easily to constructions of detail-free mechanisms. Since the analysis itself is conducted without relying on any distributional details, the resulting mechanisms are naturally detail-free. A worst-case analysis enables us to adapt the classic DGS mechanism for the dynamic case, showing its robustness to a shift from static environments to dynamic ones. Bayesian analysis, on the other hand, not only relies on the assumption that the designer knows the underlying distribution and can tailor the mechanism to the specifics of that distribution, it usually also requires the stronger assumption that all players (and not only the designer) commonly agree about the details of the underlying distribution. This assumption seems too strong for our setting.

[1]This is the "approximate auction", the second of the two mechanisms in Demange et al. (1986). The welfare obtained by this auction approaches the optimal welfare as δ goes to zero.

The auction of Demange et al. (1986) is adjusted to the online setting by assuming that at time 1, when the iterative process ends, only item 1 is sold. Then, at time 2, new players may join and the iterative process resumes, where prices start at their previous level. This continues at any time t. As previously mentioned, myopic strategies are no longer an equilibrium, as for example a player that estimates that new competitive bidders will arrive in the future may bid more aggressively for earlier items. To incorporate such considerations, we call a player *semi-myopic* if she always bids on some item with price lower than her value, but not necessarily on the item with the lowest price, as the myopic behavior requires. Thus, semi-myopic behavior is a much weaker assumption than myopic behavior: it does not specify a specific item to bid on, but rather allows to choose any item that could potentially result in a positive utility. The main requirement is only that a player will not be silent as long as there exist potentially beneficial items. This captures a wide range of strategies, that reflect different and contradicting beliefs. Our analysis shows that the Demange et al. auction, in the online (dynamic) setting, always obtains at least one third of the optimal social welfare, as long as all players are semi-myopic. Thus, even if the prior beliefs of the players are significantly different, and the "true" underlying distribution is not common knowledge, this auction mechanism will enable social coordination to some reasonable extent.

We additionally show that a second classic auction format yields similar results. This is the sequential Japanese auction: item t is sold at time t using a one-item ascending auction. We show a strategic equivalence, in our setting, between this auction and the auction of Demange et al. This equivalence enables us to properly define a myopic behavior for the sequential Japanese auction, which leads to the optimal allocation and is an ex-post equilibrium, in the offline case. Similarly to above, this equivalence also leads to a family of semi-myopic strategies, that capture players' uncertainties about the future, in the online setting. Our analysis again shows that every choice of semi-myopic strategy will always obtain at least one third of the optimal social welfare.

Why should the players be semi-myopic? Demange et al. in their original paper were satisfied to simply assume myopic behavior, and it seems to us that semi-myopic behavior is a natural extension of myopic behavior to dynamic settings. Still, it is desirable to approach this question also from a game-theoretic point of view. We seek a notion that will capture the idea that, without any knowledge about the future, we have to forecast a *set* R_i of strategies, instead of a single strategy r_i as an equilibrium point. We say that the strategy sets R_i are in a "Set-Nash equilibrium" if for any player i, and any strategy combination of the other players $s_{-i} \in R_{-i}$, player i has a best response to s_{-i} in R_i. We show that both our online ascending auctions have a Set-Nash equilibrium with strategies that are all semi-myopic. In the paper body we compare the set-Nash notion to other existing notions, for example to the "curb set" of Basu and Weibull (1991), and discuss the differences. We also provide some discussion on ways to strengthen the basic definition. We describe a hierarchy of four "set equilibria" notions, with growing strength. While, for our motivating problem, we were able to use only the basic definition, we believe that the complete hierarchy may turn out useful for other models, where robust implementation is impossible, but one still wishes to construct detail-free mechanisms and avoid unrealistic distributional assumptions.

The welfare loss caused by the transition from the offline to the online setting can be attributed to two effects: *the online effect*, which is the loss of welfare caused by the dynamic setting, since the social planner does not know the future (even if players' types were fully known), and *the strategic effect*, which is the loss of welfare caused by the players as they employ more complex strategic behavior in the presence of extreme uncertainty. Hajek (2001) shows that no algorithm for the worst-case online setting can obtain more than 62% of the optimal welfare, thus providing an estimate for the loss of welfare caused by the online effect.[2] Our results demonstrate that the additional loss due to the strategic effect is at most 30%.

Replacing the Bayesian analysis with a distribution-free analysis can be done in two ways. The first possibility is to still assume that players' types are drawn from a fixed distribution, but to use this fact only in the analysis itself. The description of the mechanism will not rely on the knowledge of the distribution (or will only partly rely on it), and the performance guarantees will hold for *any* distribution (or at least for any distribution out of a large class of distributions). This is the approach taken e.g. by Satterthwaite and Williams (2002) in the study of two sided auctions, or by McAfee (2002) in the study of market rationing. A stricter approach, which is usually the choice in ex-post implementations, is to avoid any distributional assumptions altogether. For example, the arguments of Demange et al. (1986) hold *even if players' types are chosen by an adversary*, or, in other words, a worst-case analysis. This latter worst-case approach is the one we take in this study. The former approach would have also been an interesting line of investigation in the context of our model, and we believe that the results under such assumptions would have been even tighter. In a follow-up paper, Compte, Lavi and Segev (2012) attempt such an approach, for a special case of the model that we study here.

The remainder of this paper is organized as follows. After a brief survey of related literature in Section 1.1, we formally define our model in Section 2. Section 3 describes the dynamic version of the two classic auctions, and analyzes the resulting efficiency of semi-myopic strategies. Section 4 shows that truthful deterministic auctions that guarantee a constant fraction of the optimal welfare are impossible. Section 5 describes the notion of Set-Nash that we suggest, and Section 6 shows that our two auctions indeed have a Set-Nash equilibrium which is all semi-myopic. Section 7 concludes. Additional supplementary technical details are given in the appendices.

1.1 Related Literature

Models of online auctions were studied in the context of algorithmic mechanism design, starting with Lavi and Nisan (2004). Examples for various different models can be found in the book chapter by Parkes (2007). Following the conference version of this paper (Lavi and Nisan, 2005), Hajiaghayi, Kleinberg, Mahdian and Parkes (2005) design a truthful auction for our setting that guarantees one half of the optimal welfare in our setting, assuming that players cannot declare deadlines larger than their true deadlines, and that prices are charged only after all auctions end (and depend on the entire sequence of auctions). Cole, Dobzinski and Fleischer (2008) show to keep the same efficiency guarantee of one half, while charging payments at purchase time, if both arrival times and deadlines are known to the auctioneer and cannot be manipulated by the players.

More general models of dynamic mechanism design have been studied by Parkes and Singh (2003), Athey and Segal (2013), Bergemann and Välimäki (2010), and Cavallo, Parkes and Singh (2007). They study a general multi-period allocation model, in which a designer needs to perform allocation decisions in each period and players have private values for the different allocations, that may be stochastic and time-dependent. Said (2012) shows an explicit connection between the general results of Bergemann and Välimäki (2010) and the model of sequential ascending auctions, similar to the one studied here. These results show that full efficiency can be achieved under the assumption that the mechanism has correct information about future arrivals. In Bergemann and

[2]He was only able to provide an allocation rule that achieves 50% of the welfare, though, and this gap is still an open question.

Välimäki (2010) this is achieved by assuming that players report, before they arrive, the probability of their future arrival, while Said (2012) abstracts this from the mechanism by assuming a common-prior.

Gershkov and Moldovanu (2009) study a setting where players appear according to some fixed and known stochastic process (e.g. a Poisson arrival rate), and are impatient, i.e. must either be served upon their arrival, or not be served at all. This structure enables Gershkov and Moldovanu (2009) to characterize optimal dynamic mechanisms, with respect to both the social welfare and the seller's revenue.

The algorithmic problem of online job scheduling with unit length jobs was studied by Kesselman, Lotker, Mansour, Patt-Shamir, Schieber and Sviridenko (2004). The give a simple greedy algorithm that guarantees one half of the optimal social welfare. Bartal, Chin, Chrobak, Fung, Jawor, Lavi, Sgall and Tichý (2004), give a randomized algorithm that guarantees a better bound of about 65%. Hajek (2001) shows that no deterministic algorithm can obtain a bound better than the "golden ratio".

Compte et al. (2012) study a special case of the model studied here, where buyers do not have deadlines, and are able to stay until all items are sold. They study the efficiency of a sequence of English auctions with a certain "early termination" rule, and show that this mechanism obtains in the worst-case at least 63% of the optimal welfare.

2 Model and Basic Definitions

Items: A seller sells M identical items with different expiration times. 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 (players/bidders) of the items 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 departure time, or deadline, $d(i)$. We assume that the arrival and departure times are integers (actions in a non-integral time point can be deferred to the next integral point with no effect). 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. In fact, the auctions we design also satisfy the more severe restriction that player i cannot receive an item t that expires after his deadline ($t > d(i)$). 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, as this is a matter of tie-breaking.[3]

We assume the private value model with quasi-linear utilities: player i privately obtains his variables $r(i)$, $d(i)$, and $v(i)$, and acts rationally in order to maximize his own utility: his obtained value minus his price. A player may arrive at or after his true arrival time, and declare or act as if he has any value, and any deadline.

We defer questions about the exact knowledge of the players, besides their own private parameters, until Section 5 below, where we analyze the strategic behavior.

Our goal: We aim to design allocation mechanisms that maximize the social welfare: the sum of (true) values of players that receive an item.

Basic notations: 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

[3] 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 $i \succ j$.

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 $v(X_t) = \sum_{d=t}^{M} v(X_t[d])$. In the static (offline) problem, in which all players arrive at time 1, the potential sets of winners constitute a "matroid". All properties of matroids that we require for our proofs, including the basic definition, are surveyed in Appendix A. Readers who are unfamiliar with this combinatorial structure are strongly advised to read that section first.

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 close to optimal allocation is obtained for a large, reasonable family of strategies that we term "semi-myopic". Under any such player behavior, each of our auctions belongs to a general family of semi-myopic allocation rules, that we characterize. We then show that any semi-myopic allocation rule obtains at least one third of the optimal welfare, and therefore conclude that our auctions lead to a near optimal allocation for any choice of semi-myopic strategies of the players.

In this section, we focus on the quality of allocations that the auctions achieve. Therefore we give only intuitive justifications for the player behavior that we assume. For the same reason, we also omit a few technicalities about prices and tie-breaking rules from the definitions. All these are detailed and handled with care when we analyze the strategic properties of our auctions, in the next sections.

3.1 The Online Iterative Auction

We consider an online adaptation of the iterative auction of Demange et al. (1986):

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 $t' \geq t$, 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 his name 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 a special case of the unit-demand model studied by Demange et al. (1986), Gul and Stacchetti (2000):

Definition 2 (Demange et al. (1986)) *Player i has a* **myopic strategy** *in the iterative auction if, in his turn, he always places his name on the an item $t \leq d(i)$ with the minimal price among all items. If the minimal price is at least $v(i)$, he does not bid at all.*

Lemma 1 (Demange et al. (1986), Gul and Stacchetti (2000)) *If all players are myopic and arrive at time 1 then the online iterative auction obtains the optimal allocation. Furthermore, if all other players are myopic then player i will maximize his utility by playing myopically.*

In the online setting, however, 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 bids on some item t with $p(t) \leq v(i)$ and $r(i) \leq t \leq d(i)$ (not necessarily the one with the lowest price). If there is no such item, i stops participating.*

Theorem 1 *If all players are semi-myopic then the online iterative auction achieves at least one third of the optimal welfare, up to an additive loss that tends to zero as $\delta \to 0$. More specifically:*

$$v(OPT) \leq 3 \cdot v(ON) + 2 \cdot M \cdot \delta,$$

where OPT, ON are the optimal, online allocations, respectively.

The proof is given in Section 3.3 below, where we show that, under any semi-myopic strategy, the online iterative auction follows a "semi-myopic allocation rule" (as defined in Section 3.3 below), hence obtains the desired welfare level.

3.2 The Sequential Japanese Auction

A different possibility is to sell item t at time t using a simple one item ascending auction. In a Japanese auction for a single item, a price is continuously rising, and the only action that a bidder may take is to drop out of the auction. Once the bidder drops out, he may not re-enter. When the second-to-last bidder drops out, the price ascent stops, and the last bidder to remain wins the object at that price. A natural adaptation of this auction to the online case is as follows:

Definition 4 (The Sequential Japanese Auction (intuitive version)) *The Sequential Japanese Auction sells each item t at time t, separately, using a modified Japanese auction: the participants are allowed to observe how many drop-outs occur as the price ascends and to incorporate this into their drop-out decision.*[4]

[4]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 the players to drop simultaneously. See Section 6.3.

The complete definition is deferred to Section 6.3. As before, it is useful to first consider this auction in the offline case, in which a rather surprising notion of myopic behavior leads to the optimal allocation:

Definition 5 *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.*

The logic for dropping when $d(i) - t$ 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.

Lemma 2 *If all players are myopic and arrive at time 1 then the Sequential Japanese Auction obtains the optimal allocation.* [5]

In the online setting, again, 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. All we wish is that players will not drop out "too soon":

Definition 6 *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 obtains at least one third of the optimal welfare.*

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 results in a semi-myopic allocation rule. The proof is given is Section 3.3 below.

3.3 Semi-Myopic Allocation Rules

For each combination of player strategies, the above auctions are associated with a different allocation rule. In order to analyze their performance for a family of strategies, we therefore need to characterize a family of allocation rules, that we call semi-myopic allocation rules. The main point is that *any* semi-myopic allocation rule obtains at least one third of the optimal welfare.

Specifically, fix some time 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. In Appendix A we show that these independent sets are the independent sets of a matroid (Appendix A contains all necessary background on this combinatorial structure). The *current best schedule* at time t, S_t, is the allocation with maximal value among all allocations of items $t, ..., M$ to the active players, A_t.[6] Define

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

[5] The assumption that players have different values is important here. It is not hard to verify that this lemma is actually a special case of Theorem 7 from the online strategic setting (specifically, it follows from Lemma 7). We note that myopic strategies in the offline case form an ex-post equilibrium only when using the modified prices given in Section 6.3.

[6] There exists one such allocation, by the matroid structure, and since different players have different values. See Appendix A for details.

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 (i.e. these are all the potentially *first* players). Now define the critical value at time t, v_t^*, as:

$$v_t^* = \begin{cases} 0 & S_t \text{ is independent w.r.t. items } t+1, ..., M \\ \min_{j \in f_t}\{v(j)\} & \text{otherwise} \end{cases}$$

All active players with value larger than v_t^* must belong to S_t, because of its optimality. To see this, note we can assume without loss of generality that the player in S_t that gets item t has value exactly v_t^*. Thus, if there was a higher valued player outside of S_t, we could give him item t and strictly increase the value of S_t, which is a contradiction to the optimality of S_t.

Because of this, it seems reasonable not to allocate item t to a player with value less than v_t^*, as this player cannot belong to any optimal allocation. Surprisingly, this condition is enough to obtain approximately optimal allocations:

Definition 7 (A semi-myopic allocation rule) *An allocation rule is semi-myopic if every item t is sold at time t to some player j with $v(j) \geq v_t^*$. If $v_t^* = 0$, the allocation rule may choose to keep the item unallocated.*

Lemma 3 *The Online Iterative Auction with semi-myopic players and the Sequential Japanese Auction with semi-myopic players are both semi-myopic allocation rules.*[7]

Proof: We first show the claim for the Online Iterative Auction. 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 11 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. This concludes the claim for the Online Iterative Auction.

For the Sequential Japanese Auction, we show that the winner has value at least v_t^*. 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 6 in Appendix A, $d(j) - t + 1 \leq |f_t|$. 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 and has value at least v_t^* by definition. ∎

The family of semi-myopic allocation rules can be viewed as the entire range between the following two extremes: the first is the greedy allocation rule, that always chooses the player with maximal value[8], and the second is the "myopic" allocation rule that always chooses the player that determined v_t^*. These two extremes always produce an allocation with welfare at least half of the optimal welfare (both were studied in the context of online scheduling, see e.g. Kesselman et al. (2004)). The entire family has only a slightly smaller performance guarantee:

[7]For the online iterative auction, we actually show that $v(win_t) \geq v_t^* - \delta$.

[8]Interestingly, this is a special case of the greedy algorithm of Lehmann, Lehmann and Nisan (2006) for combinatorial auctions with sub-modular valuations. They study the offline case, but it is easy to verify that their algorithm actually works online.

Theorem 3 *Any semi-myopic allocation rule obtains at least one third of the optimal welfare, and this bound is tight.*

The proof of this theorem is given in Appendix C. Note that Theorem 3, coupled with Lemma 3, is a proof for Theorems 1 and 2. The following example shows that the one-third guarantee is tight:

Example 1 *Consider the following scenario for three items. At time 1 arrive two players, j_1 has value 0 and deadline 1 and j_2 has value 1 and deadline 2. Thus, $v_1^* = 0$ and the online allocation rule may allocate 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 allocation rule may choose j_4. At time 3 no new players arrive, so item 3 remains unallocated. The resulting welfare is therefore 1, while the optimal welfare is 3.*

4 The Impossibility of Truthful Implementation

We now move to game-theoretic considerations, and start with an impossibility. Let T_i be the domain of all valid player types $(r(i), v(i), d(i))$, and let $T_{-i} = \times_{j \neq i} T_j$. By the revelation principle, it is enough to consider direct revelation mechanisms. Consider the allocation constructed by the mechanism 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 8 (Truthfulness) *A mechanism is truthful if there exist price functions $p_i : T_1 \times \cdots \times T_n \rightarrow \Re$ such that, for any i, any $b_{-i} \in T_{-i}$, any true type $b_i \in T_i$, and any $\tilde{b}_i \neq b_i$[9]:*

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

Theorem 4 *Any truthful deterministic mechanism for our online allocation problem cannot always obtain more than $1/M$ fraction of the optimal welfare (where M is the number of items).*

Remark 1: Although the proof below utilizes an extreme scenario with players with very large values, the worst-case ratio presented by the proof occurs in common, simple scenarios. In other words, the proof demonstrates that, since the mechanism defends itself against such extremes, it must make wrong decisions even in simple cases.

Remark 2: There exists a simple truthful deterministic mechanism that always obtains at least $1/M$ fraction of the optimal welfare: for any player i, set p_i to be the highest bid received in time slots $1, ..., t$, excluding i's own bid. Sell item t to player i if and only if $v(i) > p_i$, for a price of p_i. It is an easy exercise to verify that truthful-reporting is the only dominant strategy for this mechanism, and, since the player with the highest value always wins, at least a $1/M$ fraction of the optimal welfare is obtained.

Proof of Theorem 4: 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.

[9] We actually restrict the possible $\tilde{b}_i$'s such that $\tilde{r}(i) \geq r(i)$.

Claim 1 *Fix some truthful deterministic mechanism that always obtains at least $\frac{1}{c}$ fraction of the optimal welfare, for some fixed $c \geq 1$. 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 of Claim 1: Fix any combination of players that arrive at time 1, b_{-i}. Suppose first that i has deadline equal to 1. For this case, the player's type space becomes single dimensional, hence by well-known incentive compatibility arguments (for example Myerson (1981)) there exist a price function as claimed.[10]

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 mechanism not to allocate any item to i in order to maintain a fraction of the optimal welfare.[11] Therefore i must 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 a deadline of 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 the 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 mechanism can obtain welfare of at most $1+\epsilon$, while the optimal welfare is at least M, and the theorem follows. To verify that such types exist, fix $L > M$ real values in $[1, 1+\epsilon]$. Choose M values $v(i)$ uniformly at random from these L values. Then, for any given i, $Pr(v(i) = p_i(b(-i))) \leq 1/L$, as the values were drawn i.i.d. Thus, $Pr(\exists i,\ v(i) = p_i(b(-i))) \leq M/L < 1$, hence there exist a choice of values with $v(i) \neq p_i(b(-i))$ for all i. ∎

5 A Game-Theoretic Framework

Our main motivation at this point is to justify the assumption that players will behave semi-myopically. We desire a rational justification, i.e. one that shows that expected strategies are, in

[10]For the single dimensional case a function is implementable if and only if the winning probability weakly increases with the player's value (keeping the other values fixed). When the mechanism is deterministic, this essentially boils down to the fact that 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 loses and pays 0 if $v(i) < p_i^*$.

[11]More specifically, at any time point $t = 2, ..., M$, let x be the sum of values of all players that arrived at previous times $1, ..., t-1$. Then, at time t there will arrive a single player j, with value $(c+1)x$ and deadline t. If the mechanism will not assign item t to player j, no additional players will arrive in time slots $t+1, ..., M$, and the total social welfare of the mechanism can be at most x while the optimal social welfare is at least $(c+1)x$. Thus, the mechanism will fail to obtain a 1/c fraction of the optimal social welfare, and this is a contradiction. In other words, the mechanism must assign item t to player j for every item $t = 2, ..., M$, and player i will not receive any item.

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

We first describe the set equilibrium notions for games with complete information, and then explain how to extend them to a setting of incomplete information, which suits our needs here. There are n players, where each player i has a strategy space S_i. The outcome of the game is given by the n utility functions $u_i : S \to \Re$ where $u_i(s_i, s_{-i})$ denotes i's payoff when he plays strategy s_i and the others play the strategy tuple s_{-i}. The basic assumption is that, given that the other players play s_{-i}, player i will choose a strategy $s_i \in \text{argmax}\{u_i(s_i, s_{-i})\}$.

In our setting, a set of recommended strategies, R_i, is defined for each player i. The motivating scenario is where it is known that if all players play recommended strategies then the outcome is "good" in some sense. E.g., in our case, the obtained social welfare approximates the optimal one (therefore we do not put any emphasis on the minimality of the sets; see the discussion on related literature below for details). We would like to capture the notion that the sets R_i are in equilibrium. In other words, we formalize when can it be said that given that other players $j \neq i$ all play strategies in R_j, then player i also rationally plays some strategy in R_i.

We give four definitions below, all maintain the spirit of this "set equilibrium" notion, in order of increasing strength. Some of these notions have been defined before in the literature in the context of complete information games – we discuss this below in Section 5.1.1. All of the following definitions behave the same on the two extreme cases: When each R_i is a singleton set ($\forall i, |R_i| = 1$) then they are equivalent to Nash equilibrium. When R_i is the entire strategy space ($R_i = S_i$) then they are trivially satisfied.

Definition 9

1. *We say that R_i are in "Set-Nash equilibrium" (in the pure sense) if for every i, every $s_{-i} \in R_{-i}$, and every $s_i \in S_i$ there exists $r_i \in R_i$ such that $u_i(r_i, s_{-i}) \geq u_i(s_i, s_{-i})$. I.e., for every tuple of recommended strategies there exists a best response strategy in the recommended set.*

2. *We say that R_i are in "Set-Nash equilibrium" (in the mixed sense) if for every i, for every tuple of distributions π_j on R_j for all $j \neq i$, and every $s_i \in S_i$ there exists $r_i \in R_i$ such that $E_{\{\pi_j\}_{j \neq i}}[u_i(r_i, s_{-i})] \geq E_{\{\pi_j\}_{j \neq i}}[u_i(s_i, s_{-i})]$. I.e for every series of distributions on the recommended strategies of the other players there exists a best response in the recommended set. This definition captures an expected-utility scenario,* over all possible priors.

3. *We say that $\{R_i(\cdot)\}$ are in "Set-Nash equilibrium" (in the mixed-correlated sense) if for every i, for every π on $s_{-i} \in R_{-i}$, and every $s_i \in S_i$, there exists $r_i \in R_i$ such that $E_\pi[u_i(r_i, s_{-i})] \geq E_\pi[u_i(s_i, s_{-i})]$. This definition extends the previous one in the sense of allowing the other players to correlate strategies.*

4. *We say that R_i are in "Set-Domination equilibrium" if for every i, and every $s_i \in S_i$ there exists $r_i \in R_i$ such that for every $s_{-i} \in R_{-i}$, we have that $u_i(r_i, s_{-i}) \geq u_i(s_i, s_{-i})$. I.e. for every unrecommended strategy, there is a recommended strategy that is not worse-off, as long as others act as recommended.*

These definitions extend to games with incomplete information in a straightforward way. Each player i has a privately known type $t_i \in T_i$. No probability distribution is assumed on $T = T_1 \times ... \times T_n$. The utility functions now depend on the player's type, as well ($u_i : T_i \times S \to \Re$, 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 the others play the strategy tuple s_{-i}). The set of recommended strategies may now depend on the player's type, i.e., $R_i : T_i \to 2^{S_i}$. We denote $R_i(*) = \cup_{t_i \in T_i} R_i(t_i)$. All four definitions are modified so that the condition specified should now hold for all possible types t_i. In addition, the recommended sets of the other players are always taken to be $R_{-i}(*)$, while the best response r_i must exist in player i's recommended set *according to his true type*, $R_i(t_i)$. Thus, the requirement holds for all possible type realizations of the other players. For example, the first definition is altered so that the set functions $R_i(\cdot)$ are in "Set-Nash equilibrium" (in the pure sense) if for every i, every 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 $u_i(t_i, r_i, s_{-i}) \geq u_i(t_i, s_i, s_{-i})$.

In all definitions, 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 (a probability distribution on $R_i(t_i)$). It is easy to verify that this does not change the first three definitions (the best mixed strategy is always a pure one), while for the Set-Domination definition, this will weaken it to become equivalent to Set-Nash for correlated strategies (using von-Neuman's max-min principle).

The first three definitions 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 players' 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 work we use the basic definition (and drop the qualifier "in the pure sense" hereafter). In addition, all our Set-Nash strategies are undominated, and they are in fact best responses to a set of acceptable strategies wider than $R_{-i}(*)$.

5.1.1 Related notions in the Game-Theory literature

The game theory literature defines and discusses similar notions to the above set equilibria notions. Most of the works handle games with complete information, and investigate the existence and uniqueness of minimal such equilibria. We are not aware of any such study in the setting of incomplete information, where the analysis is performed in the context of mechanism design and implementation theory, where the equilibria are evaluated with respect to the quality of the outcome they yield.

Shapley (1964) defines a notion of "a saddle" for two-person zero-sum games, which is almost the same as the Set-Domination notion (but the inequalities there are strict). Shapley shows that there always exists a unique minimal saddle in a zero-sum game (the strictness of the inequalities is crucial for this), but does not address the quality of the obtained outcome. Duggan and Breton (1996, 2001) define a “mixed saddle”, which allows mixed strategies in the definition. As we note above, this is actually equivalent to the definition of Set-Nash in the correlated sense. Their results are again for the complete information case (mainly for zero-sum games, and for voting procedures). Duggan and Breton (1998) develop a general approach to construct “choice sets”. They require both an “outer stability”, which resembles our logic of constructing a set equilibrium, and also require an “inner stability”, in order to have a minimal choice set. We replace this inner stability with a requirement on the quality of the outcome. This of-course can be done in our context of implementation theory, but not in their context of normal form games with complete information. Basu and Weibull (1991) study sets of strategies that contain all their best replies (a “curb” set), a rather strong notion, and Voorneveld (2004) defines a “prep-set”, which is equivalent to our definition of Set-Nash in the mixed sense. Both works study the existence of minimal such sets in games with complete information.

Although rationalizability (Bernheim (1984); Pearce (1984)) is not perceived as an equilibrium concept, the motivation behind the definition is quite similar to ours. Indeed, this notion was successfully used to analyze first price auctions in a detail-free setting (Dekel and Wolinsky (2003); Battigalli and Siniscalchi (2003)), another example of an analysis in the context of mechanism design and implementation. It is also interesting to parallel the shift from “rationalizability” to “point rationalizability”, which Bernheim (1984) makes, to the shift from Set-Nash in the mixed sense to Set-Nash in the pure sense, that we make.

We would like to note the difference between these notions of set-equilibrium, and the analysis of “sets of Nash equilibria”. The latter analysis deals with sets of Nash equilibria, e.g. in order to determine the stability properties of an equilibrium point in the set (as in Kohlberg and Mertens (1986)), while the notions of set-equilibria are aimed to capture situations in which single-strategy tuples do not form an equilibrium at all.

5.2 Implementation in Set-Nash equilibrium

As our context is the framework of implementation theory, we wish to formally specify how the notion of Set-Nash equilibria fits in, in parallel to classical results. We do this for the basic definition of Set-Nash, but the entire discussion follows through for all four definitions in an immediate way. The setting contains a set of outcomes/alternatives, A, from which we have to choose one outcome. The choice depends on the players' types $t \in T$, according to some social choice correspondence $F : T \rightarrow 2^A$. In our example, A is the set of all valid allocations of items to players, and $F(t)$ outputs all allocations that have a social welfare of at least one third of the optimal social welfare with respect to the type t. This social correspondence represents the fact that our goal is to obtain a close-to-optimal welfare, and any allocation that obtains this will satisfy us. All the classic definitions from implementation theory can be adapted to our Set-Nash definition:

Definition 10 *Given $F : T \rightarrow 2^A$, an implementation in Set-Nash equilibrium is a mechanism with strategy sets $S_1, ..., S_n$, and an outcome function $g(s_1, ..., s_n) \in A$, such that there exists a Set-Nash equilibrium $\{R_i(\cdot)\}_i$ that satisfies that $g(s) \in F(t)$ for* **all** *$s \in R(t)$.*

Notice that we cannot hope to require that *all* equilibria will produce results according to F, as there always exists the trivial set-equilibrium that contains all strategies.

Definition 11 *A social choice correspondence $F : T \to 2^A$ is a c-approximation to the social welfare if for any $t \in T$ and any outcome $a \in F(t)$, the social welfare obtained in a is at least a $1/c$ fraction of the optimal social welfare with respect to t.*

Thus, our goal is to show that our two online auctions Set-Nash implement a 3-approximation of the social welfare.

The celebrated revelation principle states that whenever we can implement a social function in some equilibrium, we can also implement it using a direct revelation implementation, in which the strategy space of the players is simply to reveal their type. For our "set equilibrium" notion, we can have an "extended direct revelation" implementation which is "extended truthful":

Definition 12 *An implementation is an "extended direct revelation implementation" if there exist surjective functions $h_i : S_i \to T_i$ for each player i. It is "extended truthful (in Set-Nash equilibrium)" if there exists a Set-Nash equilibrium in which $h_i(R_i(t_i)) = \{t_i\}$ for all i and all $t_i \in T_i$.*

In other words, an extended direct revelation implementation has strategies that are of the form (t_i, l_i), where $t_i \in T_i$, and l_i represents some additional information. It is extended truthful (in Set-Nash equilibrium) if there exists a Set-Nash equilibrium in which $R_i(t_i) = (t_i, *)$, i.e., the player declares his true type in every one of his recommended strategies. In the context of our model, players may wish to state a deadline which is earlier than their true deadline, as they do not wish to be tentatively placed in a slot that is "too far in the future". An extended direct revelation implementation will enable the players to declare their true deadline and a desired deadline which may be closer than their true deadline. Such a mechanism is extended truthful in Set-Nash equilibrium if every recommended strategy reports the true deadline of the player, along with additional deadlines.

Proposition 1 (An extended revelation principle) *Every function $F : T \to 2^A$ that can be implemented in Set-Nash equilibrium can be implemented by an extended truthful implementation.*

Proof: Given an implementation M of F in Set-Nash equilibrium, we build an extended truthful implementation M', that encapsulates M, as follows. Let $R_i(t_i)$ be the recommended strategies of M. Then the strategy space of a player in M' is to specify his type t_i, and a strategy in $R_i(t_i)$. The mechanism then uses M' with the specified strategies to determine the result. It is immediate to verify that the sets $R'_i(t_i) = \{(t_i, s_i) \mid s_i \in R_i(t_i)\ \}$ are indeed a Set-Nash that fits the definition. ∎

It is worth pointing out that our auctions, which are not direct revelation, have an interesting extended direct revelation counterpart – we describe this in Section 6.1 below.

5.3 Ignorable Extensions of Games

This section formalizes a concept used in the proof of our main theorem, below. In the proof, we first describe an extended truthful mechanism that implements a 3-approximation to the maximal social welfare, and then show that each of our ascending auctions has "inside" it a semi-myopic mechanism. In this section, we describe this type of building block more generally.

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 13 *$\bar{G}$ is an* **ignorable extension** *of G if the recommended sets $R_i(t_i) = S_i$ (for all i) form a Set-Nash equilibrium of $\bar{G}$.*

In other words, $\bar{G}$ 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 all other players play original strategies (from S_{-i}), then player i has an original strategy (from S_i) which is a best response.

Proposition 2 *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 $\bar{G}$.*

We point out that, although these notions are related to the notion of Set-Nash equilibrium in the pure sense, there is in an immediate, similar way to define ignorable extensions to any one of the other three definitions of Set-Nash equilibria.

6 A Strategic Analysis of our Auctions

The strategic analysis of our auctions is performed in two parts. First (in Section 6.1), we describe an extended direct revelation auction which we call the "Semi-Myopic Mechanism", and show that this mechanism has a Set-Nash equilibrium which is composed of semi-myopic strategies. Thus this mechanism always obtains (in equilibrium) at least one third of the optimal welfare. Second (Sections 6.2 and 6.3), we show that this semi-myopic mechanism is "embedded" inside both our ascending auctions, in the exact sense described in Section 5.3 above. This implies that both our ascending auctions also have a Set-Nash equilibrium which is composed of semi-myopic strategies.

6.1 Semi-Myopic Mechanisms

We now devise an extended direct revelation auction with our two basic building blocks: it has a Set-Nash equilibrium, and, for these equilibrium strategies, the auction is a semi-myopic allocation rule.

Definition 14 *We define a family of* **semi-myopic mechanisms** *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 7, 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 $\max_{r(i) \leq t' \leq t} p_{t'}(i)$.

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 4 *When all players play recommended strategies according to their true types then the allocation of any semi-myopic mechanism is a semi-myopic allocation rule.*

The proof of this lemma is given in Appendix B.

Theorem 5 *The semi-myopic mechanism Set-Nash implements a 3-approximation of the welfare.*

The proof of this Theorem is given in Appendix D. In Appendix D.2 we show by an example that the recommended strategies of the semi-myopic mechanism do not contain best responses to mixed-correlated strategies. Hence, unfortunately, the semi-myopic mechanism does not have a semi-myopic Set-Nash equilibrium in the mixed-correlated sense. An interesting problem that we leave open is to devise a mechanism that Set-Nash implements (in the mixed sense) some constant approximation of the welfare.

6.2 The Online Iterative Auction

We now show that our Online Iterative Auction is an ignorable extension of a semi-myopic mechanism, thus having a Set-Nash equilibrium which approximates the welfare, according to Theorem 5. For this, we need to refine our intuitive definition:

Definition 15 (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 6 *The Online Iterative Auction is an ignorable extension of a semi-myopic mechanism.*

The proof of this Theorem is given in Appendix D.

Corollary 1 *The Online Iterative Auction Set-Nash implements a 3-approximation of the welfare.*

6.3 The Sequential Japanese Auction

To show that our Sequential Japanese Auction is an ignorable extension of a semi-myopic 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 16 (The Sequential Japanese Auction) *The basic auction structure remains the same as in Def 4. 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.*[12] *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 7 *The Sequential Japanese Auction is an ignorable extension of a semi-myopic mechanism.*

The proof of this Theorem is given in Appendix D.

Corollary 2 *The Sequential Japanese Auction Set-Nash implements a 3-approximation of the welfare.*

7 Conclusions

In this paper we have analyzed two auction structures, common both in theory and in practice, in a dynamic online setting. While, for the two auctions in the offline case, a myopic behavior leads to the optimal allocation and is in addition an ex-post equilibrium, in the online case the situation is more complex. We have focused on a detail-free worst-case analysis, showing that replacing myopic behavior with the weaker notion of semi-myopic behavior grants much freedom to the players, on the one hand, and reduces the social welfare only by a constant factor, on the

[12] 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 $i \notin D(p, n-1)$.

other hand. This notion of semi-myopic behavior encompasses a large range of player strategies, representing different and contradicting beliefs. Our results therefore show the relative robustness of the two auction formats to settings with extreme uncertainties. From a game-theoretic point of view, we have shown that there exists a "Set-Nash" equilibrium, which is all semi-myopic, in both auctions. According to this equilibrium notion, players are not expected to choose a single tuple of strategies, but rather one strategy out of a set of strategies. In our setting, we show that every strategy in this set is semi-myopic, hence this guarantees a close to optimal outcome.

References

Athey, S. and I. Segal (2013). "An efficient dynamic mechanism." *Econometrica*, **81**(6), 2463–2485.

Babaioff, M., N. Immorlica and R. D. Kleinberg (2007). "Matroids, secretary problems and online algorithms." In *Proc. of the 18th Symposium on Discrete Algorithms (SODA)*, pp. 434–443.

Bartal, Y., F. Y. Chin, M. Chrobak, S. P. Fung, W. Jawor, R. Lavi, J. Sgall and T. Tichỳ (2004). "Online competitive algorithms for maximizing weighted throughput of unit jobs." In *Proc. of the 21st Symposium on Theoretical Aspects of Computer Science (STACS'04)*, pp. 187–198.

Basu, K. and J. W. Weibull (1991). "Strategy subsets closed under rational behavior." *Economics Letters*, **36**, 141–146.

Battigalli, P. and M. Siniscalchi (2003). "Rationalizable bidding in first price auctions." *Games and Economic Behavior*, **45**, 38–72.

Bergemann, D. and J. Välimäki (2010). "The dynamic pivot mechanism." *Econometrica*, **78**(2), 771–789.

Bernheim, B. D. (1984). "Rationalizable strategic behavior." *Econometrica*, **52**, 1007–1028.

Cavallo, R., D. Parkes and S. Singh (2007). "Efficient online mechanisms for persistent, periodically inaccessible self-interested agents." Working paper.

Cole, R., S. Dobzinski and L. Fleischer (2008). "Prompt mechanisms for online auctions." In *Proc. of the 1st International Symposium on Algorithmic Game Theory (SAGT'08)*, pp. 170–181.

Compte, O., R. Lavi and E. Segev (2012). "Efficiency of sequential English auctions with dynamic arrivals (abstract)." In *Proc. of the 13th ACM Conference on Electronic Commerce (EC'12)*, p. 369.

Dekel, E. and A. Wolinsky (2003). "Rationalizable outcomes of large private-value first-price discrete auctions." *Games and Economic Behavior*, **43(2)**, 175–188.

Demange, G., D. Gale and M. Sotomayor (1986). "Multi-item auctions." *Journal of Political Economy*, **94(4)**, 863–872.

Duggan, J. and M. L. Breton (1996). "Dutta's minimal covering set and Shapley's saddles." *Journal of Economic Theory*, **70**, 257–265.

Duggan, J. and M. L. Breton (1998). "Dominance-based solutions for strategic form games." Working paper.

Duggan, J. and M. L. Breton (2001). "Mixed refinements of Shapley's saddles and weak tournaments." *Social Choice and Welfare*, **8**, 65–78.

Gershkov, A. and B. Moldovanu (2009). "Learning about the future and dynamic efficiency." *The American Economic Review*, **99**(4), 1576–1587.

Gul, F. and E. Stacchetti (2000). "The English auction with differentiated commodities." *Journal of Economic Theory*, **92**, 66–95.

Hajek, B. (2001). "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*, pp. 434 – 439.

Hajiaghayi, M., R. Kleinberg, M. Mahdian and D. Parkes (2005). "Online auctions with re-usable goods." In *Proc. of the 6th ACM Conf. on Electronic Commerce (ACM-EC'05)*, pp. 165–174.

Horowitz, E. and S. Sahni (1978). *Fundamentals of Computer Algorithms.* Computer Science press.

Kesselman, A., Z. Lotker, Y. Mansour, B. Patt-Shamir, B. Schieber and M. Sviridenko (2004). "Buffer overflow management in qos switches." *SIAM Journal on Computing*, **33**(3), 563–583.

Kohlberg, E. and J. Mertens (1986). "On the strategic stability of equilibria." *Econometrica*, **54**, 1003–1038.

Lavi, R. and N. Nisan (2004). "Competitive analysis of incentive compatible on-line auctions." *Theoretical Computer Science*, **310**, 159–180.

Lavi, R. and N. Nisan (2005). "Online ascending auctions for gradually expiring items." In *Proc. of the 16th Symposium on Discrete Algorithms (SODA)*, pp. 1146–1155.

Lehmann, B., D. Lehmann and N. Nisan (2006). "Combinatorial auctions with decreasing marginal utilities." *Games and Economic Behavior*, **55**(2), 270–296.

McAfee, R. P. (2002). "Coarse matching." *Econometrica*, **70**, 2025–2034.

Myerson, R. (1981). "Optimal auction design." *Mathematics of Operations Research*, pp. 58–73.

Parkes, D. (2007). "Online mechanisms." In N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani, eds., *Algorithmic Game Theory*, pp. 411–439. Cambridge University Press.

Parkes, D. C. and S. Singh (2003). "An MDP-Based approach to Online Mechanism Design." In *Proc. 17th Annual Conf. on Neural Information Processing Systems (NIPS'03).*

Pearce, D. G. (1984). "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, **52**, 1029–1050.

Said, M. (2012). "Auctions with dynamic populations: Efficiency and revenue maximization." *Journal of Economic Theory*, **147**(6), 2419–2438.

Satterthwaite, M. A. and S. R. Williams (2002). "The optimality of a simple market mechanism." *Econometrica*, **70**, 1841–1863.

Shapley, L. (1964). "Some topics in two person games." *Annals of Mathematics Studies*, **52**, 1–29.

Voorneveld, M. (2004). "Preparation." *Games and Economic Behavior*, **42**, 403–414.

Welsh, D. J. A. (1976). *Matroid theory.* London : Academic Press.

A Useful Properties of Offline Allocations and Matroids

This section summarizes useful matroid properties that we use throughout our proofs. For completeness, we begin with a short introductory summary of matroids and their relevant properties.

Definition 17 (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 2 (E.g., Horowitz and Sahni (1978)) *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 exist 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 by Welsh (1976).

Claim 3 *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 $B_2 \setminus j' \cup j \in I$.*

4. *Given a weight function $w(\cdot)$ that assigns a* unique *weight to each element $s \in S$, there exists a unique base B with maximal weight $\sum_{s \in B} w(s)$.*

Claim 4 *Let $X, Y \in I$, and $X \not\subseteq Y$. 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 $j' \in X \setminus Y$ 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 5 *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 that there exist at least M players in A_t whose values are all zero, and with deadline M. This frees us from the need to worry about cases where some items remain unallocated, as we can always assign unallocated items to these zero-value players.

Definition 18 (A prefix) *A subset $X \subseteq S_t$ is called a prefix if it is a prefix of every allocation $S_t[t, M]$ of S_t. I.e., every allocation of the items $t, ..., M$ to the players in S_t must assign the first $|X|$ items to the players in X.*

Claim 6 *$X \subseteq S_t$ 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 $t' \leq t + |X| - 1$. Suppose player j' is allocated item $d(j)$ (recall that we can assume that every item is allocated to some player, as we have enough zero-value players). 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 $|S_t[t, t + |X| - 1]| = |X|$ it follows that $S_t[t, t + |X| - 1] = X$, i.e. it is a prefix. ∎

Definition 19 *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 the least index such that $x_{k+1} = x_k$, and fix $P_t(d) = S_t[t, x_k]$.*

Claim 7 *$P_t(d)$ is the prefix with minimal length among all prefixes with length $\geq d - t + 1$.*

Proof: First notice that, since every item is allocated to some player, it immediately follows that $|P_t(d)| = x_k - t + 1$. Also notice that, by construction, any $j \in P_t(d)$ has $d(j) \leq x_k = t + |P_t(d)| - 1$. Therefore, by Claim 6, $P_t(d)$ 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 6. ∎

Claim 8 *$j \in P_t(d)$ 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 and $j'_i = S_t[x_i]$. 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) and assign j_2's item to j'_1, j_2 will receive item x_2 and assign j_3's item to j'_2, ... , j_k will get item x_k and j'_k (that received x_k) will get j's item. Finally, allocate j_1's item to j. Therefore we have an allocation in which j receives some item $x \leq d$, as claimed. ∎

Claim 9 *$f_t = P_t(t) = P_t(|f_t| + t - 1)$.*

Proof: If $j \in f_t$ then there exists an allocation S_t such that $j = S_t[t]$. Since $P_t(t)$ is a prefix of $S_t[t, M]$ then $j \in P_t(t)$. On the other hand, Claim 8 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 7 we now get also that $P_t(t) = P_t(|f_t| + t - 1)$, as $P_t(t)$ is a prefix with length $|f_t|$. ∎

Claim 10 *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: Fix some $j^* \in P_t(t)$. We will actually show that $min_{j \in P_{t+1}(d)}\{v(j)\} \geq min_{j \in P_t(d) \setminus \{j^*\}}\{v(j)\}$. 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 $v(j) < min_{j \in P_t(d) \setminus \{j^*\}}\{v(j)\}$. Therefore $j \notin P_t(d) \setminus \{j^*\}$. Fix some allocation of S_{t+1}, and note that $P_t(d) \setminus \{j^*\}$ and $S_{t+1}[t+1, x]$ are two bases of the matroid over items $t+1, ..., x$. Since $j \in S_{t+1}[t+1, x] \setminus (P_t(d) \setminus \{j^*\})$, there exists $j' \in P_t(d) \setminus \{j^*\} \setminus S_{t+1}[t+1, x]$ 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 11 *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 $f_t \subseteq S_t$, the right to left direction is immediate. We prove 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 zero-value players, as explained in the opening paragraph to section A.1. By the optimality of S_t it follows that $S_t \subseteq \tilde{S}_t$. 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 a zero player, and $S_t \subseteq \tilde{S}_t \setminus j$. Since $\tilde{S}_t \setminus j$ is independent w.r.t. items $t+1, ..., M$, so is S_t, as needed. ∎

Claim 12 *Let $A'_t = A_t \cup j'$ (where $j' \notin A_t$). 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 it immediately follows that, if $f_t \subseteq S'_t$ then $f'_t = f_t$, and thus the first claim follows. The right to left direction of the second claim immediately follows from the fact that $j' \notin A_t$. 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 Proof of Lemma 4

Recall the statement of the Lemma: when all players play recommended strategies according to their true types, the allocation of any semi-myopic mechanism is a semi-myopic allocation rule.

To prove this, fix any time 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. We need to show that the mechanism chooses a player with value at least $v_t^* = \min_{j \in f_t^{true}} \{v(j)\}$. (recall that the definition of recommended strategies in the semi-myopic mechanism requires the declared values to be the true ones.)

Suppose first that there is some $j \in f_t^{true} \setminus S_t$. Fix some $i \in f_t$. By definition, we can assume that i is assigned to slot t by f_t. Since $j \notin S_t$, the optimality of S_t implies that $v(i) \geq v(j)$. Since $j \in f_t^{true}$, $v(j) \geq v_t^*$. Thus, $v(i) \geq v_t^*$, as claimed.

Otherwise, $f_t^{true} \subseteq S_t$. Note that S_t is also a valid schedule with respect to the true deadlines as the true deadlines are not smaller than the tentative deadlines. By the prefix construction process described in definition 19 in Section A.1 (applied to the schedule S_t but with the true deadlines), since tentative deadlines are not larger than true ones, we immediately get that $f_t \subseteq f_t^{true}$, and the claim follows. ∎

C Proof of Theorem 3

Recall the statement of the Theorem: Any semi-myopic allocation rule obtains at least one third of the optimal welfare, and this bound is tight.

We will prove that any allocation rule that produces an allocation ON has $v(OPT \setminus ON) \leq 2\sum_{t=1}^{M} v_t^*$, where OPT is the optimal allocation. From this, the theorem follows immediately, as any semi-myopic allocation rule has $v(ON[t]) \geq v_t^*$, and therefore $v(OPT) \leq v(OPT \setminus ON) + v(ON) \leq 2\sum_{t=1}^{M} v_t^* + v(ON) \leq 2 \cdot v(ON) + v(ON) = 3 \cdot v(ON)$. We first prove two useful claims:

Claim 13 *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 of Claim 13: 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. ∎

Claim 14 *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 of Claim 14: 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 $j \in X \setminus S'$. $S' \not\subset X$, since otherwise $S = S' \cup j \subseteq X$ 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 proceed to show that any allocation rule that produces an allocation ON has $v(OPT \setminus ON) \leq 2\sum_{t=1}^{M} v_t^*$. Fix some scenario, and let OPT and ON be the optimal and online allocations for this scenario. We describe $f : OPT \setminus ON \rightarrow \{1, ..., M\}$ such that f is 2 to 1 (i.e., for any t, the set $f^{-1}(t)$ contains **at most** two elements) and $v(j) \leq v_{f(j)}^*$ for any $j \in OPT \setminus ON$. 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 $f(j) = t_j^*$.

Claim 15 *For any $j \in OPT \setminus ON$, $v_{f(j)}^* \geq v(j)$.*

Proof of Claim 15: 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 Claim 13, $j \notin S_t \setminus m_t$: 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 Claim 14. 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. ∎

Claim 16 *f is 2 to 1. I.e., f maps at most two players to t.*

Proof of Claim 16: Fix any time t. Let $j_1 \in X_{t-1}$ be the player with minimal value such that $X_{t-1} \setminus j_1$ is independent w.r.t. 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 Claim 13: set $A = OPT[1, t-1] \setminus ON$, and $B = OPT[1, t] \setminus ON$. Since Y is optimal for A (by Claim 14) 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. ∎

This concludes the proof of Theorem 3. ∎

D Proofs deferred from Section 6

D.1 Proof of Theorem 5

In order to prove that the semi-myopic mechanism Set-Nash implements a 3-approximation of the welfare, we only need to prove that the recommended strategies are a Set-Nash equilibrium:

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 equal to t^*.

Claim 17 *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 of Claim 17: 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 $v(j) \geq v_{t_2^*}^*(\sigma_1) > v_{t_2^*}^*(\sigma_2)$ 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 10 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 18 *i's price in any strategy s_i is at least $v^*_{t^*}$ (where the other players play s_{-i}).*

Proof of Claim 18: 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 Claim 12, $f_t(\sigma') = f_t(\sigma'')$. 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 Claim 12, 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 payment in σ' is at least $v(j) \geq v^*_{t_0}(\sigma'')$. By Claim 17, $v^*_{t_0}(\sigma'') \geq v^*_{t^*}$, since σ'' differs from σ only by the choice of winners from the f_t's. Thus, i's payment is at least $v^*_{t^*}$, and the lemma follows. ∎

Claim 19 *The (recommended) strategy of arriving at time $r(i)$, declaring the true value and deadline and declaring a tentative deadline equal to t^* is a best response of i against s_{-i}.*

Proof of Claim 19: If $v(i) \leq v^*_{t^*}$ then i cannot possibly gain positive utility, as Claim 18 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 Claim 10 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 17 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. ∎

This concludes the proof of Lemma 5, and hence the proof of the Theorem. ∎

D.2 Bad Examples

We show, by an example, that the recommended strategies of the semi-myopic mechanism do not necessarily contain best responses to mixed-correlated strategies, i.e., it does not necessarily contain a best response against a distribution over all $R_{-i}(*)$. We 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 semi-myopic mechanism can choose the winners (first to last) x_1, y_1, y_2, z_3, z_1. *In this case, player* x_4 *loses. 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 *wins 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 recommended.*

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

- *At time* 1 *arrive player* $(x, 1)$ *and our player* $(x_4, 4)$.
- *At time* 2 *arrives player* $(x, 2)$.
- *At time* 3 *arrives 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* 4 *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 one 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.

D.3 Proof of Theorem 6

In order to prove that the Online Iterative Auction is an ignorable extension of a semi-myopic mechanism, we first prove that the iterative auction is an extension of a semi-myopic mechanism, and then show that this extension is ignorable.

Lemma 6 *If all players* i *play strategies in* $R_i(*)$ *then the iterative auction is a semi-myopic mechanism.*

Proof: We need to map every recommended strategy of the iterative auction to a strategy of the semi-myopic mechanism, such that the result of the iterative auction (winners plus payments) will match the criteria of a semi-myopic mechanism. This is done as follows. At time t, map every player 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 $Y = Y[t, ..., M]$ 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 19. 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 (in Eq. (2) in Section 6.1) $c_t(i)$ for a player i. Those are two differently defined terms, although we will see below that they are equal, for $d = Y[i]$.

Claim 20 *$p_t(d) \geq c_t(d)$.*

Proof of Claim 20: Fix any $j \in A_t \setminus Y$ with $d(j) \geq l(d)$. If $d(j) \geq d$ then since j is myopic, $p_t(d) \geq v(j)$. Thus, assume $d(j) < d$. By the construction of $P_Y(l(d)) = P_Y(d)$, there exist players $i_1, ..., i_k$ and items $t_1, ..., t_k$ such that $t_1 = l(d) \geq t$, $t_k = d$, for any index $1 \leq x \leq k$, $i_x = Y[t_x]$, and for any index $1 \leq x \leq k-1$, $d(i_x) \geq t_{x+1}$. 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 $t_1 = l(d) \leq d(j)$, $j \in A_t \setminus Y$ (i.e., j is active at time t but was not a tentative winner at the end of the time-t iterations), and j is myopic, it follows that $p_t(t_1) \geq v(j)$, and the claim follows. ∎

Claim 21 *If $p_t(d) > p_{t-1}(d)$ then $p_t(d) \leq c_t(d)$.*

Proof of Claim 21: 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 a bid 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. Note that by the auction rules, $j \notin X$.

Since $d(j) \geq d \geq l(d)$ and $v(j) > c_t(d)$, we conclude that $j \notin A_t \setminus Y$. Thus, $j \in Y$. The maximality of d implies that $c_t(d) \geq c_t(d+1) \geq p_t(d+1)$ which is in turn larger or equal to the price of item $d+1$ at the time j placed a bid on d. Since j is myopic but did not bid on item $d+1$, we conclude that $d(j) \leq d$. This implies that $j \notin Y[d+1, M]$, and also that $d(j) = d$. Since $d(j) = d$ the prefix properties imply that $j \notin Y[t, l(d)-1]$. Thus, $j \in Y[l(d), ..., d]$.

Since $j \in Y[l(d), ..., d] \setminus X[l(d), ..., d]$, there must exist some $i \in X[l(d), ..., d] \setminus Y[l(d), ..., d]$. In a similar way to the previous paragraph, if $i \in Y[d+1, M]$ then i placed a bid 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 a bid on an item with price higher than his value, again a contradiction. Since we have reached a contradiction, the claim follows. ∎

Claim 22 *$Y = S_t$, and, for any $d \geq t$ and $i = Y[d]$, $c_t(d) = c_t(i)$ (as defined in Eq. 2).*

Proof of Claim 22: 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 $d(j) \geq l(d)$. In the other direction, if $d(j) > d$ 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 Claim 20 to obtain an allocation, when replacing i with j.

From this and Claim 20 we have that for any $d \geq t$, $i = Y[d]$, 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 semi-myopic mechanism.

Claim 23 *In the Online Iterative Auction, the winner i of time t pays* $max_{r(i) \leq t' \leq t}\{c_{t'}(i)\}$.

Proof of Claim 23: Let $p_t(i)$ be i's tentative price at time t. Let t' be such that $i = Y[t']$. By the above 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 claims, it follows that $p_t(i) = p_t(t') = c_t(t') = c_t(i)$. ∎

This concludes the proof of Lemma 6. ∎

We now continue with the proof of the theorem. By Lemma 6, the Online Iterative Auction with the sets $R(*)$ constitutes a semi-myopic mechanism. Thus, the Online Iterative Auction is an extension of the semi-myopic mechanism. It remains to argue that it is an ignorable extension, i.e., if we fix any player i and and combination of recommended strategies of the other players, $s_{-i} \in R_{-i}(*)$, then i has a best response to s_{-i} in $R_i(*)$.

Suppose i plays some strategy $\bar{s}_i$, denote this scenario by σ', and let $Y_t(\sigma')$ (for any time t) be the tentative winners at time t in scenario σ'. Let t_0 be the first time t such that $i \in f_t(\sigma')$.

Since all players besides i are myopic with tentative deadline and then with final deadline, we can map them to types (v, d) as in Lemma 6. Let σ be this scenario, where i does not show up at all. If $t_0 > 1$, $i \notin f_1(\sigma')$ and therefore $f_1(\sigma) = f_1(\sigma')$. Let σ represent the scenario in which ties are broken such that the winner of time 1 is the winner of time 1 in σ'. This holds (inductively) for all $t < t_0$. Thus, $A_{t_0}(\sigma') = A_{t_0}(\sigma) \cup i$, and t_0 is also the first time t such that $f_t(\sigma) \not\subseteq Y_t(\sigma')$.

Define t^* as in Eq. 3 of the proof of Lemma 5, and fix $j \in f_t(\sigma) \setminus Y_t(\sigma')$. Since $i \in f_{t_0}(\sigma')$, we use a chain argument similar to that of Claim 20 to conclude that i's payment is at least $v(j) \geq v^*_{t_0}(\sigma) \geq v^*_{t^*}(\sigma)$. 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 theorem follows. ∎

D.4 Proof of Theorem 7

We will prove that the Sequential Japanese Auction is an ignorable extension of a semi-myopic mechanism by the following claims:

Lemma 7 *If all players i play strategies in $R_i(*)$ then the Sequential Japanese Auction forms a semi-myopic 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.

Claim 24 *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 of Claim 24: 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. ∎

Claim 25 $p^* = max\{ \ v(j) \mid j \in A_t \setminus S_t \ \}$.

Proof of Claim 25: Let j^* be the player with maximal value among those in $A_t \setminus S_t$. By the previous 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. ∎

Claim 26 *For any* $i \in f_t$, $p^* = c_t(i)$.

Proof of Claim 26: 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 claim, and from the definition of $c_t(i)$. ∎

Claim 27 *For any* l', $|D(p^*, l') \cup \cdots \cup D(p^*, 1)| = l'$.

Proof of Claim 27: 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 $l' \geq |D(p^*, l') \cup \cdots \cup D(p^*, 1)|$. ∎

Claim 28 *If* $|D(p^*, l' + 1)| = 1$ *then* $D(p^*, l') \cup \cdots \cup D(p^*, 1)$ *is a prefix.*

Proof of Claim 28: 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 6 that $D(p^*, l') \cup \cdots \cup D(p^*, 1)$ is a prefix. ∎

Claim 29 *Let* x^* *be the critical number of drop-outs, and* $D^* = \cup_{x \leq x^*} D(p^*, x)$, *as in def. 16. Then* $D^* = f_t$.

Proof of Claim 29: 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 7, 9). Therefore $x^* = l$ and $D^* = f_t$. ∎

From all the above, the proof of the Lemma 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 semi-myopic 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^*. ∎