Expressiveness and Optimization under Incentive Compatibility Constraints in Dynamic Auctions

A dissertation presented
by

Gabriel Florin Constantin

to
The School of Engineering and Applied Sciences
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of

Computer Science

Harvard University
Cambridge, Massachusetts
May 2009

Thesis advisor
David C. Parkes

Author
Gabriel Florin Constantin

Expressiveness and Optimization under Incentive Compatibility Constraints in Dynamic Auctions

Abstract

This thesis designs and analyzes auctions for persistent goods in three domains with arriving and departing bidders, quantifying tradeoffs between design objectives. The central objective is incentive compatibility, ensuring that it is in bidders' best interest to reveal their private information truthfully. Other primary concerns are expressiveness, i.e. the richness of the effective bidding language, and optimization, in the form of aiming towards high revenue or high value of the allocation of goods to bidders.

In the first domain, an arriving bidder requests a fixed number of goods by his departure, introducing combinatorial constraints. I achieve the global property of incentive compatibility via self-correction, a local verification procedure, applied to a heuristic modification of an online stochastic algorithm. This heuristic is flexible and has encouraging empirical performance in terms of allocation value, revenue and computation overhead.

In the second domain, impatient buyers make instantaneous reservation offers for future goods. Introducing the practical ability of cancellations by the seller leads to an auction with worst-case guarantees without any assumption on the sequence of offers. A buyer whose reservation is canceled incurs a utility loss proportional to his value, but receives an equivalent cancellation fee from the seller. A simple payment scheme ensures a novel incentive compatibility concept: no bidder can profit from a lower bid while no truthful winner can profit from any different bid. I establish that no fully incentive-compatible auction can achieve similar worst-case guarantees.

In the third domain, I consider the first dynamic generalization of the classical economic model of interdependent values for a single good. In this model, a bidder's value for the good depends explicitly on other bidders' private information. I characterize incentive-compatible dynamic interdependent-value auctions and I establish that they can be reasonable if and only if no bidder can manipulate his departure. I suggest and analyze a mixed-integer programming formulation and a heuristic for designing such an auction to maximize revenue when bidders have fixed arrivals and departures.

Contents

Citations to Previously Published Work

Large portions of this thesis have previously appeared as follows:

Chapter 2, Section 2.2 "Preference-Based Characterizations of Truthfulness and the Limited Expressiveness of Order-Based Domains", by Florin Constantin and David C. Parkes. In the Proceedings of the Workshop on Preference Handling, 2005.

Chapter 3 "Self-Correcting Sampling-Based Dynamic Multi-Unit Auctions", by Florin Constantin and David C. Parkes. In the Proceedings of the ACM Conference on Electronic Commerce (EC), 2009.

Chapter 4 "An Online Mechanism for Ad Slot Reservations with Cancellations", by Florin Constantin, Jon Feldman, S. Muthukrishnan and Martin Pál. In the Proceedings of the Symposium on Discrete Algorithms (SODA), 2009.

Chapter 5 "Online Auctions for Bidders with Interdependent Values", by Florin Constantin, Takayuki Ito and David C. Parkes. In the Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2007.

"On Revenue-Optimal Dynamic Auctions for Bidders with Interdependent Values", by Florin Constantin and David C. Parkes. In the Proceedings of the Workshop on Agent Mediated Electronic Commerce (AMEC), 2007.

Acknowledgments

I have been extraordinarily fortunate to be advised by David Parkes. He will always be my academic role model – I do not expect to ever match his enthusiasm, tact, technical strength and pedagogical talent.

I hope that in the future I will live closer to my parents. Until then, I would like to first thank my mother for giving me everything that she could. The wisdom in her advice is gradually revealed to me as I become more mature. Second, I owe my career and direction in life to my father and I will be forever grateful.

For as long as I can remember, my cousin Virgil has provided me with guidance, partnership, competition and, more recently, computer parts.

Alex and Lauren have made my life more beautiful in recent years with their kindness and love. I hope I have not failed them in too many ways.

Out of all the other friends that have been close to me during my doctoral studies, I would like to acknowledge Alin, Andreea, Dan, Mihai, Răzvan and Simona.

I am grateful to Jon Feldman, Muthu Muthukrishnan, Martin Pál and Daniela Rus for their outstanding mentorship at Google and Dartmouth College.

My experience as a graduate student at Harvard would have been much poorer without my colleagues Adam Juda, Ben Lubin, Ece Kamar, Giro Cavallo, Guille Diez-Canas, Haoqi Zhang, Itai Ashlagi, Ivo Parashkevov, Jacomo Corbo, Kobi Gal, Malvika Rao, Philip Hendrix, Sébastien Lahaie, Shaili Jain and Sven Seuken.

Finally, I cannot be grateful enough to my beloved grandmother Micuţu for her love and care. Although she passed away while I was writing this thesis, I am sure that she will be watching me at yet another graduation.

In memory of Micuţu

Chapter 1

Introduction

How should goods be sold over time? Market mechanisms with temporal components are becoming widespread, but there is relatively little research on them. A canonical such mechanism is a dynamic auction, in which the population of buyers changes with time, and buyers usually have temporal constraints for allocation and payment decisions. For example, in sponsored search auctions for banner advertisements (ads) on webpages, bidders must typically know the characteristics of the ad slots ahead of time in order to develop appropriate ads. Internet auctions on sites such as `eBay.com` are quite prone to temporal *manipulations*. The best-known example is sniping, which is the practice of outbidding the current winner in the last seconds of an auction so that he does not have time to respond.

Dynamic auctions are of interest within computer science because they allow connections between online algorithms, stochastic optimization and auction theory. In this thesis, I explore this space and consider both approximation algorithms that work in a prior-free environment and sample-based online optimization that requires a generative model of the environment. In both cases, I explore the constraints that incentive compatibility (resilience to manipulations) imposes on the preferences that bidders can express and on the optimization properties of auctions.

Strategic interactions of self-interested parties (such as the seller and buyers), taking their environment as given, are analyzed in *game theory*. An auction is a canonical example of a *mechanism*, an environment for self-interested parties. *Mechanism design*, sometimes called "inverse game theory", aims to design such environments with certain objectives in mind. Mechanism design is thus a natural area for computer science approaches, since they are traditionally oriented towards building and evaluating systems. *Computational mechanism design* additionally sets typical computer science objectives, such as good performance with a minimal set of assumptions, approximation when optimality is infeasible or tackling problem complexity. Unlike in traditional computer science, algorithms for self-interested parties also require a game-theoretic perspective since design choices may influence parties' behavior.

The cross-pollination of computer science and economics has proved very fruitful, with a growing literature for computationally difficult problems of economic interest. In prediction markets, traders buy and sell "securities" on outcomes of future events, such as whether the number of cases of a certain disease will exceed 1000 by a specified deadline[1] or whether a cure for cancer will be found by the end of the year 2010.[2] Allowing expressive languages for securities such as compound securities (e.g. more than 1000 cases *and* the discovery of a cure for cancer by the end of 2010) introduces significant computational challenges for the market maker, both in determining whether to accept an order while guaranteeing no loss [75] and pricing [20] by the widely used logarithmic market scoring rule [46]. An important part of the computational mechanism design literature analyzes combinatorial auctions [26], in which a buyer can bid on combinations of items. On the

[1] `www.intrade.com`, accessed May 2009.

[2] `www.ideosphere.com/fx-bin/Claim?claim=Canc`, accessed May 2009.

practical side, Internet auctions yield billions of dollars in revenue yearly and their setup requires expertise in core computer science areas such as natural language processing or machine learning.

There is significantly less research on dynamic aspects of auctions than on static ones. This is partly due to the complex interactions that dynamic auctions entail, often leading to complex game-theoretic analysis, computational intractability or lack of closed-form solutions except in the simplest of settings.

Challenges in dynamic auctions

I now review typical challenges in a dynamic auction.

Combinatorial optimization. From an operations research perspective, auctions are knapsack problems that may have stochastic information. In a static knapsack problem with full information, one has a container of limited capacity (a knapsack) and a set of resources to choose from, each with its own volume and profit. One's goal is to select a set of resources that fit in the container and that yield a high profit. Due to their widespread applicability, there is an extensive literature on knapsack problems; see for example the survey of Kellerer et al. [51]. In combinatorial auctions, where buyers have complex valuations for sets of goods, both the seller and the buyers may face knapsack problems. In stochastic knapsack problems, some relevant quantity is not known with certainty, for example the existence of a certain resource or its reward or volume. Many dynamic auctions are naturally stochastic, with only partial information about the future. In this thesis, the most challenging knapsack problems will be encountered in Chapter 3.

Dynamic programming. A common approach to sequential optimization with a finite horizon and a model of the future is dynamic programming [11], in which one reasons backwards from the horizon to the current time period. If one determined the optimal decision at any state of the world in the next period, one could also compute the optimal decision in the current time period. "The curse of dimensionality" of dynamic programming is one of the main reasons for its intractability; Powell [76] identifies three objects leading to high-dimensionality: states, outcomes and actions. For a dynamic auction that sells identical copies of a good and faces uncertain future demand, these correspond respectively to the number of items (supply) available, the realization of demand at each future time step (bids and quantities) and the number of items to allocate. A dynamic programming approach would require computing, at each time step t, a matrix specifying the best number of items to allocate at t for each supply available and each realization of future bidders. In this thesis, the curse of dimensionality will be most apparent in Chapter 5. In Chapter 3, I escape it via an online stochastic combinatorial optimization [84] approach.

Incentives. Apart from pure optimization concerns, in any auction one must worry about manipulations by the bidders. The existing results for static auctions (I review the most important ones in Chapter 2) usually fall in two categories: impossibility results for general domains of bidder preferences [54, 77] or possibility results for small domains [2, 60]. In dynamic auctions, resilience to manipulations also poses restrictions on the use of temporal information by the seller. For example, a bidder may understate his patience if he fears that stating a high patience will expose him to more competition. Furthermore, the optimum solution may be prone to manipulations, in which case tradeoffs of optimality and resilience to manipulations must be made. Such tradeoffs will be encountered throughout the thesis.

Expressiveness. Bidders have inherently different preferences and are usually interested in convenient ways of expressing them. On the other hand, too complicated methods for expressing preferences may discourage auction participants. Having a large space of possibilities may also render difficult the prediction of bidders' behavior. While expressiveness is a concern for static auctions as well, in dynamic auctions bidders are additionally interested in having their temporal constraints met by the seller.

This thesis explores the constraints imposed by incentive compatibility on the expressiveness and optimization properties of dynamic auctions. I design and analyze dynamic auctions for three environments, highlighting along the way tradeoffs between objectives.

1.1 Roadmap and contributions

The technical preliminaries are set forth in Chapter 2, reviewing mechanism design notions and prominent results in both static and dynamic settings. In Chapter 3, I design and evaluate empirically a method to achieve resilience to manipulations via a self-correction wrapper around a state-of-the-art approximation algorithm. In Chapter 4, I initiate the study, from a worst-case perspective, of seller cancellations that negatively affect impatient buyers. Chapter 5 investigates dynamic versions of a classical static economics model in which a bidder's value depends explicitly on other bidders' private information. This thesis is concluded by the discussion in Chapter 6.

I will now present the typical setting and my contributions in more detail.

The environments of interest, to be described shortly, aim to satisfy certain objectives for a seller offering a set of durable items to a population of bidders over a finite number of time periods. Bidders may arrive and depart from the auction and may be interested in more than one item. A bidder's private information may include a form of evaluation (e.g. willingness to pay) for each set of items and his times of arrival to and departure from the auction. Typically, a bidder's value for an allocation will be the same in his arrival-departure interval, i.e. time does not affect a bidder's valuation (e.g. by discounting) other than constraining his value to be 0 outside this interval. Each bidder makes a single claim about his private information, at a time that becomes his arrival.

A common desideratum is the auctions' resilience to bidders' profitable manipulations of their values and, where applicable, of their temporal information. Typical objectives are the overall value of the allocation and seller's revenue.

Incentive compatibility via self-correction, Chapter 3

In many auctions, buyers have volume and temporal constraints, being interested in a *fixed* number of goods by a certain deadline. The work presented in Chapter 3 is the first that scalably achieves full resilience to manipulations in dynamic auctions for persistent goods sold to partially patient bidders, each demanding a fixed number of goods. Interactions in such auctions are complex, due to bidders' partially overlapping windows of activity and the combinatorial structure of feasible allocations. Optimizing dynamic auctions with complex interactions is challenging for two reasons: the large number of possible states of the market (the "curse of dimensionality") and the non-standard constraints *across* these imposed by resilience to manipulations. I circumvent these challenges via the computational self-correcting approach of output *ironing* [71], applied to the Consensus [84] online stochastic optimization algorithm. Self-correction discards decisions that are profitably manipulable by buyers, requiring a powerful underlying optimization algorithm that seldom allows fruitful manipulations.

Most successful algorithms (including Consensus) that are not concerned with incentives need to wait almost until a buyer is leaving for a maximally informed decision regarding his allocation, which leads to frequent manipulations in my model. I design a heuristic modification, NowWait, of Consensus that aggregates quantitative and temporal bid information, and establish its successful interfacing with the ironing procedure. Compared to a naive simplification of Consensus, NowWait can achieve higher social welfare, especially in low-demand markets, and has a limited overhead in both computation costs and ironing discards. I successfully incorporate in NowWait the classical [64] method of ranking bidders by virtual valuations towards the alternate goal of high revenue.

All performance measures in Chapter 3 are in expectation over a prior distribution on future buyers. I adopt a different perspective in Chapter 4, where no distributional assumption is made on future demand.

Reservations with costly seller cancellations, Chapter 4

Reservations are a widely used retail practice. In numerous markets, each buyer makes a single instantaneous offer (for example via phone or a web form) to the seller for obtaining some of the goods in the future.

Chapter 4 presents a simple model and a dynamic auction M for advance-booking of items by impatient buyers, each aiming to maximize his utility. My model's novel, practical, feature is that the seller can *cancel* at any time any earlier reservation, in which case its holder incurs a proportional utility loss, but receives an equivalent cancellation fee. Chapter 4 focuses on the case where bidders have unit-demand, briefly discussing the challenges of extending the framework beyond unit-demand.

Constant factor approximations are achieved in the worst-case when costly cancellations are allowed, but would otherwise require assumptions on bidders' values or arrival order. My auction M approximates within a constant factor both the *a posteriori* revenue of the canonical Vickrey-Clarke-Groves auction and the *a posteriori* social welfare. M can match an upper bound on the competitive ratio of *any* deterministic online algorithm if performance is measured as the sum of winning bidders' values minus the utility losses of bidders whose reservations were canceled. M's technical core is an approximation algorithm for a semi-online weighted bipartite matching problem with costly preemptions.

My auction M induces a novel game-theoretic property that is practical as a first recommendation for unsophisticated bidders: an honest winner cannot improve his utility by any other bid and any buyer always prefers bidding his true value to any lower bid. In contrast, I establish that full resilience to manipulations no longer allows constant competitiveness with respect to social welfare.

Dynamic interdependent-value auctions, Chapter 5

In a dynamic market, being able to update one's value based on information available to *other* parties currently in the market can be critical to having profitable transactions. The classical economics model of *interdependent values* (IDV) in static auctions allows such a dependence: a bidder's value is obtained via a publicly known formula from the private information (*signals*) of all bidders.

In Chapter 5, I consider the first dynamic generalization of (single-item) auctions for IDV bidders with one-dimensional signals, arrivals and deadlines. I characterize the resilience to manipulations of such auctions as equivalent to the existence, for each bidder, of a critical signal such that the bidder wins the item if and only if the bidder reports a signal at least as high as the critical one. A bidder's critical signal cannot decrease if the bidder reports a later arrival. Using this characterization, I show that reasonable auctions can be resilient to manipulations if and only if bidders' deadlines are public.

I then adopt a computational approach for the design of single-item revenue-optimal dynamic auctions for bidders with known arrivals and deadlines but private one-dimensional signals. I present a mixed-integer programming formulation of revenue-optimality that leverages the characterization of resilience to manipulations. I highlight general properties of revenue-optimal dynamic auctions in a simple parametrized example. The most promising direction is suggested by a striking resemblance between the policy obtained by numerically solving the formulation and the policy obtained by generalizing the single-item revenue-optimal auction [64] for private values to IDV bidders.

Chapter 2

Static and Dynamic Mechanisms

Abstract. I introduce mechanism design concepts, in both static and dynamic contexts, anticipating points of interest in later chapters.

In a static context, I highlight the tension between the generality of the domain of preferences and the richness of the space of incentive compatible social choice functions that are implementable on that domain. I review existing characterizations of incentive compatibility and I demonstrate through examples their limited applicability.

I define a dynamic model and a dynamic auction and review dynamic allocation mechanisms via examples and results regarding their incentive compatibility.

2.1 Static mechanism design

2.1.1 Mechanisms as games

Recall the introductory setting, in which a seller faced a decision on how to allocate goods that she owns given the values that buyers[1] associate to subsets of goods.

---

[1] I will use the terms "agent", "bidder", "buyer" or "player" interchangeably. Their gender will be assumed to be masculine, while the seller's gender will be assumed to be feminine.

The way to aggregate individual preferences and take a decision (usually with some objectives in mind), is formalized by an outcome rule (known as a social choice function when players state their preferences directly). A mechanism simply encodes the entire framework for the decision that affects players and that is taken based on players' actions.

Definition 1. *A* static *(or simultaneous or one-shot) deterministic* mechanism M *with* n *players consists of*

- Θ*, the set of each player* i*'s private information (or* type*)* $\theta_i \in \Theta$*.*
- *a set of actions* W *(same set for each player* i*)*
- *a set of outcomes* A
- *an outcome rule* $\phi : W^n \to A$*, assigning an outcome to an* n*-tuple of actions.*
- *a payoff mapping* $u : \Theta^n \times A \to \mathbb{R}^n$ *such that* $u_i(\theta, a)$ *represents player* i*'s* payoff *(or* utility*) when the outcome chosen is* a *and the private information of player* i' *is* $\theta_{i'}, \forall i' = 1..n$*.*

A mechanism is a *game*, modeling strategic interactions (through actions) among players leading to an outcome. I consider mechanisms where the same actions are available, regardless of a player's information or of his identity. The payoff of a certain outcome may however be different among players. The set of possible actions W, the set of possible outcomes A and the mapping ϕ from actions to outcomes in M are *common knowledge*: any player knows them, any player knows that any other player knows them and so on. Unless otherwise stated, I will only consider mechanisms with *incomplete information*: any other player i' has some uncertainty about any player i's *payoff* function u_i.

To formalize the fact that a player has no private information I will use the notation $\Theta = \{\perp\}$ instead of assuming that $\Theta = \emptyset$. Note that incomplete information is possible even if $\Theta = \{\perp\}$: then for any outcome a, $u_i(a)$ is a constant and for any other player j, the value of at least one of these constants is not known to j. On the other hand, complete information is possible when $\Theta \neq \{\perp\}$ if i knows all other players' payoff functions.

Note that, unless otherwise specified, I inherently assume that a player cannot rule out any other player having a certain type.[2] A player may however have *beliefs* on other players' private information given his own private information: for example the higher his own value for an item, the higher he expects others' values for the item to be (this type of correlation is known as affiliation). We adopt a belief model in Chapters 3 and 5 for each bidder i, there is a public probability distribution on i's type.

Apart from Chapter 5, a player's value for an outcome is invariant to other[3] players' information: $u_i((\theta_i, \theta_{-i}), a) = u_i((\theta_i, \theta'_{-i}), a) \,\forall\, \theta_{-i}, \theta'_{-i} \in \Theta^{n-1}$. I will adopt shorthand $u_i(\theta_i, a)$ for i's utility for an outcome a since this utility will not depend on others' private information, but only on i's private information θ_i.

Strategies encode the dependence of actions on private information.

Definition 2. *A* pure strategy s_i *for player i is a deterministic mapping from his private information to the set of actions:* $s_i : \Theta \to W$.

In the following, "strategy" may be used as shorthand for "pure strategy" when clear from context.

Example 1 serves as the running illustration for the concepts introduced. Players' payoffs are formally specified in Tables 2.1 and 2.2, in which an entry (x_1, x_2) specifies that

---

[2] A more general model is that of information sets (or partitions). Suppose that players may have some knowledge about each other's types (e.g. i knows that $\theta_{i+1} \geq \theta_{i+2}$) but any player's type offers no further information on others' types. Then each player i has an information partition $\cup_{C \in \mathbf{C}_i} \Theta \times C$, where $\cup_{C \in \mathbf{C}_i} C = \Theta^{n-1}$. In our model, there is a single partition class: $\mathbf{C}_i = \{\Theta^{n-1}\}, \forall i$. In contrast, the more general definition only requires that the partition $\mathbf{C}_i$ on others' information is the same regardless of i's private information. A strategy (see Def. 2) maps a class of the partition to an action.

[3] ζ_{-i} denotes $(\zeta_1, \ldots, \zeta_{i-1}, \zeta_{i+1}, \ldots, \zeta_n)$ for any vector (or n-tuple) $\zeta = (\zeta_1, \ldots, \zeta_n)$ and any $1 \leq i \leq n$.

(a) No player i observes his own type θ_i (value for the item) or his opponent's type.

	20	40	50	70
20	$\frac{\theta_1-20}{2}, \frac{\theta_2-20}{2}$	$0, \theta_2 - 40$	$0, \theta_2 - 50$	$0, \theta_2 - 70$
40	$\theta_1 - 40, 0$	$\frac{\theta_1-40}{2}, \frac{\theta_2-40}{2}$	$0, \theta_2 - 50$	$0, \theta_2 - 70$
50	$\theta_1 - 50, 0$	$\theta_1 - 50, 0$	$\frac{\theta_1-50}{2}, \frac{\theta_2-50}{2}$	$0, \theta_2 - 70$
70	$\theta_1 - 70, 0$	$\theta_1 - 70, 0$	$\theta_1 - 70, 0$	$\frac{\theta_1-70}{2}, \frac{\theta_2-70}{2}$

(b) Each player i observes his own type θ_i (value for the item) and his opponent's type to be 50.

	20	40	50	70
20	*15, 15*	0, 10	0, 0	0, −20
40	10, 0	*5, 5*	0, 0	0, −20
50	0, 0	0, 0	*0, 0*	0, −20
70	−20, 0	−20, 0	−20, 0	−10, −10

Table 2.1: **First-price auction** for an item: the highest bid wins (a tie is broken uniformly at random), paying his bid. Players' payoffs are shown as a pair assuming that each player can bid 20, 40, 50 or 70.

player i's payoff (value minus price) is x_i, for $i = 1, 2$.

Example 1. *Consider two buyers interested in one item; bidder i has private type (value for the item) of θ_i. Possible*[4] *bids are* $20, 40, 50$ *and* 70. *The higher bid wins; in case of a tie, the item is given with equal (50%) chance to one of the bidders.*

Two possible payment rules are considered: winner pays his bid (first-price, Table 2.1) or winner pays his opponent's bid (second-price, Table 2.2). In addition to all zero entries in the payoff matrices, the players have the following information

- *Tables 2.1(a) and 2.2(a) present a general setting: neither player knows θ_1 or θ_2.*
- *Tables 2.1(a) and 2.2(a) can also represent a private value setting: bidder i knows his own θ_i (instantiated in each table entry), but does not know the other $\theta_j (j \neq i)$.*
- *Table 2.1(b) and Table 2.2(b) present payoffs in a setting with no private information: each player knows θ_1 and θ_2.*

Given a mechanism, how can one expect it to be approached by players or, more precisely, when is a certain strategy good in a mechanism? The following definition answers

[4]Even though bids are restricted for clarity, one can imagine a situation where each bidder has two "chips" of value 20 and one of value 50. The strategy of bidding 90 is omitted for brevity.

(a) No player i observes his own type θ_i (value for the item) or his opponent's type.

	20	40	50	70
20	$\frac{\theta_1-20}{2}, \frac{\theta_2-20}{2}$	$0, \theta_2 - 20$	$0, \theta_2 - 20$	$0, \theta_2 - 20$
40	$\theta_1 - 20, 0$	$\frac{\theta_1-40}{2}, \frac{\theta_2-40}{2}$	$0, \theta_2 - 40$	$0, \theta_2 - 40$
50	$\theta_1 - 20, 0$	$\theta_1 - 40, 0$	$\frac{\theta_1-50}{2}, \frac{\theta_2-50}{2}$	$0, \theta_2 - 50$
70	$\theta_1 - 20, 0$	$\theta_1 - 40, 0$	$\theta_1 - 50, 0$	$\frac{\theta_1-70}{2}, \frac{\theta_2-70}{2}$

(b) Each player i observes his own type θ_i (value for the item) and his opponent's type to be 50.

	20	40	*50*	70
20	$15, 15$	$0, 30$	$0, 30$	$0, 30$
40	$30, 0$	$5, 5$	$0, 10$	$0, 10$
50	$30, 0$	$10, 0$	$0, 0$	$0, 0$
70	$30, 0$	$10, 0$	$0, 0$	$-10, -10$

Table 2.2: **Second-price auction** for an item: the highest bid wins (a tie is broken uniformly at random), paying the second-highest bid. Players' payoffs are shown as a pair assuming that each player can bid 20, 40, 50 or 70.

this question in two special cases. First, it identifies as unreasonable[5] a dominated strategy: a strategy that is worse than another one, regardless of other players' strategies. Second, it identifies as reasonable a dominant strategy: a strategy that is at least as good as any other one, regardless of other players' strategies. Note that for each condition in Def. 3, it is equivalent to consider all other players' actions or strategies.

Definition 3. *A strategy s_i is* dominated *if there exists a strategy s'_i that is no worse for i (i.e. dominates s_i), for any private information and any actions of other players:*

$$u_i(\theta, \phi(s_i(\theta_i), w_{-i})) \le u_i(\theta, \phi(s'_i(\theta_i), w_{-i})),\ \forall\, \theta \in \Theta^n\, \forall\, w_{-i} \in W^{n-1}$$

A strategy s_i is a best-response *to actions w_{-i} of other players if it is no worse for i than any other strategy s'_i, for any private information:*

$$u_i(\theta, \phi(s_i(\theta_i), w_{-i})) \ge u_i(\theta, \phi(s'_i(\theta_i), w_{-i})),\ \forall\, s'_i : \Theta \to W\, \forall\, \theta \in \Theta^n$$

A strategy s_i is dominant *if it is no worse for i than any other strategy s'_i, for any private information and any actions of other players:*

$$u_i(\theta, \phi(s_i(\theta_i), w_{-i})) \ge u_i(\theta, \phi(s'_i(\theta_i), w_{-i})),\ \forall\, s'_i : \Theta \to W\, \forall\, \theta \in \Theta^n\, \forall\, w_{-i} \in W^{n-1}$$

[5]Technically, a dominated strategy s_i should never be played if some s'_i is *always* better than (i.e. strictly dominates) s_i.

That is, a strategy s_i is dominant if it is a best response to any actions of other players. If, for each player i, strategy s_i is dominant then the vector $(s_1, \ldots, s_n)$ is termed a *dominant strategy equilibrium.*

In Table 2.2(b), bidding 50, i.e. one's true type, is a dominant strategy. Notice, however, that if both players play the dominant strategy then each get a payoff of 0. Surprisingly by coordinating on playing 20, each would have a higher payoff (of 15). A similar situation is captured in Prisoner's Dilemma, a classical strategic interaction in game theory. This is a conflict between incentives and optimization, a common thread throughout this thesis.

Note that if the discrete bid levels do not contain a bidder's private value (for instance, if in Example 1 a bidder's private value is 30) then the second-price auction may no longer have a dominant strategy.

It is unlikely that large games with "generic" payoffs have either a dominated or a dominant strategy. Thus a more general concept for predicting behavior is needed. Nash equilibrium is a strategic solution with wider applicability, as will be argued soon.

Definition 4. *An n-tuple of strategies $(s_1, \ldots, s_n)$ is an ex post* Nash equilibrium *if s_i is a best response to* strategies s_{-i}*, for any vector of private information:*

$$u_i(\theta, \phi(s_i(\theta_i), s_{-i}(\theta_{-i}))) \geq u_i(\theta, \phi(s'_i(\theta_i), s_{-i}(\theta_{-i}))),\ \forall\, \theta \in \Theta^n,\ \forall\, s'_i : \Theta \to W \tag{2.1}$$

In the absence of private information ($\Theta = \{\perp\}$), the concept of Def. 4 becomes the Nash equilibrium, due to Nash [67]. Each strategy s_i becomes equivalent to an action w_i and Eq. (2.1) becomes $u_i(\phi(w_i, w_{-i})) \geq u_i(\phi(w'_i, w_{-i})),\ \forall\, w'_i \in W$.

Suppose that bidders have private values: $u_i(\theta, \cdot) = u_i(\theta_i, \cdot)$. If $(s_1, \ldots, s_n)$ is an ex post Nash equilibrium and $s_{-i} : \Theta^{n-1} \to W^{n-1}$ is an onto function, spanning all possible actions, then s_i must be a dominant strategy. Thus an ex post Nash equilibrium is a dominant strategy equilibrium if, in particular, $W = \Theta$ and $s_j(\theta_j) = \theta_j$ (direct revelation mechanism, see Definition 7 below).

Note that the notions of dominated and dominant strategy as well as that of equilibrium use weak inequalities. There are variants of all three concepts with strict inequalities: strictly (as opposed to weakly (Def. 3)) dominated and dominant strategies and strict (as opposed to weak (Def. 4)) equilibrium.

Let us revisit the first-price auction (Table 2.1(b)) in Example 1. Note that bidding 70 is strictly worse than bidding anything else (in particular, bidding 70 is a strictly dominated strategy). It cannot therefore be part of a best-response. Bidding 50 is weakly dominated by bidding 20 or 40, but both players bidding 50 is a (weak) Nash equilibrium. By eliminating the dominated strategies of 50 and 70 in Table 2.1(b), one is left with two pure strict Nash equilibria: $(20, 20)$ and $(40, 40)$.

In contrast, in the second-price auction (Table 2.2(b)), each player has a dominant strategy: truthful bidding (which induces, of course, an equilibrium as well).

Thus far I have only considered deterministic strategies.

Note how in Example 1 in Table 2.1(b) at any pure strategy Nash equilibrium the two bidders need to *coordinate* on bidding the same amount: 20, 40 or 50.

Randomization can help players compromise between several strategies. Denote by $\Delta(W)$ the set of distributions over actions, i.e. vectors of $|W|$ non-negative numbers summing to 1.

Definition 5. *A* mixed strategy *s_i for player i is a mapping from his private information to the set of distributions over actions: $s_i : \Theta \to \Delta(W)$.*

For any θ_i, the support *of the mixed strategy s_i is the set of strategies with non-zero probabilities in the distribution $s_i(\theta_i)$.*

A mixed Nash equilibrium *is then a Nash equilibrium with mixed strategies.*

Contrast this with Definition 2. Clearly, a pure strategy s can be interpreted as a mixed strategy with only s in its support (thus having a probability of 1).

The power and generality of Nash equilibria is revealed by the following result

Theorem 1. *[67] Any finite matrix game (in my presentation[6], a mechanism with complete information $\Theta = \{\perp\}$ and $u_i(\perp, \phi) = \phi$) has at least one Nash equilibrium (possibly in mixed strategies).*

Consider two strategies s_i and s'_i of player i in the support of a mixed strategy s_i in a mixed Nash equilibrium where others play strategies s_{-i}. If s_i yields a strictly higher payoff to i than s'_i, then i could strictly improve upon s_i by playing s_i whenever the randomization in s_i indicates playing s'_i, thus contradicting s_i's best-response property to s_{-i}. Thus the expected payoffs of s_i and s'_i must be equal conditioned on others playing s_{-i}. This property is clearly helpful in finding the weights of strategies in a mixed Nash equilibrium's support.[7] In Example 1, Table 2.1(b), simple algebra shows that, bidding 20 or 40 with equal (50%) probability is the unique proper mixed Nash equilibrium. [8]

2.1.2 Mechanism design and desiderata

Design goals for mechanisms are often in terms of players' true types; the challenge is then in choosing the mechanism's social choice function towards such goals.

Definition 6. *A mechanism M with outcome rule ψ_M* implements *an outcome rule ϕ*

---

[6]Game theory is, naturally, the predecessor of mechanism design. Presenting a (simultaneous) game as a mechanism preserves the focus on mechanisms and attests to the generality of the mechanism concept.

[7]This is not to say that finding the support of a mixed Nash is necessarily easy. Finding such a set is achievable in polynomial time for two-player zero-sum games (in which a player's gain is the other player's loss) via linear-programming, but PPAD-complete [19, 29] for more general games with two or more players.

[8]The game in Table 2.1(b) has thus four Nash equilibria, three in pure strategies and one mixed. Wilson [86] establishes that almost all matrix games have an *odd* number of equilibria; the game in Table 2.1(b) is thus *peculiar*.

if there exists an equilibrium strategy vector $(s_1, \ldots, s_n)$ *such that for any private information vector* θ, M *chooses outcome* $\phi(\theta)$ *when the actions are* $s_1(\theta_1), \ldots, s_n(\theta_n)$, *i.e.* $\psi_M(s_1(\theta_1), \ldots, s_n(\theta_n)) = \phi(\theta_1, \ldots, \theta_n)$.

Note that the definition does not specify the exact equilibrium concept, which could be dominant strategy, ex post Nash etc.

In a *direct revelation* mechanism, each agent submits a single message privately to the mechanism: he *reports* his preferences over the alternatives.

Definition 7. *In a* direct revelation mechanism, $W = \Theta$ *i.e. the outcome rule (called the* social choice function*)* ϕ *decides an action for each vector of types:* $\phi : \Theta^n \rightarrow A$.

Note how the set of actions in a direct revelation mechanism is the set of types.

A goal for a direct revelation mechanism is that, in equilibrium, (yet unspecified, just like in Definition 6) each agent has non-negative utility by reporting truthfully.

Definition 8. *A direct revelation mechanism* M *with social choice function* ϕ *is* individually-rational *if no agent* i *obtains negative utility in equilibrium by taking part in* M *when reporting truthfully* $(s_i(\theta_i) = \theta_i)$.

A direct revelation mechanism is incentive-compatible if and only if it is an equilibrium (yet unspecified, just like in Definition 6) for each agent to report truthfully.

Definition 9. *A direct revelation mechanism* M *with social choice function* ϕ *is* incentive compatible *if it implements* ϕ *with truthful reporting* $(s_i(\theta_i) = \theta_i)$.

In this thesis I mostly focus on the strongest type of incentive compatibility, which requires that it is a dominant strategy for bidders to report truthfully.

Definition 10. *A direct revelation mechanism M with social choice function ϕ is* strategyproof *(or truthful*[9]*) if it is incentive-compatible in a dominant strategy equilibrium i.e. every agent i maximizes his utility by honestly reporting his true type θ_i:*

$$u_i(\theta, \phi(\theta_i, \theta'_{-i})) \geq u_i(\theta, \phi(\theta'_i, \theta'_{-i}))\ \forall \theta \in \Theta^n, \theta'_i \in \Theta \tag{2.2}$$

for all reports $\theta'_{-i} \in \Theta^{n-1}$ from the other agents with true types θ_{-i}.

In Eq. (2.2), the utility of each bidder is computed with respect to the true (θ), as opposed to reported (θ'), types, the latter only influencing the mechanism outcome.

Although conceptually simple, direct revelation *incentive-compatible* mechanisms are as powerful as any mechanisms for some equilibrium concepts.

Proposition 1 (Revelation principle)**.** *[37] Suppose that social choice function ϕ is implemented by a mechanism M in a dominant strategy or ex post Nash equilibrium. Then ϕ can be implemented, in the same equilibrium concept, by a direct revelation incentive-compatible mechanism M^d.*

The intuition behind the equivalence of M and M^d is that M^d can simulate bidders' strategies given their types. Therefore if no direct revelation mechanism satisfying a certain property $\mathcal{P}$ can implement a social choice function ϕ then no mechanism at all can satisfy $\mathcal{P}$ and implement ϕ. The revelation principle (Proposition 1) was introduced for implementation in a dominant strategy equilibrium by Gibbard [37].

In the rest of this thesis, I will only consider direct revelation mechanisms.

A direct revelation mechanism M with social choice function ϕ is *dictatorial* if there exists a certain player i such that ϕ always chooses an outcome that maximizes i's payoff: $\forall\, \theta_i, \forall\, \theta_{-i}, \phi(\theta_i, \theta_{-i}) \in \mathrm{argmax}_{a \in A}\, u_i(\theta_i, \theta_{-i}, a)$.

Under very general preferences, only dictatorial functions can be implemented.

[9] By a slight abuse of terminology, "incentive compatibility" may also be used to mean strategyproofness.

Theorem 2. *[37, 81] Suppose that there are at least three outcomes in A and at least two agents with unrestricted preferences: for at least two players i, $u_i(\theta, a) = \theta_i(a)$, where $\theta_i : A \to \mathbb{R}$ can be any such function. Then any social choice function ϕ that can be implemented in a dominant strategy equilibrium and that chooses each outcome for at least one set of preferences must be* dictatorial.

This result becomes less discouraging once one realizes that unrestricted preferences are much more general than any practical domain, in which assumptions on preferences (such as "I prefer more money to less") hold.

It is commonly assumed in economic theory that a special good, the *numeraire*, exists such that anyone can equate their value for a set of goods with a quantity of numeraire. While various goods have served as numeraire through history, money is the most standard.

For the rest of this thesis, any agent i's utility will be modeled as quasi-linear.

Definition 11. *A player i has* quasi-linear *utility if his utility for any outcome a and any payment vector $p = (p_1, \ldots, p_n)$ is*

$$u_i(\theta, a, p) = v_i(\theta, a) - p_i \tag{2.3}$$

where v_i is i's value *for the outcome and the type vector is $\theta = (\theta_1, \ldots, \theta_n)$.*

This definition assumes that no player's utility is affected by others' payments. Any social choice function ϕ in a direct revelation mechanism will have a *payment function* $p = (p_1, \ldots, p_n)$ as component, where $p_i = p_i(\theta, a) : \Theta^n \times A \to \mathbb{R}$. The quasi-linear utility of a player i given reports θ' is then

$$u_i(\theta, \phi(\theta'), p(\theta')) = v_i(\theta, \phi(\theta')) - p_i(\theta', \phi(\theta')) \tag{2.4}$$

Note that payments are ultimately decided by the reports θ' only.

If a direct revelation mechanism $M = (\phi, p)$ implements social choice function ϕ via payment function p then I will say that p *implements* ϕ. Suppose buyers bid for one item. The social choice function "allocate to the buyer who values the item the most" can be implemented in a dominant strategy equilibrium with a second-price payment scheme (the winner pays the next highest bid), as will be seen shortly.

In settings with numeraire, one is naturally concerned about the flow of money.

Definition 12. *Consider a direct revelation mechanism $M = (\phi, p)$ for quasilinear players that report type vector θ'. M is* weakly budget balanced *if $\sum_i p_i(\theta', \phi(\theta')) \geq 0$, i.e. M never incurs a monetary loss. M is* strongly budget balanced *if $\sum_i p_i(\theta', \phi(\theta')) = 0$, i.e. M never incurs a monetary loss or gain. M's* revenue *is the total payment made by the players $\sum_i p_i(\theta', \phi(\theta'))$.*

Note that the weak version of budget balance is less restrictive than the strong one: only a weakly budget balanced mechanism can have positive revenue.

In many mechanisms run by authorities (e.g. governments), the goal may be efficiency, i.e. maximizing players' joint value for the outcome.

Definition 13. *Consider a direct revelation mechanism $M = (\phi, p)$ for quasilinear players.*

Social welfare *for a type vector $\theta \in \Theta^n$ is the sum of players' values for the outcome chosen by ϕ: $\sum_i v_i(\theta, \phi(\theta))$.*

M is efficient *if it maximizes social welfare: $\forall\, \theta \in \Theta^n, \phi(\theta) \in \operatorname{argmax}_{a \in A} \sum_i v_i(\theta, a)$.*

A small set of objectives cannot be jointly achieved

Theorem 3. *[47] In an environment with quasilinear buyers and sellers of single units of a good, no mechanism is simultaneously efficient, weakly budget balanced and strategyproof.*

Thus, there are significant limitations on the sets of properties that a mechanism can achieve, even in simple static environments. In fact, it is quite common in mechanism design to have impossibility results for strong desiderata and then relax some of the requirements in order to obtain possibility results.

Expressiveness is an informal measure of how general a mechanism is: the more expressive a mechanism, the more varied the preferences that agents can express or the more actions they can take. For example, allowing a buyer i's value for an item to depend on other buyers' private information (interdependent values) is more expressive than i's value only depending on i's private information.

This thesis focuses on the limitations imposed by variants of incentive compatibility as strong as possible on the expressiveness and optimization properties of auctions.

There are many other desiderata for a mechanism; I briefly mention some of them that are of diminished importance in the mechanisms in this thesis.

Fairness is the property that all buyers are given, in a sense, equal opportunities and may be beneficial or detrimental to an auction. Fairness in an auction can be defined in multiple ways, see e.g. [39]. *Simplicity* is also an important desideratum since a simple mechanism may be more readily adopted by the players and is less susceptible to failures due to complex, unforeseen, interactions. A mechanism M's *complexity* is a measure of the amount of computation (or communication) that M entails; one seeks to minimize it.

Next section presents a class of functions that maximize a weighted version of social welfare and that are strategyproof on any domain.

2.1.3 Affine maximizers and the VCG mechanism

In this section I will assume that bidders have *private values* and quasilinear utilities, i.e. $u_i(\theta, a, p_i) = \theta_i(a) - p_i$, i.e. $v_i(\theta, a) = \theta_i(a)$ where $\theta_i : A \to \mathbb{R}$. The quasi-linearity of agents' utility renders any function from the following class truthful, i.e. implementable in a dominant strategy equilibrium:

Definition 14. *A social choice function* ϕ *is an* affine maximizer *(of social welfare) if it chooses an alternative that maximizes the weighted sum of players' welfare, adjusted by a constant for each alternative:* $\forall a \in A, \forall i\ \exists \alpha_i \in \mathbb{R}_{\geq 0}, \exists \beta_a \in \mathbb{R}_{\geq 0}$ *such that*

$$\phi(\theta) \in \operatorname*{argmax}_{a \in A} \left\{ \sum_{i=1}^{n} \alpha_i \theta_i(a) + \beta_a \right\} \forall\, \theta \in \Theta^n \tag{2.5}$$

A *dictatorial* social choice function is an affine maximizer with $\alpha_i = 0$ for all players i except one and all $\beta_a = 0$: the alternative chosen is always the one preferred by i, which becomes essentially a dictator.

It is well-known that any affine maximizer can be implemented by a payment function that charges each player the effect he has on the social welfare of the other players.

Proposition 2. *Any affine maximizer is truthful. If* $\phi(\theta) \in \operatorname*{argmax}_{a \in A} \left\{ \sum_{i=1}^{n} \alpha_i \theta_i(a) + \beta_a \right\}$, $\forall\, \theta \in \Theta^n$ *then* ϕ *can be implemented by payment function*

$$p_i(\theta, a) = \frac{1}{\alpha_i} \left(\max_{a' \in A} \left\{ \sum\nolimits_{i' \neq i} \alpha_{i'} \theta_{i'}(a') + \beta_{a'} \right\} \right) - \frac{1}{\alpha_i} \left(\sum\nolimits_{i' \neq i} \alpha_{i'} \theta_{i'}(\phi(\theta)) + \beta_{\phi(\theta)} \right) \tag{2.6}$$

when $\alpha_i > 0$ *and* $p_i(\theta, a) = 0$ *when* $\alpha_i = 0$.

The outcome a', the argmax in Eq. (2.6), is $\phi(\theta_{-i})$, i.e. the outcome that the affine maximizer ϕ would choose if i did not participate. In a dictatorial social choice function, the dictator's payment can be chosen as any constant (0, using Eq. (2.6)) without affecting incentive compatibility.

The **Vickrey-Clarke-Groves (VCG)** mechanism [21, 41, 85] is the most prominent affine maximizer. VCG treats equally all agents and respectively all alternatives: $\alpha_i = 1$, $\forall\, i = 1..n$ and $\beta_a = 0$, $\forall\, a \in A$. Thus player i's payment in the VCG mechanism amounts to the effect it has on others' values for the outcome chosen:

$$p_i(\theta, a) = \max_{a' \in A} \left\{ \sum_{i' \neq i} \theta_{i'}(a') \right\} - \sum_{i' \neq i} \theta_{i'}(a^*) \tag{2.7}$$

When there is one item for sale, the VCG mechanism reduces to

Second-price auction The highest bidder wins and pays the second highest bid.

Let us review an informal proof of this fact. Assume that no two bids are equal and, without loss of generality, that bids are ordered decreasingly: $\theta_1 \geq \theta_2 \geq \cdots \geq \theta_n$. There are $n+1$ outcomes: not allocating to anyone or allocating to each of the n bidders. Each bidder i has positive value (θ_i) only for the outcome in which he wins the item. Thus the VCG outcome a^* must be the outcome in which bidder 1 is allocated. Any bidder $i \neq 1$ must pay 0: a^* is also the optimum outcome without i. The optimum outcome without bidder 1 allocates to 2, the second highest bidder. Thus bidder 1, the winner, pays θ_2, since no other bidder has positive value for the VCG allocation.

The second-price auction with reserve price $\underline{p}$ is also an affine maximizer. If outcome 0 denotes not allocating the item at all then $\beta_0 = \underline{p}, \beta_a = 0, \forall\, a = 1..n$ and $\alpha_i = 1, \forall\, i = 1..n$.

What desiderata are achieved by VCG and to what extent?

Expressiveness. Players can express arbitrary private values for each outcome. There exist generalizations of VCG to interdependent values (where a bidder's value can depend on other bidders' private information) that preserve many of VCG's properties [28, 52].

Efficiency. VCG maximizes social welfare by definition.

Revenue. Krishna and Perry [53] establish that VCG obtains the most revenue among individually-rational mechanisms that implement the outcome maximizing social welfare in a Bayes-Nash equilibrium.

Individual rationality. VCG is individually rational for a truthful bidder i: i's payment in Eq. (2.7) is at most $\theta_i(a^*)$ by definition of a^*.

Truthfulness. Truthful reporting a dominant strategy equilibrium in VCG by Proposition 2. VCG is, however, vulnerable to collusion, i.e. joint manipulations by more than one player. For example, suppose that the highest three values in a second-price auction were $50, 40, 20$.

When bidders are truthful, the 50 bidder wins and pays 40 while the 40 and 20 bidders have utility 0. If the 50 bidder paid, say $5, to the 40 bidder to lower his bid to 20, the 50 bidder's price would be 20 instead of 40. Thus the 50 bidder would save $15 and the 40 bidder would receive $5, which is a strict improvement over their utilities when truthful.

Fairness. All agents have equal weight in VCG.

Simplicity. i's VCG payment can be interpreted as i's least value for a^* such that the outcome chosen remains a^*, which I consider conceptually simple. Paying the second highest bid for an item is different from the traditional payment rule of an auction, in which the winner pay his bid and thus may be considered sophisticated by some bidders.

Computational complexity. In the case of one item to sell, the second-price auction winner and his price can be easily computed in linear time. The *winner determination* problem (WDP) is concerned with the optimal way of allocating K heterogeneous items to n bidders. Then WDP is $\mathcal{NP}$-complete [59] even in the following simple case: any reported type assigns a value of 1 to exactly one bundle of items and 0 to the others and every item is contained in the bundles of exactly two bidders. WDP cannot be approximated by a polynomial algorithm [80] to a $K^{\frac{1}{2}-\varepsilon}$ factor, for any $\varepsilon > 0$, unless $\mathcal{NP} = \mathcal{ZPP}$. Recall that $\mathcal{ZPP}$ is the subclass of problems in $\mathcal{NP}$ that can be solved by a randomized algorithm with expected polynomial runtime.

In the next section I review characterizations of truthfulness.

2.2 Static characterizations of incentive compatibility

When characterizing truthfulness, there are two levers one can press. One can restrict the preference space and thus get a rich class of truthful functions or one can impose constraints on the choice function in order to obtain truthfulness on less restricted domains. Why does the domain of possible agent preferences influence the space of truthful functions? Because a truthful mechanism needs to motivate the agents to be honest and this becomes more difficult as the agents have more ways to manipulate.

I will review several results in the literature showing characterizations for certain positions of one lever or the other [42, 54, 60, 77, 78, 79]. However, the preference domains employed by these results are either very comprehensive (unrestricted preferences or order-based domains) or very restrictive (e.g. single-minded). As will be argued, many domains of practical importance are structured and are not captured by existing results.

A particular (negative) result I will focus on is presented in [54]: on order-based domains, under several technical conditions, the only truthful functions are affine maximizers. I will demonstrate the limited applicability of this result by presenting numerous examples of relevant domains that are not order-based.

One of the first results characterizing truthful functions is due to Roberts [77]:

Proposition 3. *If no restriction whatsoever is imposed on the agents' preferences (i.e. agents have* unrestricted preferences*), then truthful social choice functions and affine maximizers coincide.*

Thus, in unrestricted domains, any truthful function must be an affine maximizer. However, most real-world domains are structured: for instance, agents prefer allocated to not being allocated in an auction, which restricts the functions' domain, thus rendering Proposition 3's result inapplicable.

In the following I will only consider *deterministic* social choice functions (scf's).

A well-known characterization in terms of payment functions (see, for example, [54]) is that an scf ϕ is truthful if and only if there exists an *agent-independent* payment function that implements ϕ, i.e. given the outcome chosen, i's payment does not depend on his report.

Proposition 4. *Social choice function ϕ is truthful if and only if there is an* agent-independent *payment function p ($p_i : \Theta^{n-1} \times A \to \mathbb{R} \cup \{\infty\}$) that implements it such that for all $\theta \in \Theta$ and i, $\phi(\theta) \in \operatorname{argmax}_{a \in A}\{\theta_i(a) - p_i(\theta_{-i}, a)\}$.*

I provide an informal proof for the necessity of the existence of an agent-independent payment function, focusing on the intuition behind this result.

Since ϕ is truthful there must exist some payment function $p : \Theta \times A \to \mathbb{R} \cup \{\infty\}$ that implements it. Suppose agent i's price for θ_i is less than the price for θ'_i when the others' reports are fixed to θ_{-i} and $\phi(\theta_i, \theta_{-i}) = \phi(\theta'_i, \theta_{-i})$. Then, when the agent has preference θ'_i, he can gain utility by misreporting as θ_i. If the price vector p did not maximize i's utility then i could simply benefit by misreporting in a way that gives him the maximum utility given p.

Following [54], I denote, for agent i and alternatives $a, b \in A$,

$$\delta_{ab}[\theta_{-i}] = \inf_{\theta_i \in \Theta : \phi(\theta_i, \theta_{-i}) = a} \{\theta_i(a) - \theta_i(b)\}$$

where alternative a is assumed to be chosen for at least one type of i given others' types θ_{-i}. This quantity measures the minimal difference between the value of outcomes a and b for i. We will shortly see that these quantities play a central role in truthfulness.

Consider the complete directed graph $\Gamma_\phi[\theta_{-i}]$ that has the alternatives in A as vertices and weight $\delta_{ab}[\theta_{-i}]$ on edge (a, b) for all a, b. Call $\Gamma_\phi[\theta_{-i}]$ ϕ's *outcome graph* for agent i and θ_{-i}, reports of the other agents.

A general necessary and sufficient characterization of truthfulness is

Theorem 4. *[42] scf ϕ is truthful if and only if $\Gamma_\phi[\theta_{-i}]$ does not have negative cycles of* any *length $\forall i, \forall \theta_{-i}$.*

For a truthful ϕ, a useful quantity is the least weight of a path in Γ_ϕ between alternatives a and b:

$$\delta^*_{ab}[\theta_{-i}] = \inf_{(a_0=a,a_1,\ldots,a_k,a_{k+1}=b)} \sum_{j=0}^{k} \delta_{a_j a_{j+1}}[\theta_{-i}] \tag{2.8}$$

(I will drop the θ_{-i}'s when they are clear from the context). By Theorem 4, if ϕ is truthful and A is finite then $\forall\, a, b \in A, \delta^*_{ab} > -\infty, \delta^*_{ab} \le \delta_{ab}$ (as δ_{ab} is obtained in Eq. (2.8) for $k = 0$) and $\delta^*_{ab} + \delta^*_{ba} \ge 0$. In particular, if ϕ is truthful, outcome a can be priced as $p_i(a) = \delta^*_{aa_0}$, for an arbitrary $a_0 \in A$. With this payment function, i prefers being truthful, i.e. prefers obtaining the truthful outcome a to any other outcome b: $\theta_i(a) - \theta_i(b) \ge \delta_{ab} \ge \delta^*_{aa_0} - \delta^*_{ba_0} = p_i(a) - p_i(b)$.

The characterizations in Proposition 4 and Theorem 4 offer rather little insight into the structure of truthful functions. Other characterizations that will be reviewed (Theorem 6 and Theorem 7) are easier to work with, but have limited applicability.

It is well-known [78, 54] that truthfulness can be characterized as follows

Theorem 5. *If the scf ϕ is truthful then it is also weakly monotonic.*

Definition 15. *The social choice function ϕ satisfies* weak monotonicity *(WMON) if* $\forall i, \forall \theta_{-i} \in \Theta^{n-1}$, $\forall \theta'_i, \theta_i \in \Theta$

$$\left.\begin{array}{l} \phi(\theta_i, \theta_{-i}) = a \\ \phi(\theta'_i, \theta_{-i}) = b \end{array}\right\} \Rightarrow \theta'_i(b) - \theta_i(b) \ge \theta'_i(a) - \theta_i(a) \tag{2.9}$$

That is, if the alternative chosen by ϕ changes from a to b and only agent i has changed his report, from θ_i to θ'_i, then i's relative increase in value must be weakly higher for the new alternative b than for the old alternative a. Intuitively, this condition makes sense as it requires that if only one agent has changed the social outcome then a shift in the relative values of the old and new outcomes has occurred for that agent. In terms of $\delta_{\bullet\bullet}$'s, WMON is equivalent to Γ_ϕ not having negative cycles of length two. Thus Theorem 5 is actually a corollary of Theorem 4.

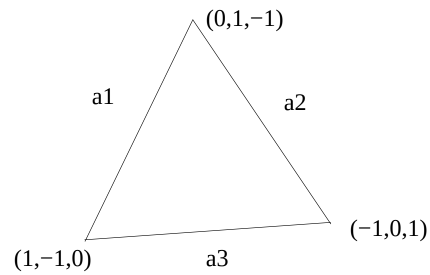

Figure 2.1: A weakly monotonic, but not truthful single-player function on a non-convex domain that is a subset of $\mathbb{R}^3$. The preference domain consists only of the line segments shown - it is therefore not convex. Any point on a line segment represents a possible preference type for the agent. There are only three alternatives: $A = \{\mathrm{a1}, \mathrm{a2}, \mathrm{a3}\}$. A point's ith coordinate gives agent's value at that preference type for alternative ai, $i = 1, 2, 3$.
The function maps each preference type into an alternative that only depends on the line segment the preference type lies, as shown. For this function, one can show that $\delta_{\mathrm{a}i,\mathrm{a}(i+1)} = 1$ and $\delta_{\mathrm{a}(i+1),\mathrm{a}i} = -1$. The function's outcome graph G has the negative weight cycle a1 $\rightarrow$ a2 $\rightarrow$ a3 $\rightarrow$ a1: $\delta_{\mathrm{a1a2}} + \delta_{\mathrm{a2a3}} + \delta_{\mathrm{a3a1}} = -3$, therefore the function is not truthful. However, $\delta_{\mathrm{a}i,\mathrm{a}j} + \delta_{\mathrm{a}j,\mathrm{a}i} = 0$ for all distinct i, j, i.e. the function is weakly monotonic.

Saks and Yu [79] establish the converse of Theorem 5 for convex domains of preferences.

Theorem 6. *Any WMON scf on a convex domain of preferences is also truthful.*

Recall that a preference assigns a real value to each outcome. The domain of preferences Θ is then convex if for any two types $\theta, \theta' : A \rightarrow \mathbb{R}$ with $\theta, \theta' \in \Theta$ any linear combination of the two is also in Θ ($\alpha\theta + (1-\alpha)\theta' \in \Theta$ for any $\alpha \in [0, 1]$).

The example in Figure 2.1 (also from [79]) shows that for non-convex domains, the result of Theorem 6 may no longer hold.

In Figure 2.2, the outer-cone illustrates the space of WMON functions. For the domain of unrestricted preferences, WMON is also sufficient for truthfulness.[10] Seemingly positive, this result which says that *all* WMON functions are truthful for unrestricted preferences, turns negative when one realizes that the space of WMON functions is *exactly*

[10]This is a corollary of the sufficiency of WMON in convex domains [79] (Theorem 6).

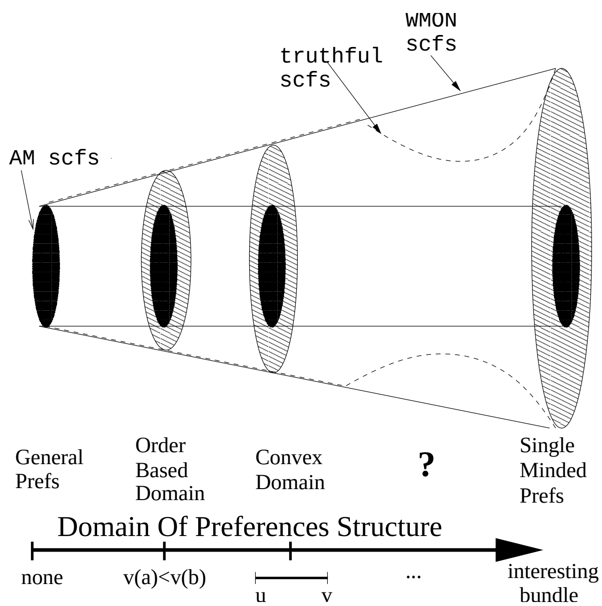

Figure 2.2: Illustration of the inclusion relations between subclasses of WMON functions on various domains. WMON is equivalent to truthful for order-based or convex domains. First example in Section 7 of [79] provides a function that is WMON but not truthful for a non-convex domain. Roberts' result [77] states that only affine maximizers (AMs) are truthful for the domain of unrestricted preferences. Lavi et al. [54] prove that if the domain is order-based then, under several technical conditions, any WMON scf must be an (almost) affine maximizer. For single-minded preferences (a particularly structured domain), [60] show that monotonicity is equivalent to truthfulness and thus the class of truthful functions is richer than the class of affine maximizers.

that of affine maximizers for unrestricted preferences (recall Prop. 3):

$$\textit{unrestricted preferences:} \quad \text{WMON} \iff \text{truthful} \iff \text{affine maximizer}$$

This is not the case for more structured preference domains.

Affine maximizers are illustrated in Figure 2.2 through the central (cylindrical) volume. So, for general (unrestricted) preferences (to the left) the transversal sections of this volume and of the outer WMON cone coincide. On the other hand, Lehmann et al. [60] have shown that while WMON was also sufficient for single-minded bidders, the class of WMON functions included non affine-maximizers for that preference domain. Thus, the WMON outer-cone is larger than the AM cylinder for single-minded preferences (to the right).

The following result follows from Proposition 2 and Theorem 5.

Proposition 5. *If scf ϕ is an affine maximizer then ϕ is also weakly monotonic.*

As an affine maximizer is a truthful function, all cycles, not only the length two ones, of its outcome graph have non-negative weight (see Theorem 4). For the domain of unrestricted preferences, one can show that all cycles of the outcome graph of an affine maximizer have weight 0.

Order-based (OB) domains, introduced by Lavi et al. [54], allow some structure in preferences. As will be seen shortly, essentially any reasonable scf on an order-based domain is an affine maximizer. This characterization is the motivation for the discussion in Section 2.3.

Definition 16. *A domain Θ is* order-based *(OB) if it is defined by a set $\mathcal{O}$ of (in)equalities of the form*

$$a \ \{<_{\mathcal{O}}, \leq_{\mathcal{O}}, =_{\mathcal{O}}\} b \quad (\textit{or } a =_{\mathcal{O}} 0), \textit{ where } a, b \in A \textit{ meaning that}$$

$$\theta_i(a) \ \{<, \leq, =\} \ \theta_i(b) \quad (\textit{or } \theta_i(a) = 0), \textit{ for all } \theta_i \in \Theta.$$

That is, all $\theta_i \in \Theta$ and only them *satisfy the (in)equalities in $\mathcal{O}$. A joint domain Θ^n is order-based*[11] *if Θ is order-based.*

Let us emphasize the point in the definition once again. $\mathcal{O}$ specifies a set of inequalities and then *all* preferences satisfying these inequalities and *only them* are in the order-based domain specified by Θ.

The domain of unrestricted preferences is trivially OB by using no constraints, i.e. $\mathcal{O} = \emptyset$.

[11]If each player i's type θ_i belongs to a different type space Θ_i then the defining inequalities may differ between players.

In a combinatorial auction (CA), the auctioneer has a set of heterogeneous indivisible items to sell to the n agents. An alternative $a = (a_1, \ldots, a_n)$ in a CA is an allocation of the goods auctioned with agent i getting bundle a_i such that the a_i's are disjoint (a_i can be the empty bundle). The domain of (general) CA preferences is order-based, being defined by the set of inequalities $\mathcal{O}_{\mathrm{CA}}$ described below:

- No externalities. ("I only care about my own bundle"): $\theta_i(a) = \theta_i(b)$, $\forall\, a, b \in A$ with $a_i = b_i$
- Free disposal. ("getting more items never hurts me"): $\theta_i(a) \leq \theta_i(b)$, $\forall\, a, b \in A$ with $a_i \subseteq b_i$
- Normalization. ("getting nothing has value 0 for me"): $\theta_i(a) = 0$, $\forall\, a \in A$ with $a_i = \emptyset$

Lavi et al. [54] establish that WMON is sufficient for truthfulness in OB domains. As already mentioned in Theorem 6, Saks and Yu [79] extend[12] this result to the more expressive class of *convex* domains[13]. Note that a domain in one of these two classes is "rich" in the sense that it guarantees the inclusion in the domain of an infinite set of preferences with certain properties. More importantly, truthful functions in order-based domains are restricted to affine maximizers.

Theorem 7. *[54] In an order-based domain, any truthful social choice function that has a dense range and satisfies unanimity, decisiveness and weak-IIA is an* almost *affine maximizer.*

[12]It is not hard to prove that any order-based domain is convex. However, the converse is not true: there are domains that are convex but not order-based (e.g. the domain of linear threshold preferences - see Subsubsection 2.3.2). In fact, if A is finite then there are only a finite number of orderings $\mathcal{O}$ on A's alternatives and thus only a finite number of order-based domains with space of alternatives A.

[13]I consider convex domains in the topological sense and not domains in which preferences that are convex functions themselves.

Theorem 7 is thus an extension of Proposition 3 since the domain of unrestricted preferences is (trivially) order-based.

I present what Theorem 7's technical conditions become for the order-based domain of combinatorial auctions preferences. An scf ϕ has a dense range if it does not always sell all items to a single bidder. An scf ϕ is unanimous if, when each player i is single-minded with interesting bundle L_i and allocating all players simultaneously is feasible ($L_i \cap L_{i'} = \emptyset$, $\forall\, i \neq i'$), ϕ allocates each player i his bundle L_i. An scf ϕ is decisive if, for any fixed types θ_{-i} of the other players and for any bundle L_i, ϕ allocates a single-minded player i his interesting bundle L_i provided i bids high enough on it. An scf ϕ satisfies weak independence of irrelevant alternatives (weak-IIA) if, when i causes a flip in ϕ's *overall* allocation due to a change in his report, it must be that i reported a different value for his old bundle or for his new bundle. That is, ϕ satisfies weak-IIA if whenever $\phi(\theta_i, \theta_{-i}) = a$ and $\phi(\theta'_i, \theta_{-i}) = b \neq a$, it holds that $\theta_i(a) \neq \theta'_i(a)$ or $\theta_i(b) \neq \theta'_i(b)$. Weak-IIA is implied by SMON, a version of WMON (also introduced in [54]) in which the inequality in Eq. (2.9) is strict.

An scf f is an *almost* affine maximizer if it is an affine maximizer for all sufficiently large values of players: there exists a threshold M and $\forall a \in A, \forall i,\ \exists \alpha_i \in \mathbb{R}_{\geq 0}, \beta_a \in \mathbb{R}_{\geq 0}$ such that if $\theta_i(a) \geq M$, $\forall\, a \in A \setminus \{0\}$ then $\phi(\theta) \in \mathrm{argmax}_{a \in A} \{\sum_{i=1}^{n} \alpha_i \theta_i(a) + \beta_a\}$

In the following section I analyze closely the notion of "order-based", finding it quite restrictive.

2.3 The limited expressiveness of order-based domains

In this section I show, by analyzing numerous domains we find important, that the assumption of order-based domains is quite limiting and thus the negative result of [54] is

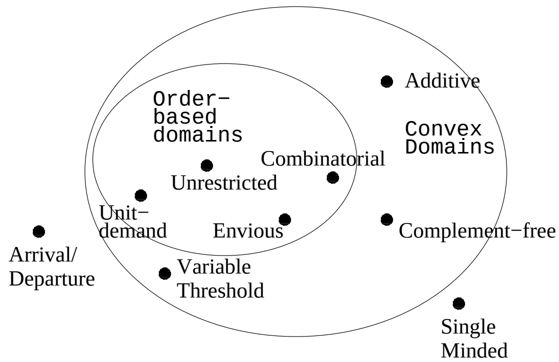

Figure 2.3: Classes of preference domains and example domains.

less powerful than it may first seem. Characterizations that go beyond order-based domains appear now even more important.

2.3.1 Combinatorial auctions with structure

In the following two subsections I investigate whether several examples of preference domains are order-based and thus illustrate the (limited) expressiveness of order-based domains. Figure 2.3 provides an overview of the considerations in this section.

I first consider domains of preferences in combinatorial auctions (CAs) but with additional preference structure. In the CA setting, a very restricted preference domain is the one of single-minded preferences. In this domain, an agent has value only for a particular set of items, the others being irrelevant.

Definition 17. *A bidder i is* single-minded *if there exists $v_i \in \mathbb{R}$ and a "interesting bundle" L such that*

$$\theta_i(a) = \begin{cases} v_i & \text{if } L \subseteq a_i \\ 0 & \text{otherwise} \end{cases} \quad \forall a \in A$$

If $|L| = 1$ then I say that i has an Unknown-Item-Unknown-Value (UIUV) *preference.*

Note that all single-minded preferences certainly satisfy the constraints in $\mathcal{O}_{\text{CA}}$.

However, the domain of single-minded preferences is not order-based.[14]

In fact, the following stronger result holds, providing a powerful tool for characterizing several of the domains of interest as not order-based

Proposition 6. *Let Θ be a preference domain that satisfies the constraints in $\mathcal{O}_{CA}$, contains all UIUV preferences but does not contain all possible CA preferences. Then Θ cannot be order-based.*

Proof. Suppose for a contradiction that Θ was order-based and let $\mathcal{O}$ be the set of constraints that defines it. Then $\mathcal{O}$ must contain all constraints for CAs $\mathcal{O}_{\text{CA}}$. However, only these are not enough as Θ is not the domain of all possible CA valuations. Thus $\mathcal{O}$ must contain at least one inequality of the form $\theta_i(a)\,\{<,\leq,=\}\,\theta_i(b)$ where $a_i \not\subseteq b_i$. Assume without loss of generality that this relation is $\theta_i(a) \leq \theta_i(b)$. This inequality is not satisfied by any UIUV preference θ for item x with $x \in a_i \setminus b_i$ that has value 0 for bundle b_i and some positive value for bundle a_i, contradiction. $\square$

The domain of UIUV preferences in the proposition can be replaced with any domain D with the following property: for any alternatives a and b there exist valuations $\theta, \theta' \in D$ such that $\theta(a) > \theta(b)$ but $\theta'(a) \leq \theta'(b)$. The key property in the proof is preserved: no ordering inequality that holds for all $\theta \in D$ can be required for any two alternatives a and b.

I use Proposition 6 extensively in the rest of the section. For now, note:

Corollary 1. *Single-minded preferences are not order-based.*

Also by Proposition 6, if some preference domain D_1 is not order-based and $UIUV \subseteq D_2 \subset D_1$, then D_2 is not order-based.

[14] It is also easy to see that the domain of single-minded preferences is not convex. Note that the domain of *known* single-minded preferences (where the interesting bundle is the same for all preferences) is order-based.

Complement-free preferences

Again in a CA setting, consider the domain of complement-free preferences, i.e. preferences which do not consider synergies between items: a set of items is never worth more than the sum of its parts (see [58] for a characterization of this domain):

$$\theta_i(a_i) + \theta_i(b_i) \geq \theta_i(a_i \cup b_i) \; \forall \, a_i \cap b_i = \emptyset \tag{2.10}$$

The domain of complement-free preferences is not order-based as it contains all UIUV preferences and Proposition 6 applies.

Additive preferences

A valuation θ_i is additive (OR) if there exist a number l and l "atoms" (bundles) $B^{(1)}, \ldots, B^{(l)}$ with values $s^{(1)}, \ldots, s^{(l)}$ such that the value of any bundle a_i is the value of the maximum packing of the atoms in a_i (or 0 if no atom is included in a_i). Formally, the evaluation function for $a \in A$ is: $\theta_i(a)$ equals 0 if $\forall \, 1 \leq h \leq l$, $B^{(h)} \not\subseteq a_i$ or $\max \sum_{j=1}^{h} s^{(j)}$ over all disjoint decompositions $B^{(1)} \cup \cdots \cup B^{(h)} \subseteq a_i$ otherwise. The domain of additive preferences is not order-based because it contains all UIUV preferences but is not fully expressive, and thus Proposition 6 applies.

Corollary 2. *Additive preferences are not order-based.*

2.3.2 Multi-unit auctions with structure

Consider a multi-unit auction (MUA) setting, where K identical items are for sale to the n agents. An allocation a here is defined by a vector of values $a_1, \ldots, a_n$ where a_i represents the number of items allocated to agent i and $a_1 + \cdots + a_n \leq K$. Out of the three axioms for CAs I will just assume normalization ("If I get no items then I have value 0 for this allocation").

Envious preferences

Consider a preference such as: "I have positive value for an allocation only if my share is the biggest", i.e. $\theta_i(a) > 0$ iff $a_i \geq a_j \, \forall j \neq i$. I call these *envious preferences.* The set $\mathcal{O}$, that establishes that envious preferences are order-based, contains all inequalities of the form:

- $\theta_i(a) = 0$ iff $\exists \, j \neq i$ s.t. $a_i < a_j$
- $\theta_i(a) > \theta_i(\emptyset) = 0$ iff $\forall j \neq i$, $a_i \geq a_j$

Free-disposal does not hold for envious preferences: if there are 3 agents and 20 items agent 1 has positive value for the allocation vector (8, 7, 5) but zero value for (9, 11, 0). However, one can add "limited" free-disposal to envious preferences, in the sense that $a_i' \geq a_i \geq a_j \ \forall j \Rightarrow \theta_i(a_i') \geq \theta_i(a_i) > 0$. Envious preferences with limited free-disposal remain order-based.

Variable-threshold preferences

Again, in a multi-unit auction, consider preferences like "I have positive value for an allocation if and only if I get at least r items", where r may vary and is not known beforehand (to the auctioneer). A formal definition of the domain THR of variable-threshold preferences is: $\theta_i \in$ THR iff $\exists \, 1 \leq r = r(\theta_i) \leq k$ such that $\theta_i(a) = 0$ if $a_i < r$ and $\theta_i(a) > 0$ if $a_i \geq r$.

The following result shows that variable-threshold preferences are not order-based. The proof is different than the immediate proofs so far and it attacks the notion of order-based directly. The proof uses the fact that all inequalities for normalization, free-disposal and no externalities should be used in defining this domain and shows that no other constraint can be added.

Proposition 7. *The domain* THR *is not order-based.*

Proof. Suppose for a contradiction that THR was order-based, defined by the set of inequalities $\mathcal{O}$. Note first that any variable-threshold valuation satisfies normalization, free disposal and no externalities so all the inequalities characterizing those properties must be in $\mathcal{O}$. However, those are not sufficient as there are many MUA preferences that are not variable-threshold. Let us analyze what kind of constraints can be added to the ones already in $\mathcal{O}$.

- Can any inequality of the form $\theta_i(a) < \theta_i(b)$ for some $a, b \in A$ be added? Because of free disposal and no externalities, $a_i < b_i$. The variable-threshold valuation that has value 1 for any allocation that assigns at least a_i items to agent i does not satisfy this inequality, so one cannot add such an inequality.

- Can one add any constraint of the form $\theta_i(a) = \theta_i(b)$ for $a_i < b_i$? No, as any variable-threshold valuation with b_i as threshold does not satisfy this constraint.

By a similar argument, no constraint can be added to $\mathcal{O}$. Therefore one cannot construct a set of constraints that defines THR, i.e. THR is not order-based. □

Thus, the negative result does not apply to resource-allocation settings with this "variable-threshold" structure.

Suppose, however, that the threshold r is fixed and does not vary from one preference to the other. For example, if $r = 1$, such preferences are known as *unit-demand* preferences. Then the domain of fixed-threshold preferences is order-based.

In this thesis, variable-threshold preferences will be encountered in Chapter 3, while unit-demand preferences will be commonplace.

2.3.3 Arrival/Departure preferences

In dynamic settings, one can define an outcome as the time when an agent is served. Then preferences like "I have positive value for outcome t if and only if $t_1 \leq t \leq t_2$", (where t_1 and t_2 are interpreted as arrival and departure times) are again natural. I call such preferences arrival/departure preferences. It is easy to prove that these preferences are not order-based: just note that for all t_1, t_2 one can find some preferences u that give positive value to t_1, zero value to t_2 and vice versa. Therefore no order-based inequalities can be put forth for this domain. As this domain is not the domain of unrestricted preferences, I conclude that it is not order-based.

2.4 Distribution-based environments and mechanisms

Thus far, I have made no assumptions about what the center or other players know about a player's private information, an important informational aspect in many settings.

A standard way to model uncertainty (players' types, in this thesis) is to assume that uncertain quantities are draws from a probability distribution. When using a probability distribution F on bidders' types (Chapters 3 and 5), I will assume that each bidder i's type is drawn *independently* from F's i-th component F_i. I will assume that each player j has the *same* prior F_i, on i's type but the priors on players' i and j types may differ. I will assume that these priors are *common knowledge*, i.e. each player knows the priors, each player knows that any other player knows the priors etc. Notationally, the expectation $\mathbb{E}_{\theta_{\mathcal{I}}}$ on $\theta_{\mathcal{I}}$ for a subset $\mathcal{I} \subseteq 1..n$ of players is taken on the joint probability distribution $F_{\mathcal{I}}$.

Given a model of uncertainty (a probability distribution on the other types in this case), a natural goal for a player is to maximize his *expected* payoff.

Definition 18. *An n-tuple of strategies $(s_1, \ldots, s_n)$ is a* Bayes-Nash equilibrium *in a dy-*

namic revelation mechanism if, for any player i, playing s_i is, in expectation *over the prior F_{-i}, a best-response to others playing strategies s_{-i}:*

$$\mathbb{E}_{\theta_{-i}}\left[u_i(\theta, \phi(s_i(\theta_i), s_{-i}(\theta_{-i})))\right] \geq \mathbb{E}_{\theta_{-i}}\left[u_i(\theta, \phi(s_i'(\theta_i), s_{-i}(\theta_{-i})))\right] \;\forall\, \theta_i, \forall\, s_i' \colon \Theta \to \Theta$$

The stochasticity here comes from players' priors; each s_i may be a pure strategy.

The revelation principle was extended to the Bayes-Nash equilibrium concept by Green and Laffont [40]. Myerson and Satterthwaite [66] give a generalization of Theorem 3's impossibility result to incentive compatibility in Bayes-Nash equilibrium.

2.4.1 Expected externality (AGV) mechanism

The Arrow-d'Aspremont-Gérard-Varet [3, 30] (AGV, also called "expected externality") mechanism achieves strong budget balance and implements the efficient outcome. It is however only incentive-compatible in a Bayes-Nash equilibrium. It charges agent i

$$p_i^{AGV}(\theta) = \left(\tfrac{1}{n-1}\sum_{j \neq i} \mathbb{E}_{\tilde{\theta}_{-j}}\left[V_{-j}(\theta_j, \tilde{\theta}_{-j})\right]\right) - \mathbb{E}_{\tilde{\theta}_{-i}}\left[V_{-i}(\theta_i, \tilde{\theta}_{-i})\right] \tag{2.11}$$

where $V_{-j}(\theta) = \sum_{h \neq j} v_h(\theta, \phi(\theta))$ and $\phi(\theta) = \mathrm{argmax}_{a \in A} \sum_i v_i(\theta, a)$ is the efficient outcome.

Suppose that players have symmetric priors ($F_i = F$ and $f_i = f$) and one item is for sale to bidders 1 and 2. Bidder 1's AGV payment is $p_1^{AGV} = \int_{\theta_2}^{\infty} r_1 f(r_1) dr_1 - \int_{\theta_1}^{\infty} r_2 f(r_2) dr_2$ while bidder 2's payment is $-p_1^{AGV}$. Thus the higher bidder effectively pays the lower bidder in this simple instance of the AGV mechanism. If there were at least three bidders (with symmetric priors) for one item, then in the AGV mechanism the lowest bidder is paid, the highest bidder pays, but other bidders may pay or be paid depending on the bids.

The AGV mechanism satisfies a very weak form (called *ex ante*) of individual rationality: a bidder's expected payment *if he did not know his or any other bidder's type* is 0 (this informational setting is illustrated in Tables 2.1(a) and 2.2(a)). The AGV mechanism strikes a different tradeoff than VCG between desirable properties. The AGV mechanism achieves strong budget balance at the cost of weakening the form of individual rationality and incentive compatibility.

2.4.2 Revenue-optimality for single-item auctions

The *virtual valuation* (introduced by Myerson [64]) of a bidder with bid r_i is

$$\tilde{w}_i(r_i) = r_i - \frac{1-F_i(r_i)}{f_i(r_i)} \tag{2.12}$$

An auction $A_{\tilde{w}}$ defined on "virtual valuations" is revenue-optimal [64] among static incentive-compatible single-item auctions that use a prior for bidders with private values.

Suppose that $\tilde{w}_i(r_i)$ is non-decreasing in r_i, for example when F_i has a non-decreasing hazard rate $\frac{f_i(r_i)}{1-F_i(r_i)}$. The revenue-optimal auction $A_{\tilde{w}}$ allocates to the bidder with the highest virtual valuation, but only if it is non-negative (otherwise the item remains unsold). This bidder pays his lowest value for which he would still win. Myerson also provides an (input) ironing technique for rendering an arbitrary $\tilde{w}_i(r_i)$ non-decreasing in r_i.

$A_{\tilde{w}}$ penalizes to a higher extent bids from *a priori* more competitive bidders. For symmetric priors, $A_{\tilde{w}}$ is a second-price auction with a reserve price of $\tilde{r}$ such that $\tilde{w}_i(\tilde{r}) = 0, \forall\, i$, i.e. the highest bidder wins provided it bids at least $\tilde{r}$ and pays the next highest bid or the reserve price, whichever is larger.

2.5 Dynamic environments

Static mechanisms faced the challenge of aggregating preferences with certain objectives in mind when all the information was available at once. New challenges appear when designing *dynamic* mechanisms:

- there is uncertainty about the future, for example via new information entering the market. Furthermore, the private information of players currently in the market may change (such a possibility is not considered in this thesis).
- bidders may have temporal manipulations available,
- computation (for example, of optimal policies) may be significantly more difficult in the dynamic case than in the static case.

I introduce now a general model for dynamic allocation environments and mechanisms.

2.5.1 Dynamic model

In the domains in this thesis I will consider a dynamic population of agents with static types over T periods of time.[15] That is, each agent arrives and departs (exactly once) in the interval $1..T$, but while present, his type does not change with time. I will identify the agents present at time t with their types (assuming for ease of notation that no two agents have the same type) and denote them by θ^t.

The supply of items will be deterministic, but may change with time, in which case the supply dynamic is known to the center.

Each agent has private information; his valuation (a publicly known mapping) aggregates all bidders' private information into a (real) value for a number of items being given to or taken away from (after being previously allocated to) the agent. This valuation on changes in allocation in each period induces a valuation over a sequence of decisions.

Myerson [65] establishes a dynamic version of the revelation principle. In a domain with bidders with private information that does not change over time and arrival-departure

[15] Chapter 4's model does not require T to be known or finite.

intervals, it states that any social choice function ϕ that can be implemented by a dynamic mechanism M can also be implemented by a direct revelation mechanism M^d, in which each agent makes a single claim about its type. The time of this claim becomes the agent's reported arrival. An agent's true arrival is construed as the time he learns about the auction or he conceptualizes his valuation for being allocated. Unless otherwise specified, I will assume that a bidder cannot report an arrival that is earlier than his true arrival. Similarly, I will assume that a bidder needs to have an allocation decision no later than his true departure and thus cannot report a later departure time.

I now formally define the models and mechanisms studied in this thesis and then exemplify how the definitions apply to several domains, in particular the settings of Chapters 3, 4 and 5.

Definition 19. *A* dynamic allocation environment *over periods* $1..T$ *consists of a seller (also called auctioneer or center) and, for each time step* $t \in 1..T$*,*

- *a supply* S^t *of indivisible* items *owned by the seller* before *time* t
- *a set of active players* θ^t *where player* i *has private information* $\theta_i \in \Theta$ *that includes his arrival* a_i *and departure* d_i*.*

  Player i *can report a type* $\hat{\theta}_i$ *to the seller at no more than one time of his choosing* $t_i \in 1..T$*, which becomes his reported arrival:* $\hat{a}_i = t_i$*. As already mentioned, a player can only report an activity interval that is tighter than his true one:* $a_i \leq \hat{a}_i \leq \hat{d}_i \leq d_i$*. The seller considers any player* active *in his reported arrival-departure interval:* $\theta^{t+1} = \theta^t|_{\hat{d} \geq t+1} \cup \theta|_{\hat{a}=t+1}$*. Thus active players are ones that have reported their arrival, but that claim to not have departed yet.*

  I denote i*'s (non-negative) number of allocated items* before *any*[16] *time* t *by* l_i^t*. No*

[16] To avoid the conceptual problem of the seller monitoring a player i before i reported anything, one can

player holds any item before he arrives: $l_i^t = 0, \forall\, i, \forall\, t = 1..\hat{a}_i$. *All allocations must be feasible; in particular, for persistent, non-reusable, items,* $\sum_{i \in \theta^t} l_i^{t+1} = |S^1| - |S^{t+1}|, \forall\, t \geq 1$.

Each player values an update between his number of allocated items l_i^t *(before* t*) and* l_i^{t+1} *(before* $t+1$*) at*

$$\delta_{a_i,d_i}(t) \cdot v_i(\theta^{1..t}, l_i^t, l_i^{t+1} - l_i^t).$$

The valuation v_i*, whose functional form is public knowledge, aggregates private information (i's own and possibly other players'). Thus i's time-independent value for an allocation decision is multiplied by a time-dependent coefficient, which I will assume to be* 0 *outside i's* true *arrival-departure interval:* $\delta_{a_i,d_i}(t) = 0, \forall\, t \notin a_i..d_i$*, modeling a bidder's interest in allocations only within this interval. It is encoded in i's valuation (and thus public knowledge*[17]*) that i is only interested in a subset* L_i *of the items* ($L_i \subseteq \cup_{t=1}^{T} S^t$)*, called his* interesting set.

The chronological order of events for Definition 19 is as follows. First, the supply S^t and holdings l^t before time t are tallied. Second, the set of active players at time t is updated from the one at time $t-1$: players with reported departures of $t-1$ exit the environment and players reporting an arrival of t enter it. Finally, updates π^t are made.

Definition 19 does not allow *allocation* externalities, that is dependence of a player's value on another player's number of allocated items. It allows however *information* externalities, where a player's value depends on the information of other players, as in Chapter 5. Fixing all players' types, a bidder i's value for an update to his number of allocated items

restrict l_i^t's definition to i's activity interval $[\hat{a}_i, \hat{d}_i]$, at the expense of making the other definitions more complicated.

[17]This informational assumption is trivially satisfied in Chapters 3 and 5. Without it, however, there are serious limitations on interesting mechanisms in Chapter 4 (see Example 9).

at time t (changing i's number of allocated items to l_i^{t+1}) is independent of previous updates and only depends on the current time via the $\delta_{a_i,d_i}(t)$ coefficient. Unless otherwise mentioned, I will assume that i is indifferent about the time of his allocation: $\delta_{a_i,d_i}(t) = 1$ for $t \in a_i..d_i$ (recall that $\delta_{a_i,d_i}(t)$ is defined as 0 outside $a_i..d_i$). Bidders could instead be modeled as having a constant discount factor δ for allocations by using $\delta_{a_i,d_i}(t) = \delta^{t-a_i}$ for $t \in a_i..d_i$.

Let Θ^* denote the collection of possible sets of active players, i.e. sets of any cardinality containing only types: $\Theta^* = \{\emptyset\} \cup \Theta \cup (\Theta \times \Theta) \cup \ldots$. Let $\mathbb{Z}_{=}^{\Theta^*}$ denote the collection of functions that associate to any $\theta \in \Theta^*$ with n_θ types a vector of n_θ integers: $\mathbb{Z}_{=}^{\Theta^*} = \bigcup_{n \geq 0} \{\psi | \psi : \Theta^n \rightarrow \mathbb{Z}^n\}$.

Definition 20. *A* direct revelation mechanism *(henceforth,* dynamic auction*) for the dynamic allocation environment over periods* $1..T$ *consists of, at each time step* $t \in 1..T$

- *a history-dependent allocation update rule*

$$\Pi^t : \underbrace{\Theta^* \times \cdots \times \Theta^*}_{t-1} \times \underbrace{\mathbb{Z}_{=}^{\Theta^*} \times \cdots \times \mathbb{Z}_{=}^{\Theta^*}}_{t-1} \rightarrow \mathbb{Z}_{=}^{\Theta^*}$$

  I will use π^t *to denote* $\Pi^t(\cdot,\cdot)(\theta^t)$*. For* reported *type vectors* $\theta^{1..t-1}$ *and updates* $\pi^{1..t-1}$ *in the previous periods,* $\Pi^t(\theta^{1..t-1}, \pi^{1..t-1})(\cdot)$ *maps the vector of types* θ^t *to a vector of integers with a coordinate for each active type at* t*. I use shorthand* π_i^t *for the integer corresponding to* $\theta_i \in \theta^t$ *in* $\pi^t(\theta^{1..t-1}, \pi^{1..t-1})(\theta^t)$*.* π_i^t *represents the* increment *for* i*'s current holding of items:* $l_i^{t+1} = l_i^t + \pi_i^t$*. In particular, if* $\pi_i^t = 0$ *then* i*'s number of allocated items is unchanged and if* π_i^t *is negative then* $|\pi_i^t|$ *items are taken away from* i*. Thus the histories of allocations and updates can be obtained from one another.*

- *a supply update rule:* $S^{t+1} = R(S^t, \cup_{i \in \theta^t} \pi_i^t)$*.*

- *for each player* i*, a payment function* p_i^t*, specifying* i*'s monetary transfer* to *the seller.*

The mechanism's social choice function is composed of the update rule and the payment function. i's utility at time t (the payoff mapping of Def. 1) is given by $u_i^t = \delta_{a_i,d_i}(t)v_i(\theta^{1..t}, l_i^t, \pi_i^t) - p_i^t(\theta^{1..t}, \cup_{t'=1}^{t} \cup_{i' \in \theta^{t'}} \pi_{i'}^{t'})$*. Thus i's utility is uniformly affected by payment decisions in any period*[18]*, as opposed to updates to his number of allocated items which only matter to i in periods within his true arrival-departure interval. i's total utility is the sum of the utilities in each period:* $u_i = \sum_{t=1}^{T} u_i^t$.

Here is how this definition supports non-identical items. If at some time t, an active bidder i finds some item in S^t not interesting ($L_i \setminus S^t \neq \emptyset$) then the mechanism must be able to assign each item x to at most one active agent i_x who finds it interesting ($x \in L_{i_x}$) such that any agent i is assigned exactly l_i^t items. If some bidder i''s number of allocated items is updated ($\pi_{i'}^t \neq 0$), then the mechanism's internal assignment of items to bidders may change.

single-minded domain, persistent identical items The domains of Chapters 3 and 5 are for bidders with single-minded preferences and no item can be reassigned or re-used: a bidder can only be allocated once and no allocation can be undone. There is an initial supply of S^1 identical items; with a slight abuse of notation also let S^t denote the number of items available before time t. $S^{t+1} = R(S^t, \cup_{i \in \theta^t} \pi_i^t) = S^t - \sum_{i \in \theta^t} \pi_i^t$.

In Chapter 3, a bidder i with type $\theta_i = (a_i, d_i, q_i, r_i)$ has a private value of r_i for q_i or more items allocated *simultaneously* once and a value of 0 otherwise. Thus $v_i(\theta^{1..t}, l_i^t, \pi_i^t) = r_i$ if $l_i^t < q_i$ and $\pi_i^t \geq q_i$ (i obtains all his desired items at once in period t) and 0 otherwise.

In Chapter 5, there is one item and a bidder's value for it depends on other bidders'

[18] In this thesis, a bidder's payments will be zero outside his reported activity interval. In general it is conceivable that a bidder may only find out his payment after his departure. Arguments for payment decisions *immediately* following allocation decisions are proposed in [22].

private information. Thus π_i^t is 1 for at most one (i,t) pair and $l_i^t = 0, \forall\, i, \forall\, t$: the auction effectively ends as soon as the item is allocated. Each bidder i has a publicly known valuation function $v_i(\theta^{1..t}, 0, 1)$; recall that i's signal θ_i is, however, private information. To model a bidder's value dependence on *all* private information before his departure, I let $\delta_{a_i,d_i}(t) = 1$ for $t = d_i$ and 0 otherwise. Consequently, the seller will only allocate a bidder upon his departure:[19] if $\pi_i^t = 1$ then $t = \hat{d}_i$. Note that the seller's decisions in Chapter 5 should depend not only on previous decisions, but also on previous types, since the values of current bidders are determined by previous bidders' signals.

multi-unit domain, persistent identical items Consider a seller owning $S^1 = K$ identical items where each bidder i has a decreasing marginal value for each additional item won. No cancellations are allowed ($\pi_i^t \geq 0$), but a bidder may be allocated at more than one time period. Bidder i's private information is a vector $(r_i^1, \ldots, r_i^K)$ where r_i^h is i's value for the h-th item won and $r_i^h \geq r_i^{h+1}, \forall\, h = 1..K-1$. i has (marginal decreasing) value for the items received at t relative to i's holding: $v_i(\theta^{1..t}, l_i^t, \pi_i^t) = \sum_{x=l_i^t+1}^{l_i^t+\pi_i^t} r_i^x$. The supply update rule remains $S^{t+1} = R(S^t, \cup_{i\in\theta^t}\pi_i^t) = S^t - \sum_{i\in\theta^t}\pi_i^t$.

single-minded domain, persistent heterogeneous items This is the setting of Proposition 8 and Proposition 9, considered in [69] and it differs from Chapter 3 only in that items are no longer identical. The value of a type $\theta_i = (a_i, d_i, (r_i, L_i))$ is defined as $v_i(\theta^{1..t}, l_i^t, \pi_i^t) = r_i$ when $l_i^t < |L_i|$ and $\pi_i^t \geq |L_i|$, but 0 otherwise. This models a bidder having value only if allocated his *entire* interesting set.

cancellations for persistent heterogeneous items The setting of Chapter 4 can be described as follows. Impatient buyers arrive one by one, each buyer i placing an

[19] i can still report an earlier departure $\hat{d}_i < d_i$ and obtain value from the item at his true departure d_i.

instantaneous bid for a future (e.g. at time $T+1$) allocation of an item from L_i. The seller must immediately accept (by *temporarily* reserving an item for) or reject (deny a reservation for) each bidder. To accommodate future valuable offers, the seller is allowed to preempt (*bump*) an already accepted bidder. In particular, any item is available at any time (by bumping the bidder the item is currently assigned to). Bumped and rejected bidders do not get a second chance to bid or to be allocated. A bumped bidder incurs a utility loss amounting to an α fraction of his value, where $\alpha \in [0, 1)$.

Chapter 4's model can be cast in the framework of Definitions 19 and 20 as follows. The supply before any t is the same, a set $S^t = K$ of heterogeneous items: $R(K, \cdot) = K$.

A bidder i's private information θ_i consists of his value for an item and also, specific to Chapter 4, of i's decision whether to *invest* if accepted. The investment decision is also immediate, upon acceptance. No bidder reports his investment decision to the seller, but without investing, the bidder has no value for gaining or losing the item (Eq. (2.13) below formalizes i's value). I construe a_i as the time i bids and assume a bidder cannot manipulate it: $\hat{a}_i = a_i$. Any bidder i can only be allocated at arrival, because of impatience: $\pi_i^t = 1$ only possibly at $t = a_i$. No bidder departs until time T (thus $d_i = T$ is technically no longer private information, just like a_i). Physical allocations are only made at $T+1$, after the mechanism ends. Any bidder i allocated (at arrival) can be bumped, i.e. have the item removed from his holding: for $t \geq a_i + 1$, $-l_i^t \leq \pi_i^t \leq 0$. There can be at most one time $t \geq a_i + 1$ when $\pi_i^t = -1$; at all other $t' \geq a_i + 1$ i's number of allocated items is not updated ($\pi_i^{t'} = 0$).

Let us define a bidder i's value, effectively rewriting Eq. (4.1):

$$v_i(\theta^{1..t}, l_i^t, \pi_i^t) = \begin{cases} \xi(\mathrm{I})r_i, & \text{if } \pi_i^t = 1 \text{ (hence } l_i^t = 0) \\ 0, & \text{if } \pi_i^t = 0 \\ -(1+\alpha)\xi(\mathrm{I})r_i, & \text{if } \pi_i^t = -1 \text{ (hence } l_i^t = 1) \end{cases} \tag{2.13}$$

Bidder i's private information is comprised of his value r_i for an item and his investment decision $\mathrm{I} \in \{\mathrm{I}_0, \mathrm{I}_1\}$, where I_1 denotes investing and I_0 denotes not investing. Updates affect i only if i invested: $\xi(\mathrm{I}_0) = 0$ and $\xi(\mathrm{I}_1) = 1$. If i invested and is bumped, the effect of Eq. (2.13) for $\pi_i^t = -1$ is to cancel i's gain in value (obtained when allocated at his arrival) and additionally inflict an α factor loss.

expiring items I illustrate Definition 20's versatility by casting the expiring items domain into its framework; I will focus on unit-demand bidders for simplicity. The natural dynamic domain of expiring items is outside the scope of this thesis. The self-correction approach in Chapter 3 builds upon a similar approach [71] used in an expiring items domain.

There is exactly one item to be allocated each period t: $S^{t+1} = \{g^{t+1}\} = R(\{g^t\}, \cdot)$; $\pi_i^t = 1$ for at most one player i and 0 for the others. Bidder i has value r_i for exactly one item in some period $t \in [a_i, d_i]$: $v_i(\theta^{1..t}, 0, 1) = r_i$ and $v_i(\theta^{1..t}, l, 1) = 0$, $\forall\, l \geq 1$.

Definition 20 can be modified to allow bidders to express combinatorial valuations on bundles of heterogeneous items; the simpler version avoids notation clutter.

The following two results establish the (almost) equivalence of monotonicity and incentive compatibility in a dynamic allocation environment with simple preferences. A unit-demand type $\theta_i = (a_i, d_i, r_i)$ *dominates* another unit-demand type $\theta_i = (a_i', d_i', r_i')$ if θ_i' bids less ($r_i' \leq r_i$) and has a tighter arrival-departure interval: $a_i \leq a_i' \leq d_i' \leq d_i$. Monotonicity essentially requires that, if a bidder i, that is allocated for type θ_i', were to report instead a type θ_i that dominates θ'_i, i would still be allocated. Monotonicity is analogous to the WMON property in Section 2.2.

Proposition 8. *[69] A monotonic, deterministic policy is implementable in a domain with (known interesting bundle) single-valued preferences.*

The following is a converse of Proposition 8.

Proposition 9. *[69] Any deterministic policy implementable without paying unallocated agents in a domain with (known interesting bundle) single-valued preferences and "reasonable" misreporting must be monotonic.*

Reasonable misreporting requires that an agent be able to report any value and any arrival-departure interval within his true arrival-departure interval. Recall that other types of temporal misreports are ruled out in Definition 19.

I now present two prominent distribution-based mechanisms; relevant distribution-free dynamic mechanisms will be reviewed in Chapter 4.

2.5.2 Dynamic distribution-based mechanisms

In this section I assume that, at any time t,

A1 there exists a distribution (which is common knowledge) on future arrivals at any time $t' \in [t+1, T]$ and this distribution is independent of past arrivals and decisions.

A2 any bidder truthfully reports his arrival: $\hat{a}_i = a_i$, $\forall\, \theta_i$.

While results in this section hold under relaxations of these assumptions, these assumptions simplify notation and exposition significantly.

Assumption A1 typically holds if buyers' bids are not influenced by previous prices or demand.

Assumption A2 allows defining a dynamic strategy for types without notation for bidders that are not active only because of delaying their arrival.

Player i's expected value for a policy of allocation updates $\Pi^{t..T}$ starting from t is

$$V_i(\theta^t, \Pi^{t..T}) = \mathbb{E}_{\theta^{>t}} \left[\sum_{\tau=t}^{T} \delta_{a_i,d_i}(\tau) \cdot v_i(\theta^{1..\tau}, l_i^\tau, \pi_i^\tau) \right] \tag{2.14}$$

where π_i^τ is the coordinate corresponding to θ_i in $\Pi^\tau(s(\theta^{1..\tau-1}), \pi^{1..\tau-1})(s(\theta^\tau))$ when players (including i) use reporting strategies s.

The within-period ex post Nash equilibrium concept is a hybrid generalization of the Bayes-Nash and ex post equilibrium concepts to dynamic settings. It requires that the equilibrium strategies are a best response for any current types and in expectation over future types. That is, provided that others play their equilibrium strategies, a player would not regret playing his equilibrium strategy even if it had access to all the information currently in the market, i.e. the types of active players.

Definition 21. *[4] A dynamic mechanism $M = (\Pi^{1..T}, p^{1..T})$ is* within-period ex post Nash incentive-compatible *if, for any player i, reporting his true type maximizes his utility in M given that others report their true types, for any current types θ_{-i}^t and* in expectation *over future players.*

$$\begin{aligned} \forall\, \theta_{-i}^t,\ & V_i(\theta_i, \theta_{-i}^t, \Pi^{t..T}(\theta_i, \theta_{-i}^t)) - p_i(\theta_i, \theta_{-i}^t, \Pi^{t..T}(\theta_i, \theta_{-i}^t)) \\ \geq & V_i(\theta_i', \theta_{-i}^t, \Pi^{t..T}(\theta_i', \theta_{-i}^t)) - p_i(\theta_i', \theta_{-i}^t, \Pi^{t..T}(\theta_i', \theta_{-i}^t)) \,\forall\, \theta_i' \in \Theta \end{aligned}$$

where $p_i(\theta^t, \Pi^{t..T})$ is i's expected total payment starting from period t, defined similarly to $V_i(\theta^t, \Pi^{t..T})$ in Eq. (2.14).

The efficient (in expectation) policy Π^* is the one that maximizes the expected social welfare of current and future players.

$$\Pi^*(\theta^t) = \underset{\text{policy } \Pi^{t..T}}{\text{argmax}} \; V^t(\Pi^{t..T}(\theta^t)) \text{ where}$$

$$V^t(\Pi^{t..T}(\theta^t)) = \left(\sum_{i \in \theta^t} V_i(\theta^t, \Pi^{t..T}) \right) + \mathbb{E}_{\theta^{>t}} \left[\sum_{\tau=t+1}^{T} \sum_{i' : \hat{a}_{i'} = \tau} V_{i'}(\theta^\tau, \Pi^{\tau..T}) \right]$$

I first review a generalization of the VCG mechanism to a dynamic allocation environment with distributional information about the future.

Definition 22. *In any period t, the* online-VCG *mechanism [72] adopts the efficient policy $\Pi^*(\theta^t)$, charging an active player i*

$$p_i^t(\theta^t) = \begin{cases} V_i(\theta^t, \Pi^*(\theta^t)) - \left(V^t(\Pi^*(\theta^t)) - V^t_{-i}(\Pi^*(\theta^t_{-i})) \right) & \text{if } i \text{ arrived at } t \\ V_i(\theta^t, \Pi^*(\theta^t)) & \text{otherwise} \end{cases} \tag{2.15}$$

Proposition 10 is an immediate corollary of results in [17, 69].

Proposition 10. *The online-VCG mechanism is within-period ex-post Nash incentive-compatible under assumptions A1 and A2.*

The online-VCG mechanism also satisfies within-period ex-post Nash individual rationality, i.e. for any set of active players, an active player has non-negative expected utility in online-VCG by reporting his true type. Thus, in the online-VCG dynamic generalization, a bidder's utility is maximized and non-negative by reporting his true type *if* the other (active) players report their true types as well and only in expectation over future players. In contrast, in the static VCG mechanism, a bidder's utility is maximized and non-negative by reporting his true type *regardless* of other players' reports.

The within-period ex post properties of online-VCG are preserved if one replaces assumption A2 with

A2') no bidder can improve the expected value of the policy Π^* by delaying his arrival.

The second mechanism introduced in this section is a generalization of the AGV mechanism (recall Section 2.4.1) to a dynamic allocation environment, reducing to it in a static environment, i.e. when $T = 1$ and bidders are all present at once and do not have arrivals or departures.

Definition 23. *In any period t, the* dynamic balanced mechanism *[4] adopts the efficient policy $\Pi^*(\theta^t)$, charging an active player i*

$$p_i^t(\theta^t) = \left(\tfrac{1}{|\theta^t|-1}\textstyle\sum_{\theta_j \in \theta^t_{-i}} \Delta_j(\theta_j, \theta^{t-1})\right) - \Delta_i(\theta_i, \theta^{t-1}) \text{ where} \tag{2.16}$$

$$\Delta_i(\theta_i, \theta^{t-1}) = \begin{cases} E(\theta_i, \theta^{t-1}) - \mathbb{E}_{\tilde{\theta}_i}[E(\tilde{\theta}_i, \theta^{t-1})] & \text{if } i \text{ arrived at } t \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \tag{2.17}$$

$$E(\theta_i', \theta^{t-1}) = \mathbb{E}_{\tilde{\theta}^{a=t}_{-i}}\left[V^t_{-i}\left(\Pi^*\left(\theta_i', \theta^{t-1}_{-i}|_{d\geq t}, \tilde{\theta}^{a=t}_{-i}\right)\right)\right] \tag{2.18}$$

Π^*'s argument in Eq. (2.18) ensures that the types at time t are correctly updated from the ones at time $t-1$: types active at $t-1$ that do not depart at $t-1$ (denoted by $\theta^{t-1}_{-i}|_{d\geq t}$) and types arriving at t (denoted by $\tilde{\theta}^{a=t}_{-i}$).

Proposition 11. *[4] The dynamic balanced mechanism is strongly budget-balanced and has truthful reporting as Bayes-Nash equilibrium.*

Apart from strong budget-balance, the dynamic balanced mechanism also preserves the ex ante individual rationality property of the static AGV mechanism.

For concreteness, I consider the online-VCG and dynamic balanced mechanisms in a very simple dynamic allocation environment.

Example 2. *Suppose that there is one item and two periods. Agents 1 and 2 arrive in* *periods I and respectively II, both departing after period II. The efficient policy* Π^* *allocates the item in period II to the highest bidder; say that a tie is broken in favor of agent* 1.

In online-VCG, payments are

$$p_1(\theta^I) = 0 - \left(\int_0^\infty \max(\theta_1, r_2) f(r_2) dr_2 - \int_0^\infty r_2 f(r_2) dr_2 \right) = - \int_0^{\theta_1} (\theta_1 - r_2) f(r_2) dr_2$$

$$p_1(\theta^{II}) = \mathbf{1}_{\theta_1 \geq \theta_2} \cdot \theta_1 \text{ and } p_2(\theta^{II}) = \mathbf{1}_{\theta_1 < \theta_2} \cdot \theta_2 - (\max(\theta_1, \theta_2) - \theta_1) = \begin{cases} 0, & \text{if } \theta_1 \geq \theta_2 \\ \theta_1, & \text{if } \theta_1 < \theta_2 \end{cases}$$

Agent 2's online-VCG payment is the same as in a second-price auction, i.e. the static VCG mechanism for one item.

In the dynamic balanced mechanism, payments are determined by

$$\Delta_1(\theta_1, \emptyset) = \int_{\theta_1}^\infty r_2 f(r_2) dr_2 - \int_0^\infty \int_{r_1}^\infty r_2 f(r_2) dr_2 f(r_1) dr_1$$

$$\Delta_2(\theta_2, \theta^I) = \mathbf{1}_{\theta_1 \geq \theta_2} \cdot \theta_1 - \int_0^{\theta_1} \theta_1 f(r_2) dr_2$$

The payments in the dynamic balanced mechanism are quite different from the ones in the static AGV mechanism in Section 2.4.1.

Summary

I have introduced standard game theory and static mechanism design concepts.

I have reviewed characterization of truthfulness in static mechanisms and I have highlighted the limited expressiveness of order-based domains, a key assumption in a state-of-the-art characterization of truthfulness for combinatorial auctions.

I have defined a dynamic allocation environment and I have shown how this definition can accommodate, among others, the environments in this thesis.

I have reviewed two prominent examples of dynamic distribution-based mechanisms.

Chapter 3

Self-Correction

Abstract. I use sample-based stochastic optimization methods for the purpose of strategyproof dynamic, multi-unit auctions. There are no analytic characterizations of optimal policies for this domain and thus a heuristic approach, such as that proposed here, seems necessary in practice. Following the framework of Parkes and Duong [71], I perform sensitivity analysis on the allocation decisions of an online algorithm for stochastic optimization, and correct the decisions to enable a strategyproof auction. In applying this approach to the allocation of non-expiring goods, the technical problem that I must address is related to achieving strategyproofness for reports of departure. This cannot be achieved through self-correction without canceling many allocation decisions, and must instead be achieved by first modifying the underlying algorithm. I introduce the `NowWait` method for this purpose, prove its successful interfacing with sensitivity analysis and demonstrate good empirical performance. My method is quite general, requiring a technical property of *uncertainty independence*, and that values are not too positively correlated with agent patience. I also show how to incorporate "virtual valuations" in order to increase the seller's revenue.

3.1 Introduction

Mechanism design addresses the problem of private information in economic environments and seeks to implement desirable outcomes despite the willingness of agents to misreport this information. Auctions present a canonical problem of mechanism design. Many important mechanism design problems are in fact dynamic, for example with a dynamic agent population and new bids arriving online [55, 69]. Consider selling theater tickets, airline seats, or banner advertisements, where bids may be expected to arrive over time and associated with bidders that require a response before all bids have been received.

I consider a very natural instance of this problem. There are S^1 units of an identical item for sale, to be sold in the course of T time periods. Each bidder (or agent) i has an arrival time $a_i \in 1..T$, departure time $d_i \in 1..T$, and value $r_i \in \mathbb{R}_{\geq 0}$ for $q_i \in \mathbb{Z}_{>0}$ units of the item. The semantics of the bidder's type, $\theta_i = (a_i, d_i, r_i, q_i)$, are that the bidder has value r_i for receiving q_i units in some period $t \in \{a_i, \ldots, d_i\}$. The arrival time models the moment at which the agent realizes his demand or learns about the existence of the auction while the departure models the latest moment at which the agent still has value for receiving the items. The agent types are identically and independently distributed with probability density function $f(\theta_i)$, and n^t agents arrive each period, with associated density function $g(t)$.

The design goal may be alternatively one of efficiency (i.e., maximizing expected value) or revenue (i.e., maximizing expected payments). I will consider both objectives in this chapter. If the goal was efficiency, then one might consider adopting the online Vickrey-Clarke-Groves mechanism [72, 73]. But this mechanism requires an optimal (or ϵ-optimal) decision policy, which is not computationally feasible in this domain. This dynamic allocation problem is inherently combinatorial because of the multi-unit demand of agents, and optimal decision policies cannot be computed (or even represented) offline because of the size of the state space, which needs to include all possible sets of undominated agent bids that can be received in a single period. Online, sample-based algorithms to compute ϵ-optimal policies [50], also quickly become intractable because the sample-tree scales exponentially in the look-ahead horizon.

Other prior work in the probabilistic, dynamic framework considers only domains that facilitate analytic characterizations of optimal policies; e.g., domains with commonly-ranked items [36], unit-demand bidders with "regular" valuation distributions [68], smoothly changing types [74] or unit-demand bidders with time-discounting [33]. I am not aware of any prior work that is able to scalably address the multi-unit, dynamic auction problem in this chapter.

In the absence of computational methods to compute optimal policies, or analytic characterizations of optimal polices, it seems necessary to adopt a heuristic approach. Parkes and Duong [71] propose "output-ironing" as a methodology to transform heuristic, online algorithms for stochastic optimization into *strategyproof* dynamic auction protocols. A strategyproof dynamic auction is one in which truthful, immediate bidding is a dominant-strategy equilibrium. It is known that strategyproof dynamic auctions in this environment require that the allocation policy is *monotonic* [43, 69]. Loosely, a monotonic policy is one in which an agent continues to be allocated as he improves his bid (appropriately defined for arrival, departure, value and quantity). The idea behind output-ironing is to verify the monotonicity of an allocation policy online, as decisions are made, and correct such decisions as necessary to make the policy monotonic.

Parkes and Duong [71] successfully apply output-ironing to environments with *expiring goods*, where one or more units must be allocated in each period; e.g., the allocation of compute time on a shared computer. But this self-correcting approach is difficult to apply in environments with non-expiring goods, such as the one considered here. The problem is with regard to establishing departure monotonicity, which requires that an agent allocated for some bid (or reported type) is also allocated for a bid with the same arrival, value, and quantity but a later departure. For example, suppose an agent i with reported type $(1, 5, \$10, 2)$ is allocated in period 5. It must be verified that the decision policy will continue to allocate i (in some period) for all reports of type $(1, d', \$10, 2)$ with $d' > 5$. But with non-expiring goods, any reasonable policy will wait until i's departure to decide whether to allocate i so that it maximizes the available information about other bids. This in turn makes it impossible to verify departure monotonicity with respect to later reports of departure because the bids that will arrive and affect allocation decisions are as yet unknown (e.g., at time period 5 it is unknown whether or not bids with later reported departures will be allocated). Thus, output ironing would need to cancel all allocations except those to bids for which there are no possible types with later departures (i.e., maximally-patient agents). This would result in an implemented policy with very low efficiency.

My contribution. I design `NowWait`, a heuristic modification of the CONSENSUS algorithm [84] for online stochastic optimization, that is provably departure monotonic and thus precludes the need for output ironing with respect to departure. When coupled with output ironing in the other dimensions of a bidder's type, it provides a strategyproof and scalable dynamic multi-unit auction. I also establish that a simplified form of output ironing, referred to here as *adjacency ironing* is sufficient to establish monotonicity. This significantly improves the scalability of the methodology. `NowWait` balances the immediate reward from accepting a bid with the estimated opportunity cost from waiting to a future period. Empirical analysis demonstrate higher than 90% efficiency with respect to the offline optimum, and higher than 98% efficiency with respect to the online optimum when this benchmark is available. This is achieved with around a 20x slow down due to using computational ironing, over-and-above the underlying method of stochastic optimization, and a per-period solve time of approximately 40 seconds in instances with more than 100 arrivals and 100 goods to allocate. The approach is very flexible, and can be applied to inputs that are first transformed to "virtual valuations" as in Myerson's revenue-optimal offline auction [64]. Experimental results show that this can boost revenue significantly in environments with low demand relative to supply, as would be expected. The approach can also be combined with learning of the underlying distribution on agent types, because the incentive properties do not rely on having a correct probabilistic model.[1]

Other related work.

Boutilier et al. [14] apply online stochastic combinatorial optimization to clearing expressive banner ad auctions, but without consideration of incentive issues. Hajiaghayi et al. [45] adopt a dynamic-programming approach to design simple dynamic auctions for settings with unit-demand bidders, with values drawn from a known prior but whose number may not be known in advance. Parkes [69] provides a general survey of online mechanisms, including both probabilistic and prior-free approaches.

The agenda of *automated mechanism design* (AMD) [24] shares the goal of creating a mechanism automatically, but differs from the approach adopted here in that it adopts optimization to design a functional description of all decisions that will be made by a mechanism, rather than seeking to adapt an existing decision algorithm, such as in the approach adopted here. This makes AMD very difficult to scale.

The work presented here conforms to the agenda of *heuristic mechanism design*, recently advocated by Parkes [70]. This stipulates that a problem in computational mech-

[1] The incentive problems that can occur because of informational externalities when learning in the context of dynamic auctions (see Gershkov and Moldovanu [35]) can be avoided by precluding the use of a new report of a type θ by bidder i for the purpose of updating the center's model about future reports until the agent has departed.

anism design be considered solved when a state-of-the-art algorithm for solving a problem with cooperative agents can be adopted, "with small modification" to solve the problem with self-interested agents. Output ironing is a small modification to Consensus in this sense, because it retains the vast majority of the decisions and also the same underlying computational approach in making decisions.

Incremental mechanism design [25] also modifies the rules of a mechanism to remove opportunities for manipulation. Unlike this work, it requires an iterative approach because fixing one opportunity may introduce another, and is described only for offline mechanisms where the complete type profile is known. Lavi and Swamy [57] provide a general procedure to transform approximate VCG mechanisms into truthful-in-expectation mechanisms for static environments. But it is not apparent how to apply the approach to dynamic policies, where only part of an allocation is available at any point.

Outline. I first define the set-up more formally, and define the Consensus algorithm for online stochastic combinatorial optimization. In Section 3.2, I also define output ironing and establish a result about the sufficiency of adjacency-ironing. I continue by defining the important property of *departure obliviousness*, and noting the flexibility of my approach in reference to virtual valuations and learning. Sections 3.3 and 3.4 define variants on Consensus that are adapted to my problem and explain how to perform sensitivity analysis. In Section 3.5 I present experimental results, before concluding. All proofs are deferred to the appendix.

3.2 Consensus and ironing

Recall that in my model, the type of an agent (a_i, d_i, r_i, q_i) specifies a value r_i for an allocation of q_i units in some period $\{a_i, \ldots, d_i\}$. I adopt the standard quasi-linear utility model, in which an agent that pays p for q_i units has utility $r_i - p$. I refer to $d_i - a_i$ as an agent's *patience*, and assume this is bounded with $d_i - a_i \leq \overline{\Delta}$. In the auctions I consider, an agent can make to the auctioneer a single claim about his type. All misreports of type are possible except for reports of *early arrivals*, i.e., I assume that an agent cannot claim to have an earlier arrival than his true arrival. To motivate this, recall that the arrival period models the period at which an agent learns of his own demand or learns about the existence of the mechanism. It is reasonable to restrict an agent's strategy from bidding before this period.[2]

The efficient policy maximizes total expected value:

$$V^*(h^1) = \max_{k^1 \in \mathcal{K}(h^1)} \mathbb{E}_{\theta^{2..T}} \left[\max_{k^2 \in \mathcal{K}(h^2)} \mathbb{E}_{\theta^{3..T}} \left[\ldots \max_{k^T \in \mathcal{K}(h^T)} v(k, \theta) \right] \right], \tag{3.1}$$

where k^t is the allocation decision taken at t, state $h^t = (S^t, A^t)$ denotes the number of available items S^t and the current set of *active* agents A^t (agents with $t \in \{a_i, \ldots d_i\}$), $\mathcal{K}$ defines the set of feasible allocation decisions in period t, θ^t is the set of types that arrive at t, and $v(k, \theta)$ is the total value to agents allocated by decisions $k = k^{1..T}$ given types $\theta = \theta^{1..T}$. I write $i \sqsubset k$ for "agent i is allocated by decision k".

A dynamic auction $M = (\pi, x)$ defines a decision policy $\pi = \{\pi^{1..T}\}$ and a payment policy $x = \{x^{1..T}\}$. The decision and payment policy may be randomized and depend on random events $\omega = \{\omega^{1..T}\}$, for example random samples of future bids. With this, the decision policy π induces decisions $k^t := \pi^t(S^t, A^t, \omega^{1..t})$ and collects payment $x_i^t(S^t, A^t, \omega^{1..t}) \in \mathbb{R}_{\geq 0}$ from each active agent. As useful shorthand, let $\pi_i(\theta, \omega) = 1$ denote that agent i is allocated for reported types θ given uncertain events ω, with $\pi_i(\theta, \omega) = 0$ otherwise.

Define the *critical-value* for agent i given policy π and reports θ_{-i} of other agents as $v^c_{(a_i,d_i,q_i)}(\theta_{-i}, \omega) = \min\{r'_i \text{ s.t. } \pi_i((a_i, d_i, r'_i, q_i), \theta_{-i}, \omega) = 1\}$, or ∞ if no such r'_i exists.

[2] This assumption is adopted in many other papers, including Hajiaghayi et al. [43] and Pai and Vohra [68].

This is the smallest bid value for which agent i would still win, all else unchanged.Define a partial order on types: $\theta_i \preceq_\theta \theta'_i \equiv (a_i \geq a'_i) \wedge (d_i \leq d'_i) \wedge (r_i \leq r'_i) \wedge (q_i \geq q'_i)$. Type θ'_i is higher than θ_i if it offers the seller more flexibility, i.e. it has higher reward, demands less units and has a larger availability interval.

Monotonicity requires that an allocated agent would still be allocated if his type were higher, all else unchanged:

Definition 24. *Policy π is monotonic if $(\pi_i(\theta_i, \theta_{-i}, \omega) = 1) \wedge (r_i > v^c_{(a_i,d_i,q_i)}(\theta_{-i}, \omega)) \Rightarrow \pi_i(\theta'_i, \theta_{-i}, \omega) = 1$ for all $\theta'_i \succeq_\theta \theta_i$, for all θ_{-i}, ω, and all agents i.*

Monotonicity is sufficient, and essentially necessary (if losing agents receive no payment) for strategyproofness in this environment [69]. A mechanism with a monotonic decision policy is made strategyproof by defining a payment policy that charges each allocated agent his critical value. The critical value can be computed upon the departure.

Surprisingly, optimal policies need not be monotonic:

Example 3. *[71] There are 3 units to allocate and 2 periods. In period 1, agent 1 has type $(1, 1, \$5, 1)$ and agent 2 has type $(1, 2, \$500, 2)$. In period 2, with probability $1 - \gamma$ an agent will arrive with type $(2, 2, \$1000, 3)$ and with probability γ an agent will arrive with type $(2, 2, \$5000, 1)$, for some small $\gamma > 0$. The optimal policy must make a decision in period 1 about agent 1 because agent 1 will depart in this period. Agent 1 is not allocated because this would preclude the ability to allocate to type $(2, 2, \$1000, 3)$ that will arrive with high probability in period 2. Then in period 2, suppose the unlikely event occurs and an agent with type $(2, 2, \$5000, 1)$ arrives and the optimal policy allocates to agent 2 and also this new arrival.*

Now consider what happens were agent 2 to bid \$1000 rather than \$500. The optimal decision in period 1 is now to allocate to agent 1 for \$5 and expect to allocate to *agent 2 in period 2. It is better to get certain value of \$5 from agent 1 than expected value* $\gamma 5000$ *from the unlikely type in period 2 (for a small enough* γ*). But now the same unlikely event occurs, and an agent with type* $(2, 2, \$5000, 1)$ *arrives in period 2. The optimal policy allocates to this agent and with 1 unit left is now unable to allocate to agent 2. This is a failure of monotonicity: agent 2 increases his value but went from winning to losing, fixing the types of other agents.*

It is easy to see that this same failure of monotonicity will occur with the policies constructed using sample-based stochastic optimization algorithms such as Consensus [84], described next. It is this failure of monotonicity that sets up the problem addressed in this paper.

3.2.1 The Consensus algorithm

The Consensus (**C**) algorithm, proposed by van Hentenryck and Bent [84] for online stochastic optimization, is illustrated in Figure 1, together with the additional step of output ironing in determining decision $\breve{k}^t$ in period t.

Algorithm 1 Consensus algorithm with ironing at time t.

votes(k):=0 for each allocation k of up to S^t items to A^t
σ^j:=GetSample(t) for each $j = 1..|\Sigma|$; $\Sigma = \{\sigma^{1..|\Sigma|}\}$
for each $j = 1..|\Sigma|$ **do**
  $\alpha^j := \mathrm{Opt}(S^t, A^t, \sigma^j) \cap A^t$ `// active agents only`
  $\alpha^j_{\mathtt{s}}$:=`Select`$(\alpha^j, \Sigma, S^t, A^t)$
  votes($\alpha^j_{\mathtt{s}}$):=votes($\alpha^j_{\mathtt{s}}$)+1
end for
$k^t := \arg\max_k$ votes(k)
$\breve{k}^t := \{i \sqsubset k^t : \text{ not isIroned}_{A,D,Q}(\theta_i, t, (S, A)_{a_i..t}, \Sigma)\}$
return $\breve{k}^t$

A *scenario* σ^j in period t is a sample of a possible future: σ^j defines the types $\theta^{t+1..T}$ for periods $t+1$ through T. Given a scenario σ^j, and the current state (S^t, A^t), there is a well-defined offline optimization problem $\mathrm{Opt}(S^t, A^t, \sigma^j)$. This is a weighted knapsack problem: find the subset of bids $A^t \cup \sigma^j$ that maximize the total value allocated without exceeding the capacity S^t. The **C** algorithm constructs samples, solves this offline optimization problem for each sample, and of the agents allocated in this offline problem picks out as winning agents only the active agents (i.e., discarding future, sampled, agents). Denote this set of winning agents in scenario σ^j as α^j. This set may then be additionally "filtered" via a `Select` function to give set $\alpha^j_{\mathtt{S}}$. The set of active agents $\alpha^j_{\mathtt{S}}$ then receives one vote, the one for scenario σ^j. It is the `Select` function that will be modified to make **C** departure monotonic. $\mathbf{C} \oplus$ `Select` specifies that **C** is used together with the `Select` function.

The proposed decision k^t for current time t is the one with most votes (breaking ties at random). The final decision $\breve{k}^t$ results from discarding all ironed agents.

The **C** algorithm is applicable in domains satisfying *uncertainty independence*, i.e. in which the distribution of future agents is independent of past and current decisions:[3]

$$\mathbb{P}(\theta^{t+1..T}|k^{1..t}) = \mathbb{P}(\theta^{t+1..T}) \tag{3.2}$$

for all t, all $k^{1..t}$. This property requires that future demand is exogenous, and independent of current and past allocation decisions. For example, when selling airline tickets, it requires that bids for seats arrive irrespective of the number of seats remaining for sale. Uncertainty independence ensures that scenarios are valid for any decision path and allows for the same $|\Sigma|$ scenarios to be valid whatever the decision made now and in future periods.

[3]All my results remain valid if the uncertainty independence requirement (3.2), is also conditioned on past and current arrivals: $\mathbb{P}(\theta^{t+1..T}|k^{1..t}, \theta^{1..t}) = \mathbb{P}(\theta^{t+1..T}|\theta^{1..t})$. However, at each time step, new scenarios would have to be generated. The results further extend if the supply is stochastic, satisfying a similar uncertainty independence property.

3.2.2 Output ironing

Output ironing proceeds as follows (the "isIroned" function). Policies π and $\breve{\pi}$, and decisions k^t and $\breve{k}^t$, denote respectively the policies and decisions before and after ironing. Let $t_i^\pi(\theta, \omega) \in T \cup \{\infty\}$ denote i's allocation time (∞ if none exists) when reported types are θ, for random events ω.

Definition 25 (ironing)**.** *Given decision k^t, the* ironed *decision $\breve{k}^t$ only keeps those $i \sqsubset k^t$ for which*

$$t_i^\pi(\theta_i'', \theta_{-i}, \omega) \leq t_i^\pi(\theta_i', \theta_{-i}, \omega), \tag{3.3}$$

for all $\theta_i'' \succeq_\theta \theta_i' \succeq_\theta \theta_i$. If (3.3) *fails, i's allocation is canceled.*

The ironing step is performed in a period t in which $\mathbf{C}$ proposes to allocate an agent. Eq. (3.3) requires that an allocation to agent i is canceled unless an allocation to the same agent would have also been made, and in a monotonically-earlier period, for all higher reported types of agent i. When an allocation is canceled the items that were to be allocated are discarded and agent i is never allocated.[4]

Ironing requires not only that an agent is provably allocated for all higher reports, but that this occurs in *monotonically earlier* periods. With this, it is never the case that a type survives ironing, while a higher type would not and one can make an inductive argument to establish that

Theorem 8. *[71] Ironed policy $\breve{\pi}$ is monotonic.*

The uncertainty-independence property facilitates ironing, because it enables the

[4]The allocated goods are discarded, rather than returned, when the decision is canceled in order to prevent a knock-on effect, wherein a decision to iron the decision of one agent would have a ripple effect on the decision of the "base policy" π and thus whether another agent should be ironed.

simulation of counterfactual states as the type of an agent is varied.[5]

In fact, it is sufficient to perform a simplified form of ironing. Let $\theta_i'' \in \theta_i'{+}{+}$ if θ_i'' is a higher type than θ_i' but θ_i'' strictly improves over θ_i' in at most one dimension of (a, d, r, q). I henceforth fix θ_{-i} and ω, and omit them from the t_i^{π} notation.

Definition 26. *Given decision k^t,* adjacency-ironing *only keeps those $i \sqsubset k^t$ for which, for all $\theta_i' = (a_i', d_i', r_i', q_i') \succeq_\theta \theta_i = (a_i, d_i, r_i, q_i)$, with $r_i' = r_i$, it holds that*

$$t_i^{\pi}(\theta_i'') \leq t_i^{\pi}(\theta_i'),\ \forall\, \theta_i'' \in \theta_i'{+}{+} \text{ with } r_i'' = r_i' \text{ and} \tag{3.4}$$
$$t_i^{\pi}(\langle a_i', d_i', r_i''', q_i'\rangle) \leq t_i^{\pi}(\langle a_i', d_i', r_i'', q_i'\rangle), \forall r_i''' \geq r_i'' \geq r_i'$$

If (3.4) *fails, i's allocation is canceled.*

Theorem 9. *Adjacency-ironing is equivalent to ironing.*

Proof. Suppose $t_i^{\pi}(\theta_i) = t$. Both $\overline{\pi}$ and $\breve{\pi}$ can only cancel decisions of the same base policy π: $t_i^{\overline{\pi}}(\theta_i), t_i^{\breve{\pi}}(\theta_i) \in \{\infty, t\}$.

Suppose $t_i^{\breve{\pi}}(\theta_i) = t$. If $t_i^{\overline{\pi}}(\theta_i) = \infty$ then by Eq. (3.4), $\exists \theta', \theta_i'' \in \theta_i'{+}{+}$ with $t_i^{\pi}(\theta_i'') > t_i^{\pi}(\theta_i')$. But then $t_i^{\breve{\pi}}(\theta_i) = \infty$ by definition (Eq. (3.3)), contradiction. Thus $t_i^{\overline{\pi}}(\theta_i) = t$.

Suppose now that $t_i^{\breve{\pi}}(\theta_i) = \infty$. By Eq. (3.3) there exists $\theta_i'' \succeq_\theta \theta_i' \succeq_\theta \theta_i$ such that $t_i^{\pi}(\theta_i'') > t_i^{\pi}(\theta_i')$. But then there must exist $\tilde{\theta}_i'$ and $\tilde{\theta}_i'' \in \tilde{\theta}_i'{+}{+}$ with $\theta_i'' \succeq_\theta \tilde{\theta}_i'' \succeq_\theta \tilde{\theta}_i' \succeq_\theta \theta_i' \succeq_\theta \theta_i$ such that $t_i^{\pi}(\tilde{\theta}_i'') > t_i^{\pi}(\tilde{\theta}_i')$: on the lattice of types, $\tilde{\theta}_i''$ to $\tilde{\theta}_i'$ is just a step of the walk from θ_i'' to θ_i'. Then $t_i^{\overline{\pi}}(\theta_i) = \infty$ as well, by Eq. (3.4): from violating the first condition if $\tilde{r}_i'' > \tilde{r}_i'$ and the second one if $\tilde{r}_i'' = \tilde{r}_i'$. □

Algorithm 2 performs adjacency-ironing, following the prescription of this definition. Algorithm 3 provides pseudo-code for isIroned$_R$, which checks the first condition in

[5]If the realization of new bids was dependent on policy decisions, then the effect of some earlier change in decision could not be simulated because the future after that change in decision would not be known.

Algorithm 2 isIroned$_{A,D,Q}(\theta_i, t_i^*(\theta_i), (S^t, A^t)_{a_i \le t \le t_i^*}, \Sigma)$: online verification of monotonicity with respect to all higher types θ_i' for an agent of type θ_i allocated in period t_i^*. New allocation time $t_i^*(\theta_i')$ is found by simulating π starting at a_i, where i's type is replaced by θ_i'. The breakpoints computed in isIroned$_R$ for each θ_i' allow determining efficiently whether $t_i^\pi(\theta_i'') > t_i^\pi(\theta_i')$ for any $\theta_i'' \in \theta_i'$++.

for each $\theta_i' = (a_i', d_i', r_i, q_i')$, $(a_i' \le a_i, d_i' \ge d_i, q_i' \le q_i)$ **do**
  if isIroned$_R(\theta_i', t_i^*(\theta_i'), (S^t, A^t)_{a_i \le t \le t_i^*(\theta_i')}, \Sigma)$ **or**
   for any $\theta_i'' \in \theta_i'$++ with $r_i'' = r_i'$, $t_i^\pi(\theta_i'') > t_i^\pi(\theta_i')$ **then**
    return `true //` i `ironed`
  end if
end for
return `false //` i `not ironed`

Eq. (3.4). It tracks the changes in **C** decisions for values higher than i's reported value, r_i, fixing the rest of i's type and all other agent types. For each scenario in each time period in $\{a_i, \ldots, t_i^*\}$ it identifies values at which the set of agents selected to be allocated in the offline allocation in that scenario would change: these are the *scenario breakpoints.* It does so via the BrkPts function, which determines the set of all (time, scenario, value) triples at which the set of agents selected to be allocated changes. This function also determines the "before" and "after" decision as the value is increased past the scenario breakpoint, denoted respectively $\alpha_{\mathsf{s}}^{<}(\beta)$ and $\alpha_{\mathsf{s}}^{>}(\beta)$ for breakpoint β. An example of Algorithm 3 for **C** when `Select` is the `IgnoDep` function is presented in Example 4.

Example 4. *Say 3 items are for sale for* 2 *time periods. Agents* $X_{1,2}$ *arrive at time* 1*:* X_i *has* $a_i = 1, d_i = i, q_i = i, r_i = i, i = 1, 2$*. There are 7 time* 2 *scenarios* $\sigma^{1..7}$*, each with one agent with quantity* $2, 2, 3$ *and value* $3, 4, 10$ *on scenarios* $\sigma^{1,2}$*,* $\sigma^{3,4}$ *and* $\sigma^{5,6,7}$ *respectively. Votes are* $\{X_1\}$ *and* $\emptyset$ *for* $\sigma^{1..4}$ *and* $\sigma^{5..7}$ *respectively; hence* X_1 *is allocated. Agent* X_3 *arrives at time* $2 : a_3 = d_3 = 2, q_3 = 1, r_3 = 0.5$*. As a result,* X_2 *is allocated at time* 2*:* $t_2^* = 2$*.*

Let us follow value output ironing for X_2*, tracking decision changes from time*

$1 = a_2$ *to* $2 = t_2^*$ *as 2's value is increased. Time 1 scenario breakpoints are* $3, 4, 9$ *on* $\sigma^{1,2}, \sigma^{3,4}$ *and* $\sigma^{5,6,7}$ *respectively. All 7 time 2 breakpoints are at 0.5.*

Denote by $\mathbf{C}_t$ *the* $\mathbf{C}$ *decision at* t*. If* X_2 *had value 3,* $\mathbf{C}_1$ *would change to* $\emptyset$ *as* $\sigma^{1,2}$ *votes become* $\{X_1, X_2\}$*. Using Algorithm 3's notation,* $\{X_1\} = \mathbf{C}(\mathrm{votes}(\Sigma^{\neq j_\beta}), \alpha^<_{\mathtt{s}}) {\neq} \emptyset = \mathbf{C}(\mathrm{votes}(\Sigma^{\neq j_\beta}), \alpha^>_{\mathtt{s}})$ *at* $t^\beta = 1$*, for* $j_\beta = 1$ *and 2,* $r_\beta = 3$*,* $\alpha^<_{\mathtt{s}} = \{X_1\}$ *and* $\alpha^>_{\mathtt{s}} = \{X_1, X_2\}$*. All 7 time 2 breakpoints are now at 0 since* $X_{2,3}$ *can both be allocated.* X_2 *is still allocated at time 2, surviving ironing so far.*

If X_2 *had value 4,* $\mathbf{C}_1$ *would change again to* $\{X_1, X_2\}$ *as* $\sigma^{3,4}$ *votes become* $\{X_1, X_2\}$*;* $t_2 = 1$ *and any time 2 breakpoint is discarded since no items are left. Last breakpoint is at 9: all votes are now for* $\{X_1, X_2\}$*. As allocation times never increase as* X_2*'s value increases, isIroned returns* `false`*:* X_2 *survives ironing and is allocated at time 2.*

In this example, t^β *always equaled* t_i*, precluding the need for updating* $\vec{S}, \vec{A}$ *and breakpoints.*

Consider an agent i that can be allocated, i.e. with $q_i \leq S_i$. For simple `Select` methods, such as `OnlyDep` which selects only those agents that are departing in the current period, the only breakpoint on scenario j in period t for agent i with type $\theta_i = (a_i, d_i, r_i, q_i)$ is given by:

$$r_o^j(i) = V(S^t, A^t \backslash \{i\}, \sigma^j) - V(S^t - q_i, A^t \backslash \{i\}, \sigma^j) \tag{3.5}$$

where by $V(S, A, \sigma^j)$ I denote the value of the solution of the offline optimization problem $\mathrm{Opt}(S, A, \sigma^j)$. This follows from the simple combinatorics of the offline weighted knapsack problem.

I will soon introduce more nuanced `Select` methods in which there can be multiple scenario breakpoints, with the decision that receives a vote changing more than once. This makes sensitivity analysis, and thus ironing, a bit more tricky.

The payments of agents must be computed as the critical value of the ironed policy. For this, the procedure outlined above for ironing is essentially reversed: one steps down lower values until the agent would be unallocated in policy π, or allocated but fail the ironing test and thus be unallocated in ironed policy $\breve{\pi}$ [71].

3.2.3 Departure obliviousness and myopic monotonicity

The obvious concern with ironing, which cancels decisions and discards resources, is that it may establish monotonicity at the expense of destroying the value of a policy by canceling almost all decisions.

In fact, if this was a problem of stochastic optimization with cooperative agents then it would be optimal to delay any allocation decision until an agent's departure, since this is without cost to the agent and allows the center to receive more information about agent types. This is encapsulated in the `OnlyDep` select method: only allocate to those agents that are in the majority vote decision and depart now.

But as explained earlier, this would lead to a very low quality policy when coupled with output ironing. Output ironing would fail to establish the monotonically-earlier property of Eq. (3.3) for any types except maximally-patient agents, and cancel most allocation decisions.

I will focus on *departure oblivious* `Select` methods, i.e. invariant to an allocated agent's delay of departure:

Definition 27. *Policy π is* departure-oblivious *if for any agent i allocated in period t_i^*, the decisions made by the policy for periods $a_i \leq t \leq t_i^*$ do not change for any reported departure $d_i' > d_i$, holding all other inputs unchanged.*

This property trivially implies monotonicity with respect to departure. In combination with ironing with respect to arrival, value and quantity only ("(a, r, q)-ironing"), i.e. checking Eq. (3.4) only for types θ_i'' and θ_i' that differ from θ_i in these attributes, this provides full monotonicity.

Proposition 12. *For a departure-oblivious policy, (a, r, q)-ironing is equivalent to ironing.*

Proof. Let $\check{\pi}$ denote the policy obtained by (a, r, q)-ironing departure-oblivious policy π.

Fix θ_{-i} and ω, and consider type θ_i with departure d_i such that $t_i^{\pi}(\theta_i) = t$.

Both $\check{\pi}$ and $\breve{\pi}$ can only cancel decisions of the same base policy π: $t_i^{\check{\pi}}(\theta_i), t_i^{\breve{\pi}}(\theta_i) \in \{\infty, t\}$. If $t_i^{\breve{\pi}}(\theta_i) = t$ then $t_i^{\check{\pi}}(\theta_i) = t$ as well, as the set of checks for $\check{\pi}$ is a subset of the one for $\breve{\pi}$.

Suppose now that $t_i^{\breve{\pi}}(\theta_i) = \infty$. There exists then $\theta_i'' \succeq_\theta \theta_i' \succeq_\theta \theta_i$ such that $t_i^{\pi}(\theta_i'') > t_i^{\pi}(\theta_i')$. Let $\tilde{\theta}_i''$ and $\tilde{\theta}_i'$ equal θ_i'' and θ_i', but with $\tilde{d}_i'' = \tilde{d}_i' = d_i$. By departure obliviousness, $t_i^{\pi}(\tilde{\theta}_i'') \in \{\infty, t_i^{\pi}(\theta_i'')\}$. If $t_i^{\pi}(\tilde{\theta}_i'') = \infty$ then $t_i^{\check{\pi}}(\theta_i) = \infty$ as well, since $\tilde{\theta}_i''$ and θ_i have the same departure. Suppose that $t_i^{\pi}(\tilde{\theta}_i'') = t_i^{\pi}(\theta_i'')$ and, similarly, that $t_i^{\pi}(\tilde{\theta}_i') = t_i^{\pi}(\theta_i')$. Then $t_i^{\check{\pi}}(\theta_i) = \infty$ as i would be (a, r, q)-ironed: $t_i^{\pi}(\tilde{\theta}_i'') > t_i^{\pi}(\tilde{\theta}_i')$. □

Note the different decompositions of ironing across the dimensions of a bidder's type: Algorithms 2 and 3 separate value from the other dimensions, whereas departure obliviousness precludes the need for departure ironing.

In terms of constraints imposed on the policy, *myopic* monotonicity is at the opposite end of the spectrum from obliviousness (see Table 3.1 for a summary of the various monotonicity notions, when restricted to departure).

Definition 28. *A* myopically monotonic *policy π is such that, if i, who is allocated at t_i^* for θ_i, reports $\theta_i' \succeq_\theta \theta_i$ for which actions under π are identical with actions for θ_i until time $t_i^* - 1$, then π still allocates i at t_i^* for the different report:*

$$\textit{if } \pi^t(\theta_i') = \pi^t(\theta_i)\, \forall\, t \le t_i^* - 1 \textit{ and } \theta_i' \succeq_\theta \theta_i \textit{ then } i \sqsubset \pi^{t_i^*}(\theta_i')$$

Property	Requirement
Myopic departure-monotonicity	if $\pi^t(d_i) = \pi^t(d'_i)\,\forall\, t \leq t^*_i - 1$ then $i \sqsubset \pi^{t^*_i}(d'_i)$
Monotonicity (Def. 24)	$t^\pi_i(d'_i) \in [a_i, d'_i]$
Ironing condition (Eq. (3.3))	$t^\pi_i(d'_i) \leq t^*_i = t^\pi_i(d_i)$
Departure obliviousness	$\pi^t(d_i) = \pi^t(d'_i)\,\forall\, t \leq t^*_i$

Table 3.1: Requirements under increasingly stronger notions of departure monotonicity on policy π when agent i allocated at time t^*_i changes his reported departure from d_i to $d'_i > d_i$. Monotonicity does not strictly imply myopic departure-monotonicity.

Note that myopic monotonicity only requires that i is allocated; the other active agents' allocation decisions may change. If all agents are impatient then myopic monotonicity is identical to (anytime) monotonicity. I will often restrict myopic monotonicity to one dimension of a bidder's type: θ_i and θ'_i of Def. 28 are then identical in the other dimensions.

Clearly, if π is departure-oblivious then it also is myopically departure-monotonic.

While necessary for anytime-monotonicity, myopic monotonicity does not preclude output ironing if decisions earlier than t^*_i change for θ'_i. Without any information on the policy, if some $\theta'_i \succeq_\theta \theta_i$ is not checked, for θ'_i the agent may be allocated later than for θ_i, or not allocated at all. However, all `Select` methods I introduce are departure-oblivious, and therefore anytime departure monotonic. Hence, none of them needs departure output-ironing; only method `OnlyDep` will employ it. They do however need (a, r, q)-ironing.

3.2.4 Flexibility: virtual values and learning

In the context of single item, static auctions, Myerson [64] proved the revenue-optimality of an efficient auction defined on "virtual valuations". The virtual valuation of a bidder whose bid r_i is drawn with probability density function (pdf) f and cumulative distribution function (cdf) F is:

$$\tilde{w}(r_i) = r - \tfrac{1-F(r_i)}{f(r_i)} \tag{3.6}$$

My approach to dynamic auction design provides considerable flexibility. For example, one can apply the algorithm essentially unchanged, except for substituting valuations with virtual valuations. Each reported type is converted into a virtual valuation by retaining (a_i, d_i, q_i) but replacing r_i with $\tilde{w}(r_i)$. All computation is then performed with respect to virtual valuations: samples are taken from the distribution on virtual valuations induced by the distribution on valuations, ironing is performed with respect to a partial order defined on virtual valuations, and the payment is first determined as the critical "virtual value" and then transformed into the corresponding actual value.

When the distribution f has a non-decreasing hazard rate (as required by Myerson [64]), the ironed policy remains monotonic and thus strategyproof. Without this property, then it would be necessary to also adopt Myerson's notion of ironing to first transform the virtual valuation function into a monotone increasing function, and adopt this as the mapping from values into virtual values. Myerson's transformation could be termed "input ironing" whereas the ironing adopted here is "output ironing."

As another indication of this approach's flexibility, the self-correcting methodology advanced here does not require that the mechanism has correct information about the underlying distribution on types. The distribution can simply be learned over time, for example through a non-parametric approach that samples from the past such as that proposed by van Hentenryck and Bent [84]. In order to retain strategyproofness, it is necessary to preclude the reported type of an agent until the agent has itself departed from the system. This way, the report of an agent cannot affect the mechanism's distributional model while the agent still cares about the model.

3.3 Basic Select methods

In this section, I introduce some basic Select methods that are departure oblivious and therefore useful together with output ironing for the design of strategyproof, dynamic multi-unit auctions. An already encountered method is

$$\texttt{OnlyDep}\colon \texttt{Select}(\alpha^j, \Sigma, S^t, A^t) = \alpha^j|_{d=t}$$

`OnlyDep` is clearly not departure oblivious.

For a straw man method that is departure oblivious, I use the identity `Select` method that ignores all departure information:

$$\texttt{IgnoDep} : \texttt{Select}(\alpha^j, \Sigma, S^t, A^t) = \alpha^j \tag{3.7}$$

`HazRate` only selects bidders somewhat likely to leave soon:

$$\texttt{HazRate} : \texttt{Select}(\alpha^j, \Sigma, S^t, A^t) = \left\{ i \sqsubset \alpha^j \colon \frac{1-F_i^D(d_i)}{f_i^D(d_i)} < c \right\} \tag{3.8}$$

where departures have pdf f_i^D and cdf F_i^D, and $c \in (0,1)$ is a parameter. Agent i is retained in α_{s}^j by `HazRate` iff his reported departure d_i is late enough.

`HROrRew` also selects bidders with high value-per-item:

$$\begin{aligned}\texttt{HROrRew} &: \texttt{Select}(\alpha^j, \Sigma, S^t, A^t) = \\ &\{ i \sqsubset \alpha^j \colon \left(\frac{1-F_i^D(d_i)}{f_i^D(d_i)} < c \right) \vee \left(\mathbb{P}[\frac{R}{Q} > \frac{r_i}{q_i}] < w \right) \}\end{aligned} \tag{3.9}$$

where parameters $c \in (0,1)$ and $w \in (0,1)$, and R and Q are random variables denoting an agent's value and quantity. That is, if i is "too good to miss" then he is selected even if his departure does not satisfy Eq. (3.8). Parameters c and w can be optimized for the distribution on agent types to maximize the performance of **C**$\oplus$`HazRate` or **C**$\oplus$`HROrRew`.

Lemma 1. **C**$\oplus$`IgnoDep` *is departure-oblivious. If* F_i^D *has a monotone non-decreasing hazard rate (i.e. it is* regular*) then* **C**$\oplus$`HROrRew` *and* **C**$\oplus$`HazRate` *are departure-oblivious.*

Proof. The claim for select method `IgnoDep` is immediate. For $\mathbf{C}\oplus$`HROrRew` and $\mathbf{C}\oplus$`HazRate`, let us assume that agent i is allocated at t_i^* when reporting departure d_i. If d_i satisfies the departure condition in Eq. (3.8) and Eq. (3.9), and F_i^D is regular then all $d_i^+ > d_i$ also satisfy this condition. As decisions before t_i^* are unchanged, the t_i^* decisions for d_i^+ and d_i will be identical since for any scenario j at time t, the event $i \sqsubset \mathrm{Opt}^j$ is independent of i's departure. Therefore `HROrRew` and `HazRate` are departure-oblivious. $\square$

Just as with `OnlyDep`, the selected subset of agents $\alpha_{\mathtt{s}}^j$ in scenario σ^j will change at most once as the value of agent i is increased with the `HazRate` method. The change, if any, occurs if $q_i \leq S^t$ and at bid value $r_o^j(i) = V(S^t, A^t \setminus \{i\}, \sigma^j) - V(S^t - q_i, A^t \setminus \{i\}, \sigma^j)$. Agent i is in $\alpha_{\mathtt{s}}^j$ for $r_i \geq r_o^j(i)$ if and only if Eq. (3.8) is satisfied.

When the departure condition in `HROrRew` is satisfied then this behaves as `HazRate` and there is one breakpoint at $r_o^j(i)$ for an agent that can be feasibly allocated. But otherwise, there can be two breakpoints when the value $r_j^c(i)$, at which $\mathbb{P}[\frac{R}{Q} > \frac{r_j^c(i)}{q_i}] = w$, is greater than $r_o^j(i)$. In this case, for $r_i \in [r_o^j(i), r_j^c(i))$ agent i is in α^j but not selected and then for $r_i \in [r_j^c(i), \infty)$ agent i is also selected. More importantly, within ironing, `HROrRew` only yields one breakpoint: the condition in Eq. (3.9) is independent of time and is satisfied by all higher types if satisfied by a certain type.

Table 3.2 shows that all simple methods have a *single* value breakpoint on a scenario j at t for an agent i that is allocated at $t_i^* \geq t$. Since i is allocated at t_i^*, if $r_i \geq r_o^j(i)$ then i is also in j's vote at t: the additional `Select` test is independent of time for both `HazRate` and `HROrRew`. Recall that only breakpoints for an agent that is allocated matter.

I denote by $\mathrm{Opt}_{j,\pm\infty} = \mathrm{Opt}^j(S^t, A^t|_{r_i:=\pm\infty})$ the offline optimum for scenario j at t given supply S^t and active agents A^t when i's value is changed to $\pm\infty$. As sets, $\mathrm{Opt}_{j,+\infty} = \{i\} \cup \mathrm{Opt}^j(S^t - q_i, A^t \setminus \{i\})$ and $\mathrm{Opt}_{j,-\infty} = \mathrm{Opt}^j(S^t, A^t \setminus \{i\})$.

Method `OnlyDep` is myopically-monotonic in arrival and value but not departure.

Value r_i	$[0, r_o^j(i))$	$[r_o^j(i), \infty)$
Interpretation	$i \not\sqsubset \text{Opt}^j$	$i \sqsubset \text{Opt}^j$
Vote	$\text{Opt}_{j,-\infty}$	$\text{Opt}_{j,+\infty}$

Table 3.2: Possible votes of `IgnoDep`, `HazRate` or `HROrRew` on scenario j at time t for different values of agent i *allocated* at time $t_i^* \geq t$, given others' reports.

None of the methods are myopically quantity-monotonic, because reporting a lower quantity by a winning agent $i \sqsubset \text{Opt}^j$ may induce a vote change on j to some other action $i \sqsubset \alpha'$ possibly resulting in a change of the Consensus decision.

The following Proposition provides a necessary condition for a violation of myopic quantity monotonicity in the basic `Select` methods. The rather involved nature of the necessary condition offers a non-rigorous explanation for the scarcity of ironing with respect to quantity.

Proposition 13. *Consider a time t and method* `S` *being one of* `OnlyDep`, `IgnoDep`, `HazRate` *or* `HROrRew`. *Bidder i, who is in the* **C** *decision $\pi^t_{\mathbf{C}\oplus\mathtt{S}}$ at t for θ_i, reports $q_i' < q_i$ for which actions under $\mathbf{C} \oplus \mathtt{S}$ are unchanged until time $t-1$. Suppose that i is no longer allocated at t for $q_i' : i \not\sqsubset \pi^t_{\mathbf{C}\oplus\mathtt{S}}$. Then there must exist a partition $\Sigma_a \cup \Sigma_b \cup \Sigma_c \cup \Sigma_d = \Sigma$ of scenarios such that*

- *on each $\sigma^{j_a} \in \Sigma_a$, the voted allocation is the same ($\alpha^a_{\mathtt{S}}$) and includes i for q_i and q_i'. Any active agent in exactly one of the offline optima α^a and $\alpha^{a\prime}$ must not be selected by* `S`.

- *on each $\sigma^{j_b} \in \Sigma_b$, the voted allocation when i reports q_i is $\alpha^a_{\mathtt{S}}$ (the same as for Σ_a). The voted allocation for q_i' contains at least one active agent not allocated by α^a; in particular, $v(\alpha^a \backslash \{i\}) + V(S^t - \#(\alpha^a \backslash \{i\}), \emptyset, \sigma^{j_b}) < V(S^t - q_i', A^t \backslash \{i\}, \sigma^{j_b})$*

- *on each $\sigma^{j_c} \in \Sigma_c$, i is not in the voted allocation when i reports q_i or q_i'. That is, $r_i + V(S^t - q_i', A^t \backslash \{i\}, \sigma^{j_c}) < V(S^t, A^t \backslash \{i\}, \sigma^{j_c})$; otherwise i would be selected if in the* *offline optimum for either quantity.*

- *on any* $\sigma^{j_d} \in \Sigma_d$*, the offline optimum is neither* α^a *nor* α^c *and it may or may not allocate* i *for* q_i *or* q'_i
- $|\Sigma_a| \leq |\Sigma_c| \leq |\Sigma_a| + |\Sigma_b|$

Proof. Take $\alpha^a_{\mathtt{s}}$ and $\alpha^c_{\mathtt{s}}$ to be the $\mathbf{C} \oplus \mathtt{S}$ decisions for q_i, respectively q'_i. Take Σ_a (respectively Σ_c) to be the scenarios on which the voted allocation is $\alpha^a_{\mathtt{s}}$ (respectively $\alpha^c_{\mathtt{s}}$) for both q_i and q'_i. Take Σ_b to be the scenarios on which the voted allocation is $\alpha^a_{\mathtt{s}}$ *only* for q_i. On any Σ_b scenario, i must be allocated for q'_i: i remains in the offline optimum and still passes the `Select` test for the lower quantity. Finally, take Σ_d as all scenarios not in Σ_a, Σ_b or Σ_c.

Otherwise the voted allocation would be either $\alpha^a_{\mathtt{s}}$ nor $\alpha^c_{\mathtt{s}}$. □

A bidder i may no longer in the $\mathbf{C} \oplus \mathtt{S}$ decision after lowering his quantity even if all other bidders have unit-demand and $q'_i = 1$:

Example 5. $q_i = 2, q'_i = 1, S^t = 3, r_i = \$6, A^t = \{\$2, i\}$*. 5 copies of* σ_{j^a} *with* $\{\$7, \$3\}$ $(\alpha^a = \alpha^{a\prime} = \{i\})$*, 5 copies of* σ_{j^b} *with* $\{\$7\}$ $(\alpha^{b\prime} = \{\$2, i\})$*, 8 copies of* σ_{j^c} *with* $\{\$7, \$7, \$7\}$ $(\alpha^c = \alpha^{c\prime} = \emptyset)$ *and 2 copies of* σ_{j^d} *with* $\emptyset$ $(\alpha^d = \alpha^{d\prime} = \{\$2, i\})$*. If allocated in the offline optimum then* $\$2$ *or* i *are selected.*

3.4 The NowWait heuristic

In this section I describe `NowWait`, a departure-oblivious `Select` method that makes an explicit tradeoff between the value of allocating to an agent that could disappear and the likely benefit of waiting for other opportunities.

The `NowWait Select` method filters α^j down to $\alpha^j_{\mathtt{s}}$ by retaining those agents for which the estimated value from allocating now is greater than the estimated value from waiting, considering that i's value may be lost because he may depart:

$$\texttt{NowWait}: \texttt{Select}(\alpha^j, \Sigma, S^t, A^t) = \{i \sqsubset \alpha^j: \mathrm{now}_i^t(\alpha^j, r_i) \geq \mathrm{wait}_i^t(\alpha^j, r_i)\}$$

For estimating the future value on a scenario j', I make the pessimistic assumption that all agents present in A^t (except i) either depart or are allocated in this period so that the future demand is represented only by that in each scenario $j' \in \Sigma$. The global estimate is simply an average over the per-scenario estimates.[6]

Let $\alpha^j_- = \alpha^j \setminus \{i\}$ and $v(\alpha^j_-)$ denote the total value to the agents allocated in α^j_-. Let $\#(\alpha)$ denote the number of items allocated by action α. The value obtained by allocating to agent i with value r_i in period t is estimated as:

$$\mathrm{now}_i^t(\alpha^j, r_i) = r_i + v(\alpha^j_-) + \tfrac{1}{|\Sigma|} \textstyle\sum_{j' \in \Sigma} V(S^t - \#(\alpha^j), \emptyset, \sigma^{j'}),$$

Let $\rho = \rho_t$ be the probability that agent i will still be present in the next period $t+1$ given type θ_i but ignoring his reported departure (to provide departure obliviousness):

$$\rho = \mathbb{P}[D > t | D \geq t, a_i, r_i, q_i] \tag{3.10}$$

The value for waiting to allocate i is estimated as:

$$\begin{aligned}\mathrm{wait}_i^t(\alpha^j, r_i) = v(\alpha^j_-) &+ (1-\rho) \tfrac{1}{|\Sigma|} \textstyle\sum_{j' \in \Sigma} V(S^t - \#(\alpha^j_-), \emptyset, \sigma^{j'}) \\ &+ \rho \tfrac{1}{|\Sigma|} \textstyle\sum_{j' \in \Sigma} V(S^t - \#(\alpha^j_-), \{i\}, \sigma^{j'})\end{aligned}$$

Since for any i and σ^j, α^j, $\mathrm{now}_i^t(\alpha^j, r_i)$ and $\mathrm{wait}_i^t(\alpha^j, r_i)$ are independent of i's reported departure one gets

Proposition 14. **C**$\oplus$`NowWait` *is departure-oblivious.*

[6] If I were to retain an unallocated agent i' and consider the presence of this agent when computing opportunity costs for i in scenario j', then ironing for i' would also need to analyze the effect of i' raising his value on the allocation decision for i and maybe other agents at t because of this coupling effect through the `NowWait` select rule. I wish to avoid this additional complication.

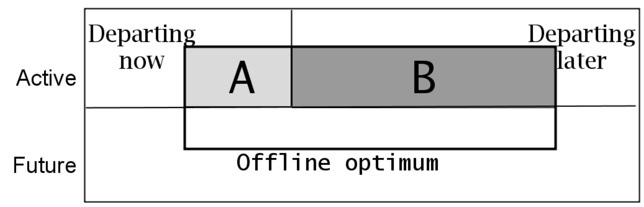

Figure 3.1: Subsets of agents for `Select` on a scenario.

Figure 3.1 highlights the subsets of interest among active and future (sampled) agents on a scenario. `Select` methods select subsets of active agents in the offline optimum: A and B for `IgnoDep`, A (only departing now) for `OnlyDep` and subsets of A and B (*a priori* "urgent") for `NowWait`.

3.4.1 Sensitivity analysis

`NowWait` requires sensitivity analysis of the resulting **C** algorithm. For this it is crucial to be able to compute scenario breakpoints to determine when the decision $\alpha^j_{\mathtt{S}}$ changes as agent i's value varies. There may be two breakpoints per scenario. The first occurs at $r^j_O(i)$ when agent i enters α^j and the second at $r^j_C(i)$ when agent i is retained by the `Select` method.

To better understand `NowWait`'s behavior, use shorthand $v_{j'} = V(S^t - \#\alpha^j, \emptyset, \sigma^{j'})$ and $c_{j'} = V(S^t - \#\alpha^j_-, \emptyset, \sigma^{j'}) - V(S^t - \#\alpha^j, \emptyset, \sigma^{j'})$. $c_{j'}$ is the opportunity cost incurred by allocating to i if scenario j' was the actual future. If $\alpha^j = \{i\}$ and $j' = j$ then $c_j = r^j_O(i)$. With this, $\text{now}^t_i(\alpha^j, r_i) = r_i + v(\alpha^j_-) + \frac{1}{|\Sigma|}\sum_{j'\in\Sigma} v_{j'}$ and $\text{wait}^t_i(\alpha^j, r_i) = v(\alpha^j_-) + (1 - \rho)\frac{1}{|\Sigma|}\sum_{j'\in\Sigma}(c_{j'} + v_{j'}) + \rho\frac{1}{|\Sigma|}\sum_{j'\in\Sigma}(\max(r_i, c_{j'}) + v_{j'})$, where the final term comes from recognizing that $V(S^t - \#(\alpha^j_-), \{i\}, \sigma^{j'}) = \max(r_i + v_{j'}, c_{j'} + v_{j'}) = \max(r_i, c_{j'}) + v_{j'}$. Simplifying, agent i, allocated in α^j, is retained if and only if:

$$r_i|\Sigma| \geq (1 - \rho)\sum_{j'\in\Sigma} c_{j'} + \rho\sum_{j'\in\Sigma}\max(r_i, c_{j'}) \tag{3.11}$$

Note that, apart from the number of items (implicit in $c_{j'}$) allocated by α^j_-, the condition in Eq. (3.11) is independent of scenario j.

I claim that if, for some scenario $\sigma^{j'}$, $\alpha^{j'}$ allocates the same number of items as α^j and allocates i then $r_i \geq c_{j'}$. Indeed,

$$r_i + v(\alpha^{j'}_-) + V(S^t - \#\alpha^{j'}, \emptyset, \sigma^{j'}) \geq v(\alpha^{j'}_-) + V(S^t - \#\alpha^{j'}_-, \emptyset, \sigma^{j'})$$

and $c_{j'}$ is the difference of the two V terms, since I assumed that $\#\alpha^{j'} = \#\alpha^j$.

In general, the value and patience of an agent may be correlated according to type distribution $f(\theta_i)$. This will complicate sensitivity analysis for $\mathbf{C}\oplus$`NowWait`. For now I assume that patience and value are independent, and thus ρ does not depend on an agent's value r_i.

Given this independence assumption, there is a threshold $\tau^*_\rho(i)$, for i in α^j to be selected by `NowWait` such that:

$$\mathrm{now}^t_i(\alpha^j, r_i) \geq \mathrm{wait}^t_i(\alpha^j, r_i) \text{ if and only if } r_i \geq \tau^*_\rho.$$

Theorem 10 will prove a more general statement.

Denote the second scenario breakpoint by $r^j_c(i) = \max(r^j_o(i), \tau^*_\rho)$. It may be the case that $\tau^*_\rho > r^j_o(i)$ and thus $r^j_c(i) > r^j_o(i)$ if, for example, scenario j predicts less demand than the other scenarios and i is likely very patient (ρ is close to 1). For values in the range $[0, r^j_o(i))$, neither α^j nor $\alpha^j_{\mathtt{S}}$ include agent i. For values in the range $[r^j_o(i), r^j_c(i))$, i is included in α^j, but is removed by `NowWait`: $i \not\subset \alpha^j_{\mathtt{S}}$. For values in the range $[r^j_c(i), \infty)$, both α^j and $\alpha^j_{\mathtt{S}}$ include i.

Table 3.3 summarizes `NowWait` breakpoints. Out of $\alpha^-_{\mathtt{S}}, \alpha^+_{\mathtt{S}} \setminus \{i\}$ and $\alpha^+_{\mathtt{S}} \cup \{i\}$, one is the actual vote on j; only $\alpha^-_{\mathtt{S}}$ and $\alpha^+_{\mathtt{S}} \setminus \{i\}$ can coincide as only $\alpha^+_{\mathtt{S}} \cup \{i\}$ allocates i.

Let $\phi_\rho(r_i) = r_i|\Sigma| - (1-\rho)\sum_{j' \in \Sigma} c_{j'} - \rho \sum_{j' \in \Sigma} \max(r_i, c_{j'}) = \mathrm{now}^t_i(\alpha^j, r_i) - \mathrm{wait}^t_i(\alpha^j, r_i)$.

Value r_i; interpretation	$[0, r_o^j(i))$ $i \not\sqsubset \text{Opt}^j$	$[r_o^j(i), r_c^j(i))$ $i \sqsubset \text{Opt}^j$	$[r_c^j(i), \infty)$ $i \sqsubset \text{Opt}^j$
Reference α^j	$\text{Opt}_{j,-\infty}$	$\text{Opt}_{j,+\infty}$	$\text{Opt}_{j,+\infty}$
Vote (action); interpretation	$\alpha_{\mathtt{S}}^-$ ignore i	$\alpha_{\mathtt{S}}^+ \setminus \{i\}$ wait for i	$\alpha_{\mathtt{S}}^+ \cup \{i\}$ allocate i

Table 3.3: Possible votes of `NowWait` for different values of agent i, when the reports of other agents are fixed. $\text{Opt}_{j,\pm\infty} = \text{Opt}^j(S^t, A^t|_{r_i:=\pm\infty})$ denotes the offline optimum for scenario j given supply S^t and active agents A^t when i's value is changed to $\pm\infty$.

Proposition 15 establishes the intuitive property that the more likely an agent i is to still be present in the next period, the higher value i needs to have in order not to be filtered from the decision by the `NowWait` `Select` method.

Proposition 15. *As ρ increases from 0 to 1, the threshold τ_ρ^* at which an agent $i \sqsubset \alpha^j$ is selected by* `NowWait` *weakly increases from the average to the maximum of the costs $\{c_{j'} : j' \in \Sigma\}$, when agent patience is independent of value.*

Proof. One gets $\tau_\rho^* \leq \tau_{\rho'}^*$ for $\rho \leq \rho'$ as it is easy to check that for all r_i and all ρ, $\phi_\rho(r_i)$ is non-decreasing in ρ.

For $\rho = 0$, the threshold τ_0^* is the average of marginal costs $\{c_{j'} : j' \in \Sigma\}$: $\tau_0^*|\Sigma| = \sum_{j' \in \Sigma} c_{j'}$. For $\rho = 1$, $\tau_1^*|\Sigma| = \sum_{j' \in \Sigma} \max(\tau_1^*, c_{j'})$ and thus $\tau_1^* = \max_{j'} c_{j'}$. □

Say there are only two scenarios $\sigma_{1,2}$ with opportunity costs $c_1 < c_2$. Consider the computation of agent i's τ_ρ^* on scenario σ_1. Given that i is allocated on scenario σ_1, $r_i \geq c_1$. Suppose that $r_i \leq c_2$; otherwise r_i clearly satisfies Eq. (3.11). Eq. (3.11) amounts to $2r_i \geq (1-\rho)(c_1 + c_2) + \rho(r_i + c_2)$ i.e. $r_i \geq \frac{(1-\rho)c_1 + c_2}{2-\rho}$ and Proposition 15 can be readily verified.

My implementation of sensitivity analysis accounts for the fact that, by varying a_i (and thus ρ) or q_i, the threshold τ^* for `NowWait` may change.

r_i	2 copies of σ^j	4 copies of σ^{j_4}	5 copies of σ^{j_5}	**C**
$[r_c^j, \infty)$	Opt: $\alpha_4 \cup \{i\}$ vote: $\alpha_4 \cup \{i\}$	Opt: any vote: any	Opt: $i \sqsubset \alpha_5$ vote: $i \sqsubset \alpha_5$	α_5 or $\alpha_4 \cup \{i\}$
$[r_o^j, r_c^j)$	Opt: $\alpha_4 \cup \{i\}$ vote: α_4	Opt: $i \not\sqsubset \alpha_4$ vote: $i \not\sqsubset \alpha_4$	Opt: $i \sqsubset \alpha_5$ vote: $i \sqsubset \alpha_5$	α_4
$[0, r_o^j)$	$i \not\sqsubset$Opt=vote=α^j	Opt=vote=α_4	Opt=vote=α_5	α_5

Table 3.4: An example with two distinct breakpoints for `NowWait`. When i's value is r, decisions α_5, α_4 and α^j receive 5, 4 and 2 votes respectively: the **C** decision is α_5 which allocates i (thus $r \geq r_c^{j_5}(i)$). When i's value is $r' > r$, α_5 and α_4 receive 5 and 6 votes respectively: the **C** decision is α_4, not allocating i. However, $i \sqsubset \mathbf{C}$ for values higher than $r_c^j(i)$.

3.4.2 Monotonicity properties of `NowWait`

I unveil now a potential failure of myopic value monotonicity in `NowWait`. Consider two values $r < r'$ for which decisions on scenarios other than σ^j are identical. If $i \sqsubset \text{Opt}(S, A, \sigma^j)$ for r then $\text{Opt}(S, A, \sigma^j)$ will not change for r'. Thus, if $i \sqsubset \mathbf{C}$ for r, then the same holds for r'. If $i \not\sqsubset \text{Opt}(S, A, \sigma^j)$ for r' then the σ^j vote is the same for r or r': i does not influence any bidder's r_c^j.

However, suppose that $i \sqsubset \mathbf{C}$ for r and $r < r_o^j(i) < r' < r_c^j(i)$, i.e. $i \sqsubset \text{Opt}(S, A, \sigma^j)$ only for r' but even then i is not "urgent" enough (this is the setting of Table 3.4). Then i may no longer be in the winning Consensus allocation for r'. Therefore, as defined, `NowWait` is not myopically value-monotonic. By value output ironing, i will only be selected by `NowWait` when his value is at least

$$v^{\texttt{NW}}(i) := \inf_r \{\, \forall\, r' \geq r, i \sqsubset \mathbf{C}(\{\text{brkPt}_j(i)\}_{\sigma^j \in \Sigma}, r_i := r')\}$$

This potential lack of monotonicity is only for a bounded range of values: if $i \sqsubset \text{Opt}^j(S, A)\, \forall\, j$ and $r_i \geq \max_{j,j'} c_{j'}(\sigma^j)$ on any scenario j then $i \sqsubset \mathbf{C}$ for any ρ. That is, $v^{\texttt{NW}}(i) \leq \max_{j,j'} c_{j'}(\sigma^j)$. Furthermore,

Proposition 16. `NowWait` *is myopically value-monotonic for an agent i whose $r_i \geq v^{\texttt{NW}}(i)$.*

As Section 3.5 shows, `NowWait`'s ironing cancellations are limited. I attribute this success to myopic monotonicity:

Proposition 17. `NowWait` *is myopically arrival-monotonic if the patience distribution is regular.*

Proof. Consider the effect of i reporting arrival $a_i - 1$ instead, such that $\pi^t(a_i - 1) = \pi^t(a_i)\,\forall\, t \leq t_i^* - 1$.

I claim that the vote on a scenario σ^j at t_i^* can only change to an action including i. Whether i is in Opt^j is independent of i's arrival. If the vote changes then $r_c^j(i)$ must have changed. Let ρ (respectively ρ^+) denote $\mathbb{P}[\Delta > t - a | \Delta \geq t - a]$ when a is $a_i - 1$, respectively a_i. Since the patience distribution is regular, $\rho \leq \rho^+$. By Proposition 15, $\tau_\rho \leq \tau_{\rho^+}$. Therefore the actions must be $\alpha_{\mathtt{S}}^+ \setminus \{i\}$ for arrival a_i and $\alpha_{\mathtt{S}}^+ \cup \{i\}$ for $a_i - 1$, for some action $\alpha_{\mathtt{S}}^+$.

Let $v_+^{\mathtt{NW}}(i)$, respectively $v^{\mathtt{NW}}(i)$ be the $v^{\mathtt{NW}}$ thresholds for a_i respectively $a_i - 1$. I claim $v_+^{\mathtt{NW}}(i) \geq v^{\mathtt{NW}}(i)$. Suppose that the contrary was true; let $r \in (v_+^{\mathtt{NW}}(i), v^{\mathtt{NW}}(i))$: if i had value r, the Consensus decision would include i for $a_i - 1$, but not for a_i. As the only vote counts that can increase are for decisions that allocate i, i must also be in the Consensus decision for $a_i - 1$, contradicting $r < v^{\mathtt{NW}}(i)$. □

Surprisingly, I observed more ironing experimentally with respect to reward rather than quantity. The following Proposition establishes a partial monotonicity with respect to quantity of `NowWait`. It suggests a reason for the relative scarcity of ironing with respect to quantity observed.

Proposition 18. *Fix quantities $\underline{q}_i < q_i$ and a scenario j such that i is selected by* **C**$\oplus$`NowWait` *(i.e. $r_i \geq \tau_\rho^*$) for q_i. Let $\underline{\alpha}^j$ be the allocation to active agents in the offline optimal solution when i reports the lower quantity $\underline{q}_i$. Clearly $\underline{\alpha}^j$ allocates i.*

Suppose that one *of the following conditions is satisfied*

A. *all other bidders have unit demand*

B. $\underline{\alpha}^j$ *allocates the same number of items as* α^j

C. $\underline{\alpha}^j_-$ *allocates the same number of items as* α^j_-

Then $c_{j'} \geq \underline{c}_{j'}, \forall \sigma^{j'} \in \Sigma$ *and* i *is also selected by* $\mathbf{C} \oplus \texttt{NowWait}$ *for* $\underline{q}_i$ *(i.e.* $r_i \geq \underline{\tau}^*_\rho$*).*

Proof. I will focus on proving that opportunity costs cannot increase when i reduces his quantity. The fact that i is still selected by $\mathbf{C}\oplus\texttt{NowWait}$ follows after recalling that i being selected amounts to $r_i|\Sigma| \geq (1-\rho)\sum_{j'\in\Sigma} c_{j'} + \rho \sum_{j'\in\Sigma} \max(r_i, c_{j'})$ and that $c_{j'}$ is defined as $V(S^t - \#\alpha^j, \emptyset, \sigma^{j'}) - V(S^t - \#\alpha^j_-, \emptyset, \sigma^{j'})$.

A. Suppose that all others have unit demand.

   I claim that if some active bidder i' (other than i) is allocated by α^j then i' is also allocated by $\underline{\alpha}^j$. Indeed, $\mathrm{Opt}(S^t, A^t, \sigma^j)$ contains i and the highest $S^t - q$ other active and sampled bidders, where q can be q_i or $\underline{q}_i$. Thus $\#\alpha^j_- \leq \#\underline{\alpha}^j_-$.

   I now claim that $\underline{\alpha}^j$ cannot allocate more *items* than α^j. Suppose the contrary: let i' be the active bidder (apart from i) with the lowest bid among all active bidders allocated by $\underline{\alpha}^j$ and i_σ be the *sampled* bidder with the *highest* bid among all sampled bidders that are *not* allocated by $\mathrm{Opt}(S^t, A^t, \sigma^j)$ when i reports $\underline{q}_i$. i' must be preferred to i_σ (in particular if $r_{i'} > r_{i_\sigma}$). As $\underline{\alpha}^j$ allocates more items than α^j, but the supply is the same, α^j cannot allocate i' and it must allocate i_σ. i' *cannot* be preferred to i_σ (in particular if $r_{i'} < r_{i_\sigma}$), a contradiction. Thus $\#\alpha^j \geq \#\underline{\alpha}^j$.

B. $\#\alpha^{j'}_- \leq \#\underline{\alpha}^{j'}_-$ since $\underline{\alpha}^j$ allocates the same number of items as α^j $(\#\alpha^{j'} = \#\underline{\alpha}^{j'})$.

The proof for case C is similar and omitted. □

Table 3.5 summarizes monotonicity properties of `NowWait` and the basic methods.

	OnlyDep	IgnoDep	HROrRew	NowWait
Value	myopic	myopic	myopic	myopic*
Departure	$\emptyset$	oblivious	oblivious	oblivious
Arrival	myopic	myopic	myopic	myopic
Quantity	none†	none†	none†	myopic‡ if unit-demand competitors

Table 3.5: Degrees of monotonicity of Select methods with respect to type components.
* if bidder's value is higher than all his opportunity costs (by Proposition 16).
† violations of myopic monotonicity appear unlikely by Proposition 13.
‡ by Proposition 18

3.4.3 Correlation of value and patience

I now consider the consequence of allowing an agent's patience to be correlated with his value. For computational tractability, it will again be important to identify a single threshold τ^* at which the agent will pass the NowWait Select test. The estimated probability ρ that an agent i will still be present in the next period becomes a function $\rho(r_i)$, that depends on i's value and therefore varies as i's value is adjusted in performing sensitivity analysis.

Theorem 10. *If $r_i \cdot (1 - \rho(r_i))$ is non-decreasing in r_i, then $\exists$ threshold τ^* with:* $\mathrm{now}_i^t(\alpha^j, r_i) \geq \mathrm{wait}_i^t(\alpha^j, r_i) \iff r_i \geq \tau^*$.

Proof. Assume $c_1 \geq \ldots \geq c_{|\Sigma|}$. Let $c_0 = \infty$ and $c_{|\Sigma|+1} = 0$.

Fix k in $0..|\Sigma|$ and let $s_k = \sum_{j'=1}^{k} c_{j'}$. I show that $\phi(r)$ (that becomes $|\Sigma| r - \sum_{j'=1}^{|\Sigma|} c_{j'} + \rho s_k - k\rho r$) is non-decreasing in r on the interval (c_k, c_{k+1}). $\phi(\cdot)$'s global monotonicity follows from its continuity. $\phi(\cdot)$'s monotonicity clearly holds for $k = 0$ since $s_0 = 0$. Let $k \geq 1$, $r' > r$ and $x = \frac{1}{k}(\phi(r') - \phi(r) - (r' - r)(|\Sigma| - k))$.

$$x = (r - \tfrac{s_k}{k})(\rho(r) - \rho(r')) + (r' - r)(1 - \rho(r')) \text{ and}$$

$$x = r'(1 - \rho(r')) - r(1 - \rho(r)) + \tfrac{s_k}{k}(\rho(r') - \rho(r))$$

If $\rho(r) \geq \rho(r')$ then $x \geq 0$ from the first line. Otherwise, as $r'(1-\rho(r')) - r(1-\rho(r)) \geq 0$, one gets $x \geq 0$ from the second line. In either case, $x \geq 0$ implying $\phi(r') - \phi(r) \geq 0$. □

Negative or moderately positive correlation between value and patience is thus sufficient for the existence and uniqueness of a threshold τ^*, such that $\text{now}_i^t(\alpha^j, \tau^*) = \text{wait}_i^t(\alpha^j, \tau^*)$, with $\text{now}_i^t(\alpha^j, r_i') \geq \text{wait}_i^t(\alpha^j, r_i')$ if and only if $r_i' > \tau^*$.[7] For $r_i(1-\rho(r_i))$ to be non-decreasing, it is sufficient for example that: (i) $\rho(r_i)$ is independent of r_i, (ii) $\rho(r_i)$ is non-increasing with r_i, or (iii) $\rho(r_i)$ is not increasing too quickly with r_i, such that $\partial\rho(r_i)/\partial r_i \leq \frac{1-\rho(r_i)}{r_i}$ for all r_i.

As the smallest solution of $\phi(\cdot) = 0$ with weakly monotone ϕ (see Theorem 10's proof), τ^* can be found via binary search in $[\frac{1}{|\Sigma|}\sum c_{j'}, \max c_{j'}]$ (recall Proposition 15). Assume for simplicity that $c_1 \geq c_2 \geq \cdots \geq c_{|\Sigma|}$ and let $c_0 = \infty$ and $c_{|\Sigma|+1} = 0$. With no correlation, the binary search can be sped up by finding c_l such that $r_c^j(i) \in [c_{l+1}, c_l]$; Eq. (3.11) becomes then a linear equation.

For the remainder of this subsection, denote possible patiences by $\{\delta_1 < \ldots < \delta_{n_\Delta}\}$. The following Proposition provides a sufficient condition for $\rho(r_i)$ to be non-increasing in r_i (case (ii)).

Proposition 19. *If patience and value are negatively correlated in the following sense: $\forall\, l = 1..n_\Delta - 1$, the ratio*

$$\frac{\mathbb{P}[\Delta = \delta_l, R = r_i]}{\mathbb{P}[\Delta = \delta_{l+1}, R = r_i]} \tag{3.12}$$

is non-decreasing in r_i, $\forall\, l = 1..n_\Delta - 1$ then $\rho(r) \geq \rho(r')$, $\forall\, r < r'$.

Proof. The fact that the agent's ρ does not increase is equivalent to $\mathbb{P}[\Delta > \delta_l | \Delta \geq \delta_l, r]$ being non-increasing with r, for all $1 \leq l \leq n_\Delta$. Thus $\mathbb{P}[\Delta > \delta | \Delta \geq \delta, r] = \frac{\mathbb{P}[\Delta > \delta, R = r]}{\mathbb{P}[\Delta \geq \delta, R = r]}$.

[7] In a unit-demand domain, Pai and Vohra [68] also require negative correlation of value and patience, in the form of a decreasing hazard rate condition in one parameter out of value, arrival and departure when the other two are fixed.

Denoting $p_l(r) = \mathbb{P}[\Delta = \delta_l, R = r]$

$$\mathbb{P}[\Delta > \delta_l | \Delta \geq \delta_l, r] = \frac{\sum_{h=l+1}^{n} p_h(r)}{\sum_{h=l}^{n} p_h(r)} = 1 - \frac{p_l(r)}{\sum_{h=l}^{n} p_h(r)}$$

I assumed that for all l, $\frac{p_{l+1}(r)}{p_l(r)}$ is non-increasing in r. This implies that so is $\frac{p_h(r)}{p_l(r)}$ for any $h \geq l+1$ or still that so is $\frac{\sum_{h=l+1}^{n} p_h(r)}{p_l(r)}$ and therefore $\frac{p_l(r)}{\sum_{h=l}^{n} p_h(r)}$ is non-decreasing. □

If patience and value are independent then the Eq. (3.12) ratio equals $\frac{\mathbb{P}[\Delta=\delta_l]}{\mathbb{P}[\Delta=\delta_{l+1}]}$, constant in r. Proposition 20 quantifies the negative correlation implied in general by the Eq. (3.12) condition.

Proposition 20. *If the domain of possible values* $R = [\underline{r}, \overline{r}] \subseteq \mathbb{R}_+$ *then the Eq.* (3.12) *condition implies* $\mathbb{E}[r|\delta_l] \geq \mathbb{E}[r|\delta_{l+1}]$.

Proof. $\mathbb{E}[r|\delta_l] - \mathbb{E}[r|\delta_{l+1}] = \int_{\underline{r}}^{\overline{r}} r(f_l(r) - f_{l+1}(r))dr$ where $f_h(r)$ denotes the density of value given patience δ_h.

I prove $f_l(\underline{r}) \leq f_{l+1}(\underline{r})$. Suppose the contrary and let $1 < c = \frac{f_l(\underline{r})}{f_{l+1}(\underline{r})}$. That implies $\frac{f_l(r)}{f_{l+1}(r)} \geq c \, \forall \, r \in [\underline{r}, \overline{r}]$ and

$$1 = \int_{\underline{r}}^{\overline{r}} f_l(r)dr \geq \int_{\underline{r}}^{\overline{r}} c f_{l+1}(r)dr \geq c > 1$$

Similarly, $f_l(\overline{r}) \geq f_{l+1}(\overline{r})$. By Eq. (3.12), for $r_0 = \inf\{r : \frac{f_l(r)}{f_{l+1}(r)} \geq 1\}$

$$f_l(r) \leq f_{l+1}(r) \text{ if } r < r_0 \text{ and } f_l(r) \geq f_{l+1}(r) \text{ if } r > r_0$$

Letting $z(r) = f_l(r) - f_{l+1}(r)$,

$$\begin{aligned} \int_{\underline{r}}^{\overline{r}} rz(r)dr &= \int_{\underline{r}}^{r_0} rz(r)dr + \int_{r_0}^{\overline{r}} rz(r)dr \\ &\geq \int_{\underline{r}}^{r_0} r_0 z(r)dr + \int_{r_0}^{\overline{r}} r_0 z(r)dr \qquad (3.13) \\ &= r_0 \int_{\underline{r}}^{\overline{r}} (f_l(r) - f_{l+1}(r))dr = r_0 - r_0 = 0 \end{aligned}$$

where Eq. (3.13) is implied by $z(r) \leq (\geq) 0$ if $r < (>) r_0$. □

3.4.4 Example: one item, two impatient bidders

I illustrate `NowWait` and `IgnoDep`, in a simple optimal stopping environment with one item for sale, two periods and one impatient, unit-demand, bidder per period with value i.i.d. sampled from distribution F. Method `OnlyDep` does not remove any agent from α^j on scenario σ^j (i.e. it is identical to `IgnoDep`) since agents are completely impatient.

Denote by v_c^1 the first bidder's critical value. As the number of scenarios tends to ∞, `OnlyDep`'s v_c^1 is the *median* of F: half of the second period draws need to be higher for the first bid to be rejected. In comparison, `NowWait`'s v_c^1 is the *mean* of F: $\rho = 0$, each $c_{j'}$ is drawn from F and $\tau^* = \frac{1}{|\Sigma|} \sum_{j'} c_{j'}$. Thus, `NowWait` appears better placed for average-case performance. Section 3.5.1's experimental data show the effects of different critical values on allocative efficiency for two or more periods keeping the unit-supply constraint.

3.4.5 Likelihood of being selected if in offline optimum

i's reward r_i is greater than the opportunity cost c_j on scenario j if i is in the offline optimum on scenario j. Intuitively, one expects that, over all scenarios j', r_i will *often* be greater than $c_{j'}$'s, i.e. $r_i \geq c_j$ is unlikely to be an "accident". In particular, if i is likely to depart in the next period (has low ρ), then it becomes likely that i will be selected by `NowWait` given that he is in the offline optimum on one scenario (recall Proposition 15).

I will quantify this intuition by abstracting r_i and $c_{j'}$'s. Suppose r_i is a random variable R with cdf F and pdf $f = F'$ and $c_{j'}$ is a random variable C with cdf G and pdf $g = G'$, both F and G being defined on positive values. Given that r_i is higher than a draw from C, if $|\Sigma| - 1$ more draws are taken from C (each denoted by C') then one expects that r_i will be greater than $c_{j'}$ in $(|\Sigma| - 1)\mathbb{P}[R \geq C' | R \geq C]$ draws. Note that this expectation is before taking any draws from R or C. Thus

$$\mathbb{P}[R \leq r | R \geq C] = \frac{\int_0^r f(\alpha) \int_0^\alpha g(\beta) d\beta}{\int_0^\infty f(\alpha) \int_0^\alpha g(\beta) d\beta} \text{ and } f[R = r | R \geq C] = \frac{f(r) \int_0^r g(\beta) d\beta}{\int_0^\infty f(\alpha) \int_0^\alpha g(\beta) d\beta}$$

$$\mathbb{P}[R \geq C' | R \geq C] = \int_0^\infty f[R = r | R \geq C] \int_0^r g(\beta) d\beta dr = \frac{\int_0^\infty f(r) G^2(r) dr}{\int_0^\infty f(r) G(r) dr}$$

Suppose there were one item for sale and one bidder in the current period and each of the next k periods. Then $G = F^k$: $c_{j'}$, the opportunity cost on scenario j', is the highest order statistic of k draws from F. Simple calculus shows that $\mathbb{P}[R \geq C' | R \geq C] = \frac{k+1}{2k+1} \geq \frac{1}{2}$. That is, one expects *a priori* that if r_i is in the offline optimum on scenario j then r_i is greater than $c_{j'}$ on average in $1 + (|\Sigma| - 1)\frac{k+1}{2k+1}$ scenarios, i.e. more than half of them.

3.5 Experimental evaluation

I analyze in turn the allocative efficiency, revenue and runtime of this ironing-based approach to the design of dynamic, multi-unit auctions. Unless otherwise mentioned, the **C** algorithm uses 50 scenarios (samples of possible futures) and a bidder's quantity and patience are uniform in 1..5 and his value distribution is *Exp*(0.1) times his quantity.

3.5.1 Allocative efficiency

Each of the 124 points in Figure 3.2 represents an average over at least 20 runs of `NowWait` and `IgnoDep`'s relative efficiencies on a domain where a bidder's value is his quantity times an exponentially distributed variable. In such domains, I varied supply, demand, the number of time periods, the exponential parameter λ or bidders' maximum quantity or patience. `IgnoDep` performs at least 9% worse in about a quarter of the domains, 5% worse *on average* and never better than 9% in comparison with `NowWait`.

I go on to study more closely an auction with 10 items, 5 time periods, and 2 bidders arriving in each period. The `Select` method `HazRate` is of course dominated by

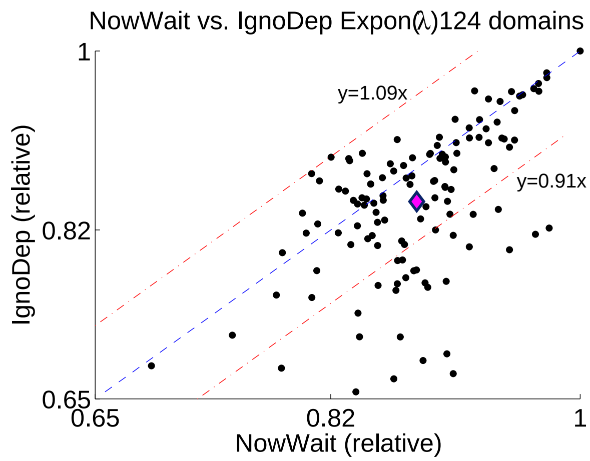

Figure 3.2: `NowWait`'s versus `IgnoDep`'s relative efficiency (offline optimum = 1) for 124 domains with exponentially distributed values with independence of value and patience. The pink diamond, at (0.882, 0.849), represents the average of all 124 points.

the more general `HROrRew`, and only results for this second method are presented. The parameters c and w for the `HROrRew` method were optimized offline, but no setting was better than that for `IgnoDep` (i.e., allowing all allocated, active agents to remain in the selected set of winning agents). Table 3.6 presents allocative efficiencies as an average over 200 trials, divided by the average offline efficiency, i.e. the value that would be achieved if all bids were available in period 1.

Whereas ironing destroys the efficiency of `OnlyDep` (as expected, because all allocations except those to maximally patient agents must be canceled), `NowWait` still yields an efficiency of 0.895 after ironing. Note that the overall performance of `NowWait` with ironing is 5% better than that of both `IgnoDep` and `HROrRew`. Standard deviations for Table 3.6 are around 0.15 for all entries but `OnlyDep` with ironing, for which it is around 0.3. For all entries except `OnlyDep` with ironing, the 95% confidence intervals have a radius of 0.02, confirming the statistical significance of my results. The `Fixed` method is less sophisticated than the other methods. It optimally (offline) tunes a per-item price p and allocates any bidder whose bid amounts to at least p per item. This method's average allocative efficiency

Ironing	`NowWait`	`OnlyDep`	`HROrRew`	`IgnoDep`	`Fixed`	Opt
No	0.915	0.952	0.860	0.860	0.815	1
Yes	0.895	0.526	0.852	0.852	0.815	1

Table 3.6: Allocative efficiency normalized to offline efficiency in a dynamic auction for 10 items when value is distributed $Exp(0.1)$ and is thus not correlated with patience.

Ironing	`NowWait`	`OnlyDep`	`HROrRew`	`IgnoDep`	Opt
No	0.944	0.941	0.881	0.881	1
Yes	0.918	0.245	0.874	0.874	1

Table 3.7: Allocative efficiency normalized to offline efficiency in a dynamic auction for 10 items when value is distributed $Exp(0.1 \cdot (d_i - a_i))$ and is thus correlated with patience.

is 0.815, which further highlights the extent of `NowWait`'s (efficiency 0.895) improvement over `IgnoDep` (efficiency 0.852).

For comparison, `NowWait`'s and `IgnoDep`'s efficiencies are very similar if per-item values are U(0,1) instead. This is again an effect of each policy's approach: `NowWait` (respectively `IgnoDep`) aims for good mean (respectively median) performance (see Sec. 3.4.4). The mean and median are equal for the uniform, but not for the exponential distribution, as used in Table 3.6.

Table 3.7 considers the effect of allowing for negative correlation between value and patience, when the exponential distribution parameter is proportional to a bidder's patience. Before ironing, `NowWait`'s allocative efficiency is slightly better than `OnlyDep`'s, that has the advantage of waiting until a bidder's departure. Ironing is now even more destructive on `OnlyDep`. This is because high-value bidders, the ones selected by the offline knapsack problem, tend to have *small* patiences, and thus are often ironed.

For all methods except `OnlyDep`, cancellations were very infrequent. There were no ironing cancellations if all bidders have unit-demand. This confirms the intuition that even though possible, instances of ironing are rather rare and caused by combinatorial peculiarities. For `NowWait`, more than half the cancellations were due to ironing in value rather than arrival or quantity.

One item, impatient bidders

I now consider the simple domain of a single unit of supply and one impatient bidder per period. This domain, considered in Section 3.4.4 for two periods, is appealing due to the availability of an optimal policy (Gilbert and Mosteller [38], henceforth `GM`), that is monotonic and thus strategyproof, providing an optimal *online* benchmark.

As bidders are impatient, there is no need for monotonicity with respect to departure. Thus, this is a setting in which the additional complexity of `NowWait` should *not* be expected to be worthwhile over-and-above the simplicity of `OnlyDep` (identical here to `IgnoDep`). Furthermore, unit-supply and impatience render all `Select` methods monotonic.

The optimal policy `GM` is defined by a sequence R_n of posted prices (where n bidders are yet to arrive). The critical value R_n also represents the expected efficiency after the n-th remaining bidder arrives. `GM` is monotonic and truthful for impatient agents because R_n is independent of the reported value of an agent (and no temporal strategies are available for impatient agents). For the Exp(λ) distribution, $R_0 = 0$ (allocate last bidder if no earlier winner) and $R_{n+1} = R_n + \frac{1}{\lambda} e^{-\lambda R_n}$ [38].

In Table 3.8 I compare the efficiency of `NowWait` and `OnlyDep` with `GM` and the offline optimum. The results for `NowWait` and `OnlyDep` are averaged over 100,000 trials. For small horizons (and hence small numbers of agents), the `NowWait` method actually outperforms `OnlyDep`. But the simpler, `OnlyDep` method does better for larger horizons and more agents. I explain this by noting that if n bidders are yet to arrive, `NowWait` (respectively `OnlyDep`) sets as critical values the mean (respectively median) of the highest order statistic of n iid $Exp(0.1)$ variables.

In summary, it is encouraging to us that the sample-based stochastic optimization

Horizon	$\mathbb{E}$[NowWait]	$\mathbb{E}$[OnlyDep]	$\mathbb{E}$[GM]	$\mathbb{E}$[Opt]
2	0.911	0.897	0.911	1
4	0.871	0.867	0.873	1
8	0.855	0.859	0.863	1
16	0.854	0.863	0.867	1
32	0.858	0.871	0.877	1

Table 3.8: Relative (offline optimum=1) allocative efficiency in a unit-supply domain with impatient bidders. `GM` is the maximally-efficient, monotonic policy.

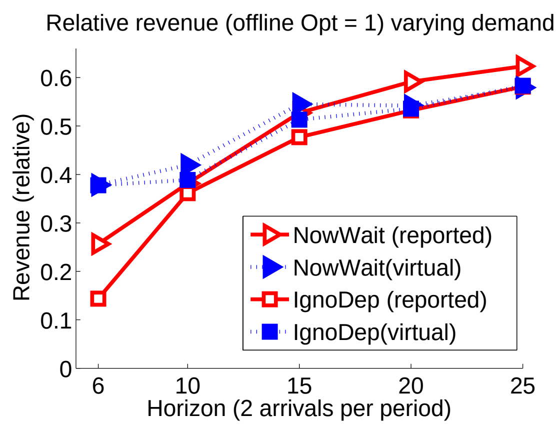

Figure 3.3: Revenue normalized to offline efficiency in a dynamic auction for 30 items, varying the number of time periods (and thus the level of competition). Examines the effect of adopting virtual valuations.

methods can come within 97.8% of the value of the optimal online policy in this environment (this is the relative performance of **C** with `NowWait` at a horizon of 32), while being flexible and general enough to extend to multi-unit demand environments.

3.5.2 Boosting revenue

I also performed experiments with virtual valuations in place of agent valuations in order to test the effectiveness of this approach to boosting revenue.

Figure 3.3 plots the revenue with and without virtual valuations. The revenue metric is normalized with respect to the total value from the efficient offline allocation. In this auction, there are 30 items available and vary the number of periods from 6 to

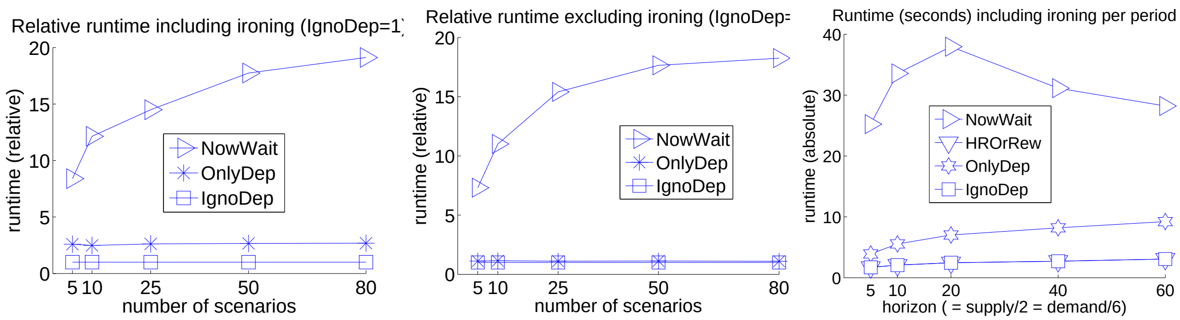

Figure 3.4: Runtimes in a dynamic auction. The left-hand side, respectively center, experiment shows relative runtime including, respectively excluding, ironing for an auction with 20 items and 10 time periods. The x-axis varies the number of scenarios. In the right-hand side experiment, showing seconds of runtime per time period, the expected demand:supply ratio is kept constant at 3 by adjusting them together with the number of time periods, as shown on the x-axis.

25, with 2 bidders arriving per period. As the number of periods increases the competition increases: the expected demand ranges from low-demand (36) to high-demand (150), i.e. the demand:supply ratio trends from 1.2x to 5x. Virtual valuations have a significant positive effect on revenue for low demand environments. For example, in the case of 6 periods (and thus a low demand:supply ratio of 1.2x), adopting virtual valuations provides a boost of as much as 169% for `IgnoDep` and 49% for `NowWait`. On the other hand, virtual valuations can also be detrimental to revenue properties for high demand environments (a demand:supply ratio of 4 or more). Thus, it would be important for a designer to understand the type of environment before adopting virtual valuations. The revenue statistics of `NowWait` generally dominate those of `IgnoDep`, both with and without ironing. `OnlyDep`'s revenue (not shown in Figure 3.3) is always significantly below that of other methods due to extensive cancellations by ironing.

3.5.3 Computational scalability

All experiments were run on a CentOS 8-node Pentium 4 at 3GHz cluster with 512 MB of RAM. Figure 3.4 summarizes the results, averaged over 50 trials, for a dynamic auction with 2 bidders arriving each period. In the set of experiments for the left-hand side and the center plot in Figure 3.4, there are 20 items and 10 periods and I increase the number of scenarios sampled within **C**. In the experiments for the right-hand side plot in Figure 3.4, I increase the supply of items while holding the expected demand:supply ratio constant at 3, by increasing the time horizon and thus the total expected demand as the supply increases. In addition to looking at the scalability of the system, I am also interested in the overhead that is imposed by the need to perform computational ironing.

All three plots show that `NowWait`'s computational overhead when compared to the other `Select` methods is limited and reasonable. For example, as the number of scenarios increases, `NowWait`'s overhead grows sub-linearly. This is despite `NowWait`'s theoretical quadratic (as opposed to linear) dependence on the number of scenarios (for all $j, j' \in \Sigma$, the $c_{j'}$ costs must be computed for each scenario (σ^j) by solving two offline optimization problems). I believe that this is due in part to some `NowWait`-specific improvements that I have made, for example caching the offline optimization results for the $c_{j'}$ costs (see Sec. 3.4.1). For most experimental settings considered, on average, **C** and ironing take around 15 times more than **C** alone for all methods except `OnlyDep`, for which the overhead is around 40x.[8]

The right hand side experiment in Figure 3.4 measures my methods' *absolute, per period* runtime (in seconds) of CONSENSUS and ironing, as the relevant characteristics (number of time periods, supply and expected demand) of the domain are scaled proportionally. It is quite encouraging, though perhaps surprising, that `NowWait`'s per-period runtime is

[8] `OnlyDep`'s higher overhead of ironing can be noticed in Figure 3.4: the ratio of the runtimes of `OnlyDep` and `IgnoDep` is close to 1 without ironing, but about 2.5 with ironing. There exists (a slightly smaller) difference in ratios even if a, d and q are kept constant in $\text{isIroned}_{A,D,Q}$ (see Algorithm 2). The reason for this discrepancy is revealed by recalling that isIroned_R (see Algorithm 3) computes breakpoints at all times between a bidder's arrival and time of allocation. `OnlyDep` allocates bidders at their departure, which is significantly later on average than the other methods. Consequently, breakpoints for significantly more time periods are computed in isIroned_R for `OnlyDep`.

decreasing for the longer horizons.

In closing the discussion on `NowWait`'s scalability, note that all `NowWait`'s theoretical properties would still hold if any agent-independent opportunity costs, not necessarily the set of $c_{j'}$ used by $\text{now}_i^t()$ and $\text{wait}_i^t()$, were plugged into `NowWait`. More sophisticated, agent-dependent, estimates may also render output-ironing feasible. Alternate estimates would be particularly appealing if they led to even better scalability, while preserving the quality of optimization.

Summary

I presented the first application of stochastic optimization to dynamic, incentive-compatible multi-unit auctions with patient bidders that demand multiple units of an item. Method `NowWait` is used to modify the CONSENSUS algorithm [84] and evaluate opportunity costs when deciding whether to retain a vote to allocate to a particular agent given a particular scenario. Self correction by output ironing yields truthfulness, and one can aim for either good efficiency or revenue. The results show excellent efficiency and scalability, with a sub-linear computational overhead of `NowWait` with respect to CONSENSUS, and also demonstrate the effect of revenue boosting via the use of virtual valuations.

Algorithm 3 isIroned$_R(\theta_i, t_i^*(\theta_i), (S^t, A^t)_{a_i \le t \le t_i^*}, \Sigma)$: verification of value-monotonicity for an agent i with type θ_i allocated in period t_i^*. The set of breakpoints $B = \{\langle t^\beta, \sigma^{j_\beta}, r_\beta, \alpha_{\mathtt{s}}^{<}(\beta), \alpha_{\mathtt{s}}^{\ge}(\beta)\rangle\}_\beta$ track the changes in scenario votes. As i's value increases past r_β, the vote on scenario σ^{j_β} at time t^β changes from $\alpha_{\mathtt{s}}^{<}(\beta)$ to $\alpha_{\mathtt{s}}^{\ge}(\beta)$. If changing only the vote on scenario j_β causes the $\mathbf{C}$ decision to flip, then π must be simulated to allow checking that i's allocation time (stored by t_i) for the higher (slightly above r_β) reward is no later than for his lower (below r_β) reward. Finally, as the $\mathbf{C}$ decision changes at t^β, breakpoints between periods $t^\beta + 1$ to t_i (where $t^\beta \le t_i$) must be updated. The counterfactual sets $\{\vec{S}, \vec{A}\} = (S^t, A^t)_{a_i \le t \le t_i^*}$ are maintained, being initialized to the actual ones determined by π for i's reported value, which is lower than all r_β's.

$B := \bigcup_{t=a_i}^{t_i^*} \mathrm{BrkPts}^R(\theta_i, t, \vec{S}, \vec{A}, \Sigma)|_{r \ge r_i}$
$t_i := t_i^*$
while $B \ne \emptyset$ **do**
  Let $\beta := \langle t^\beta, \sigma^{j_\beta}, r_\beta, \alpha_{\mathtt{s}}^{<}, \alpha_{\mathtt{s}}^{\ge}\rangle$ such that $r_\beta \le r_b \,\forall\, b \in B$.
  if at t^β, $\mathbf{C}(\text{votes}(\Sigma^{\ne j_\beta}), \alpha_{\mathtt{s}}^{<}) \ne \mathbf{C}(\text{votes}(\Sigma^{\ne j_\beta}), \alpha_{\mathtt{s}}^{\ge})$ **then**
    `// Simulate` π `until time` $\min\{t_i,$ `time` i `wins`$\}$
    increase i's value to slightly over r_β
    $\breve{t} := t^\beta$
    repeat
      $\pi^{\breve{t}} := \pi^{\breve{t}}(A^{\breve{t}}, S^{\breve{t}})$
      update active agents: $A^{\breve{t}+1} := \theta^{\breve{t}+1} \cup A^{\breve{t}}|_{d \ge \breve{t}+1} \setminus \pi^{\breve{t}}$
      update supply: $S^{\breve{t}+1} := S^{\breve{t}} - \#(\pi^{\breve{t}})$
      $\breve{t} := \breve{t} + 1$
    until $i \sqsubset \pi^{\breve{t}-1}$ or $\breve{t} > t_i$
    if i has not won **then**
      return `true //` i `ironed`
    else
      $t_i := \breve{t} - 1$ `// allocation time for new value` r_β
      $B := \big(B|_{t \le t^\beta} \setminus \{\beta\}\big) \cup \bigcup_{t = t^\beta+1..t_i} \mathrm{BrkPts}^R(\theta_i, t, \vec{S}, \vec{A}, \Sigma)|_{r \ge r_\beta}$
    end if
  end if
end while
return `false //` i `not ironed`

Chapter 4

Cancellations for Impatient Buyers

Abstract.

Reservations, also known as advance-booking, are a standard selling practice. Sellers with a varying inventory and audience (for example, in sponsored search) are interested in automatic, simple yet powerful market mechanisms for the allocation of reservations. I introduce a simple model for reservations, in which impatient private-value unit-demand bidders arrive sequentially. The seller can cancel at any time an earlier reservation, resulting in a utility loss to the reservation holder of a fraction of his value.

My main result is an online mechanism $\mathtt{M}_\alpha(\gamma)$ with many desirable game-theoretic and optimization properties. Winners have an incentive to be honest and bidding one's true value dominates any lower bid. $\mathtt{M}_\alpha(\gamma)$'s revenue is within a constant fraction of the a posteriori revenue of the Vickrey-Clarke-Groves mechanism. $\mathtt{M}_\alpha(\gamma)$'s efficiency is within a constant fraction of the a posteriori optimal solution. If efficiency also takes into account the utility losses of bidders whose reservation was canceled, $\mathtt{M}_\alpha(\gamma)$ can match an upper bound on the competitive ratio of any deterministic online algorithm. Several extensions are considered, including matroids and bidders with value for more than one item.

4.1 Introduction

In a system for booking items at a specific time in the future, buyers are typically interested in a reasonable guarantee in the present. For example, in the airline industry, a ticket is very likely, but not guaranteed, to allow its owner to board the plane. In particular, airlines often use *overbooking* since statistically a few passengers will miss their flight. In the unlikely scenario that more passengers than seats arrive, the airline typically offers compensation to passengers who will not board the plane (volunteers or randomly chosen).

My motivation arises from automatic systems enabling such reservations. In particular, my focus is on large systems that have to manage items in many different environments. One such example is sponsored search, where advertisers place advertisements (*ads*, henceforth) in response to users' web search queries, or at predetermined slots on publishers' web pages. These web pages differ wildly in their traffic, targeting and effectiveness and the overall inventory levels are massive. Not all publishers can estimate their inventory accurately: traffic to websites responds, among others, to time-dependent events. Most web publishers are not able to estimate accurately a price for an ad slot, or provide sales agents to negotiate terms and would like automatic methods to price ad slots. Thus, what is desirable is a simple, effective, automatic, online[1] market-based mechanism to enable advance booking over such varied, massive inventory.

Inspired by these considerations, I investigate worst-case mechanisms for advance-booking (reservations). My contribution is to propose a simple model, to design a suitable mechanism and to analyze its properties. In more detail, my contributions are as follows.

Model. I propose the following simple model for advance booking. An auction starts at time 1; the seller has a set of *items* for sale that will be available at time T. Bidder i may

[1] I use the word *online* as in *online algorithm*—i.e., the input arrives over time, and the algorithm makes sequential decisions —I do not mean "on the Internet."

arrive at some time $a_i < T$, having a value $v(i)$ for exactly one item out of a subset of items $L(i)$. Upon his arrival, i places a bid $w(i)$ (which results in $L(i)$ becoming known to the seller) and requests an immediate response. Bidder i is either accepted or rejected. If accepted, the bidder faces a second strategic decision: whether to *invest* or not in the item. If accepted (regardless of the investment decision), he may be removed (*bumped*) later by the seller. If the bidder did not invest, being bumped or winning an item have no effect on his utility. If however he invested and is bumped, he incurs a loss of an α fraction of his value and may be compensated with a *bump payment*. At time T, each accepted bidder i that has not been bumped obtains one item from $L(i)$ (his items of interest). The model is presented formally in Section 4.2.

This model lets the seller accept a reservation at time t for an item available at a later time T, and lets the buyer get a reasonable guarantee. However, crucially, it lets the seller *cancel* the reservation at a later time. Cancellation is necessary for the seller to take advantage of a spike in demand and rising prices for an item and not be forced to sell the item below the market because of an *a priori* contract. In addition, in a pragmatic sense, cancellation is crucial: for example, a seller might overestimate her inventory for a later date and accept reservations, but as time progresses, her inventory may be smaller, and the seller will not be able to honor all the accepted reservations from the past. Finally, cancellations are very much part of the business with advance bookings, both within advertising and beyond such as in airline bookings. At the same time, it comes at a cost, which is the bumped bidders' utility loss. The seller benefits from the reduction in uncertainty, and pays for this via bump payments.

Mechanism. I present an efficiently implementable mechanism $\mathtt{M}_\alpha(\gamma)$ for determining who is accepted, who is bumped and also the prices and bump payments. The parameter γ represents how much higher a new bid has to be in order to bump an older bid. A bumped bidder will be paid an α fraction of his bid, making up for his utility loss.

Properties of My Mechanism. I show a number of important strategic as well as efficiency- and revenue-related properties of $\mathtt{M}_\alpha(\gamma)$. First, the strategic properties:

- $\mathtt{M}_\alpha(\gamma)$ is individually rational and winners have an incentive to bid truthfully while losers should bid at least their true value.
- I study *speculators*, that is, bidders who do not invest. Speculators are only interested in earning the bump payments as opposed to winning items. I show several game theoretic properties about the behavior of the speculators, including bounding their overall profit.

Next, optimization properties:

- With respect to the bids received, the *efficiency* (value of assignment) of $\mathtt{M}_\alpha(\gamma)$ is at least a constant factor (depending on γ and α) of the offline optimum. Under mild player rationality assumptions, I also show that my mechanism is competitive with respect to the optimum offline efficiency on bidders' true *values*.
- I prove similar bounds under the notion of *effective efficiency* which interprets social welfare as the sum of the winners' bids minus bumped bidders' losses. For suitable $\gamma(\alpha)$, $\mathtt{M}_\alpha(\gamma)$ is *optimal*: its effective efficiency matches a numerically obtained upper bound on the effective efficiency of *any* deterministic algorithm.
- Under a very slight rationality assumption, the *revenue* of $\mathtt{M}_\alpha(\gamma)$ is at least a constant factor (dependent on γ and α) of that of the VCG mechanism on all received bids.

To the best of my knowledge my results are the first about mechanisms with strong game-theoretic properties for advance booking (more generally, online weighted bipartite matching) with a costly cancellation feature. I make no assumptions on the arrival order of the bidders or on their values. All results extend to a more general setting where the items for sale are elements of a matroid; I discuss this in Section 4.6.6.

There are specific examples of systems that implement advance booking with cancellations, not necessarily through an automatic mechanism. For example, this is common in the airline industry, where tickets may be booked ahead of time, and customers may be bumped later for a payment. In the airline case, the inventory is mostly fixed, sophisticated models are used to calculate prices over time, and often negotiations are involved in establishing the payment for bumping, just prior to boarding (time T). In some cases, bump payments may even be larger than the original bid (price) of the customer. Likewise, in offline media such as TV or Radio, advance prices are negotiated by humans, and often if the publisher does not respect the reservation due to inventory crunch, a payment is *a posteriori* arranged including possibly a better ad slot in the future. These methods are not immediately applicable to the auction-driven automatic setting that I consider.

From a technical point of view, one can view my model as an online weighted bipartite matching problem (or more generally, an online maximum weighted independent set problem in a matroid). On one side there are items known ahead of time. The other side contains buyers whose bids (weighted nodes) arrive online. My goal is to find a "good" weighted matching in the eventual graph. Each time a buyer appears, he must be retained or discarded; retaining him may lead to discarding a previously retained bidder. My mechanism builds on an one-pass matching algorithm [62] to determine a suitable bump payment and prices.

I have initiated the study of mechanisms for advance reservations with cancellations. A number of technical problems remain open, within my model as well as in its extensions, which I describe later for future study.

4.1.1 Related work

Several papers consider similar *offline* settings with no cancellations. Bikhchandani et al. [12] present an ascending, efficient, truthful in equilibrium auction for selling elements of a matroid to patient bidders. Cary et al. [16] show that a random sampling profit extraction mechanism approximates a VCG-based target profit in an offline procurement setting on a matroid.

Feige et al. [32] study an *offline* weighted bipartite matching problem where the seller can partially satisfy a bidder's request at the cost of paying a proportional penalty. They show that approximating the optimal solution with respect to effective efficiency (see Section 4.2) within any constant factor is NP-hard, but they provide a bi-criterion approximation result for an adaptive greedy algorithm.

There has been extensive work in the field of revenue management for advance sale of goods (with or without cancellations), but only under a probabilistic distribution of bidders' values or arrivals [83]. In particular, Gallien and Gupta [34] exhibit symmetric Bayes-Nash equilibria in online auctions with buyout prices.

The results in this chapter are in contrast to the impossibility, when no cancellations are allowed, of achieving a constant factor approximation with respect to the bids received when there is no prior information on the number of bidders, even if a prior distribution on bidders' values is known [45].

In a worst-case approach (like the one here), any nontrivial online result must make additional assumptions; in my case, I overcome these impossibilities by allowing cancellations. In contrast, in secretary problems, bids may be arbitrary but their order is assumed to be uniformly random (cannot be specified by an adversary): Dimitrov and Plaxton [31] provide an algorithm with a constant efficiency competitive ratio. Generalizing their setting to matroids, Babaioff et al. [6] provide a $\log r$-competitive algorithm where r is the rank of the matroid. Both these algorithms observe half of the input and then set a threshold price. A different assumption is that of bounded values: Lavi and Nisan [55] show that a simple online posted-price auction based on exponential scaling is optimal among online auctions for identical goods without cancellations.

Cole et al. [22] study a model where the patience levels of items and bidders are almost opposite to my model's: items are instantaneous (one per period), but bidders are partially patient, with arrivals and departures. They present an optimal constant-competitive *prompt* mechanism for the case where bidders' arrival-departure intervals are known. In contrast, they obtain lower bounds on the competitive ratio of mechanisms if both value and departure time are private information.

Lu et al. [61] study online worst-case revenue-optimal auctions with instantaneous bidders. They provide an essentially optimal auction whose revenue is within a $\ln \mathrm{Opt} \cdot \ln\ln \mathrm{Opt} \cdot \ldots \cdot (\ln^{(k)} \mathrm{Opt})^{1+\varepsilon}$ factor of the optimum Opt, for any positive integer k and any $\varepsilon > 0$. Typical revenue-optimality results claim bounds with respect to weaker goals, such as the second-price or the optimal revenue from selling at least two units.

In Hajiaghayi et al. [44]'s model, unit-demand bidders have arrivals and departures, but their values come from a (possibly) unknown distribution. The authors present a constant-competitive mechanism that is incentive-compatible with respect to both temporal and value misreports.

Biyalogorsky and Gerstner [13] study *contingent pricing* for a seller offering for an item over two time periods. There is an impatient buyer L in period 1 and possibly a buyer H with higher value in period 2. Contingent pricing is the practice of offering a discount, respectively a consolation reward, to the impatient buyer L if the higher-value buyer H arrives, respectively does not arrive, in period 2. The structure of optimal contingent pricing is dependent on L's attitude towards risk (of not being allocated the item).

Independently and concurrently, Babaioff et al. [8] study the same problem as this chapter, but from an algorithmic perspective only, leaving incentives and revenue considerations aside. The algorithms and efficiency results coincide. They focus on effective efficiency, for which they analytically prove an upper bound on any deterministic algorithm's competitive ratio (I only present numerical results in Figure 4.1 strongly suggesting this bound). Unlike us, they go on to study costly cancellations ("buyback") in knapsack problems. They provide an algorithm similar to $\mathtt{M}_\alpha(\gamma)$ and prove a bi-criterion approximation result, an informative bound since they also prove that no deterministic algorithm has a constant competitive ratio.

4.2 Auction model

I define an *online reservation auction* as follows. There is a seller who has a finite set of non-identical *items*, which will be allocated at some future time $T + 1$. The seller runs a continuous, online auction beginning at time 0, and ending at time T.

Each bidder i arrives online, at a unique time $a_i \in [0, T]$ and he reports a *choice set* $L(i)$ of items he is interested in, as well as a *bid* (positive amount) $w(i)$, demanding an immediate response (i is not allowed to bid again later). He is instantly *accepted* (i.e. promised an item from $L(i)$), or *rejected*. However, at any point between time a_i and T, the seller may choose to *bump* an accepted bidder i, in which case a *bump payment* $\hat{p}_i$ is given to the bumped bidder. Any rejection, at arrival or by being bumped, is definitive. At time T, there must be a matching of items to accepted bidders that have not been bumped such that each such bidder i receives one item from his choice set $L(i)$; each such i is then charged a price p_i.

A *mechanism* for the advance-booking problem defines the actions of the seller:

whether to accept/reject incoming bidders, when to bump accepted bidders, and how to set bump payments and prices.

4.2.1 Bidder model

I assume a *private value* model for the bidders contingent on a private investment decision. If accepted, a bidder can privately choose whether to "invest" or not. i obtains positive value from being allocated or incurs a utility loss (defined below) when bumped only if he chose to invest.

Suppose that bidder i chooses to invest. He then has a private value $v(i) \geq 0$ for being allocated (at time $T + 1$) any single item from his choice set $L(i)$. The bidder may not report this value as his bid. Additionally, the bidder's cost for being bumped is modeled as a negative value $-\alpha v(i)$, that is, an α fraction of his value for being allocated. I will require[2] that any mechanism pays back $\alpha w(i)$ to a bumped bidder i, making up for his utility loss when bumped. The parameter $0 \leq \alpha < 1$ modeling the negative bump utility will play a central role in my mechanism and analysis. I formally model bidder i's *utility* as quasilinear in money:

$$\textit{utility}(i, \mathrm{I}) = \xi(\mathrm{I}) \cdot v(i) - x(i), \quad \text{where} \tag{4.1}$$

- $\xi(\mathrm{I}_0) = 0$ for any decision; $\xi(\mathrm{I}_1) = 0$ if i is rejected, $\xi(\mathrm{I}_1) = 1$ if i is accepted and granted an item from $L(i)$, and $\xi(\mathrm{I}_1) = -\alpha$ if i is bumped;
- $x(i)$ is i's *money transfer* to the seller: $x(i) = 0$ if i is rejected, $x(i) = p_i \geq 0$ if i is accepted and allocated, and $x(i) = -\alpha w(i) \leq 0$ if i is bumped.

For a mechanism run on bids $\mathbf{w}$, I will denote by $S = S(\mathbf{w})$ the set of *survivors* (bidders still accepted after time T) and by $R = R(\mathbf{w})$ the set of bumped bidders.

[2] I impose this constraint primarily to ensure that honest bidders have non-negative utility. See Example 8 for further motivation.

4.2.2 Efficiency and revenue.

How to measure the quality of an outcome? The *efficiency* (or *social welfare*) of an auction is the total value derived by the bidders participating in the auction. Usually in a mechanism design setting, efficiency is the sum of the valuations of the bidders who were allocated items:

$$\textit{efficiency} = \sum\nolimits_{i \in S} v(i).$$

However there is another interpretation in my model since the bidders lose value if they are accepted then bumped; thus I will also consider the notion of *effective efficiency*:

$$\textit{effective efficiency} = \sum\nolimits_{i \in S} v(i) - \sum\nolimits_{i \in R} \alpha v(i).$$

A mechanism's *revenue* is its total monetary gain/loss:

$$\textit{revenue} = \sum\nolimits_{i \in S} p_i - \sum\nolimits_{i \in R} \alpha w(i)$$

I would like a mechanism that scores favorably in all of these metrics on all instances of the auction, under hopefully mild conditions on strategic behavior. Following the logic of *competitive analysis*, I will compare my mechanism to the standard offline solution: the VCG mechanism [52]. Mapped to my setting (but offline), this amounts to finding a maximum matching of bidders to items, and charging prices that induce truthfulness.

Note that it is essential to the novelty of this model that the bidders derive negative utility from having their allocation promise revoked. Indeed if $\alpha = 0$, one could simply accept all bidders as they arrive, and then after time T run the VCG mechanism (giving no bump payments) which makes being truthful a dominant strategy.

4.3 Main results

In this section I will state my main results (without proof) and highlight the significance of each. In the next section, I will define my mechanism $\mathtt{M}_\alpha(\gamma)$. The mechanism is parametrized both by the model parameter α as well as an additional parameter $\gamma > 0$ that can be set arbitrarily as long as $0 < \alpha < \frac{\gamma}{1+\gamma}$. I will state my main results in terms of these two parameters.

The algorithmic game-theoretic perspective requires reasoning about the strategic behavior of bidders in order to motivate the results' preconditions. One basic property that any reasonable mechanism must have is that it is *individually rational*, which simply means that participating in the auction is always a rational thing to do (i.e. participating is never worse than not participating). A common relaxation of this definition, that I will use, is the following: if a bidder bids his true value (sets $w_i = v_i$) and invests, then his utility is always non-negative.

Another desirable property of an auction mechanism is for it to be *truthful*, which means that the optimal strategy for participating bidders is always to report their true value. In my model it is also desirable that any accepted bidder invests. Unfortunately with bump payments (which I just argued were necessary) I cannot hope to have a truthful mechanism since anyone with no interest in any allocation (i.e., $v_i = 0$) can bid without investing hoping to get a bump payment. So, given that I cannot assume bidders will be honest, the natural thing to do is analyze the efficiency and revenue of the mechanism in a Nash or other form of equilibrium. Unfortunately, a (pure-strategy) Nash equilibrium does not always exist, as I will argue in Section 4.6.2. However I can still argue that my mechanism has some strong incentive properties.

Recall the following standard game-theoretic terminology from Chapter 2: a bid *dominates* another bid if it is at least as good a strategy given *any* bids by other players; a *best-response* is the best possible bid given a particular set of others' bids.

Theorem 11. [Basic Incentive Properties] *In $\mathtt{M}_\alpha(\gamma)$,*

1. *A truthful, investing, bidder cannot run a loss (i.e. $\mathtt{M}_\alpha(\gamma)$ is individually rational),*
2. *For a fixed investment decision, truthful bidding dominates any lower bid,*
3. *If truthful and investing, any survivor is best-responding, and*
4. *Bidding truthfully is a best-response unless a higher bump payment can be achieved with a higher bid.*

This theorem (proved in Section 4.6.1) establishes individual rationality, but more importantly it rules out the possibility that the bids will be lower than the values. To the best of my knowledge, this is a novel form of incentive compatibility: while it does not make truthfulness a dominant strategy, it ensures that competition is no less than if every bidder were truthful. Furthermore, it highlights truthfulness as a simple viable strategy from a practical point of view, or for unsophisticated bidders. The only reason for not bidding truthfully is the prospect of a higher bump payment.

I can now state the efficiency and revenue bounds for $\mathtt{M}_\alpha(\gamma)$ assuming in particular that no bidder underbids, which one would expect by Theorem 11. The only other assumption I make for my efficiency bounds is that the set of bidders does not jointly incur a loss, which is quite mild an assumption. Indeed, truthful, investing, bidders never incur a loss and if other bidders incur a loss and can collude then they would be better off not participating at all.

For any vector $\mathbf{w} = (w(1), \ldots, w(n))$ of bids, I let $\mathrm{Opt}[\mathbf{w}]$ denote the weight of the optimal matching. Note that $\mathrm{Opt}[\mathbf{v}]$ then gives the optimal efficiency and effective efficiency of an offline mechanism, achieved by VCG. On bids $\mathbf{w}$ I denote the VCG revenue by $\mathtt{REV}_{vcg}[\mathbf{w}]$.

Theorem 12. **[Efficiency]** *Let* $\mathbf{w}$ *be a set of bids such that all bidders bid at least their true value and bidders' total utility is non-negative. Then* $\mathtt{M}_\alpha(\gamma)$ *has*

$$\begin{aligned}\text{efficiency} &\geq \frac{\frac{1}{1+\gamma} - \frac{\alpha}{\gamma}}{(\frac{2}{1+\gamma} - \frac{\alpha}{\gamma})(1+\gamma)} \cdot \mathrm{Opt}[\mathbf{v}] \quad \textit{and} \\ \text{effective efficiency} &\geq \frac{\frac{1}{1+\gamma} - \frac{\alpha}{\gamma}}{2-\alpha} \cdot \mathrm{Opt}[\mathbf{v}].\end{aligned}$$

Theorem 13. **[Revenue]** *Let* $\mathbf{w}$ *be a set of bids such that all bidders bid at least their true value. Then* $\mathtt{M}_\alpha(\gamma)$ *has*

$$\text{revenue} \geq (\tfrac{1}{1+\gamma} - \tfrac{\alpha}{\gamma}) \cdot \mathtt{REV}_{vcg}[\mathbf{w}] \geq (\tfrac{1}{1+\gamma} - \tfrac{\alpha}{\gamma}) \cdot \mathtt{REV}_{vcg}[\mathbf{v}]$$

Note that a limit on manipulation is needed for a lower bound on true efficiency: if low value bidders bid really high, being allocated all the items and preventing the rightful winners from being allocated, the true efficiency of the mechanism is very low. I further discuss manipulations in Section 4.6.2, where I also give additional results on speculator strategies.

Efficiency bounds that leave incentives aside (Theorem 14 and Corollary 3 in Section 4.5) are tighter than the bounds in Theorem 12, and can be obtained more easily.

In Section 4.7.4 I give an *upper bound* on the effective efficiency (in terms of bids) of any deterministic algorithm, for which my allocation algorithm is tight (only when $\alpha < 0.618$ and for a certain γ that depends on α).

For example, suppose $\alpha = \frac{1}{4}$. Then for $\gamma = 1$, the constants become: $\frac{1}{6}$-competitive on efficiency, $\frac{1}{7}$-competitive on effective efficiency, and $\frac{1}{4}$-competitive on revenue. For γ specified in Section 4.7.4 (about 0.809), the mechanism is about 0.382-competitive on effective efficiency.

I defer the proofs of Theorems 12 and 13 to section 4.7. As a revenue benchmark, I use the offline VCG mechanism that charges each bidder his externality on the other bidders. Theorem 13's proof uses Lemma 2 and the following facts: a winning bidder's VCG payment is a losing bid and the VCG revenue can only increase if some bids are increased.

4.4 $\mathtt{M}_\alpha(\gamma)$: an online mechanism

I present my advance-booking online mechanism $\mathtt{M}_\alpha(\gamma)$ in this section. The allocation algorithm follows the *Find-Weighted-Matching* algorithm in [62][3], that uses an unconstrained *improvement factor* $\gamma > 0$. I require $\alpha < \frac{\gamma}{1+\gamma}$ i.e. $\gamma \in (\frac{\alpha}{1-\alpha}, \infty)$ (recall that $0 \leq \alpha < 1$) for non-negative lower bounds in Theorems 12 and 13.

My mechanism $\mathtt{M}_\alpha(\gamma)$ (given formally in Algorithm 4) maintains a set of accepted bidders for which there exists a matching of bidders to items. For each new arriving bidder i bidding $w(i)$, $\mathtt{M}_\alpha(\gamma)$ adds i to the current matching if it can do so without bumping a currently accepted bidder. Otherwise, $\mathtt{M}_\alpha(\gamma)$ looks for some bidder j in the accepted set with $w(j) < \frac{w(i)}{1+\gamma}$ such that replacing j by i maintains the existence of a matching. If such a bidder exists, the mechanism accepts i and bumps j^*, the lowest weight such bidder, who is paid the bump payment $\alpha w(j^*)$. After time T, when all bidders have been processed, the accepted bidders become the survivors. Each survivor is allocated an item from his choice set using an arbitrary matching and makes a payment that I define below.

Eq. (4.2) below establishes a survivor's payment to the seller, and requires the following definitions.

Definition 29. *Let i be a bidder and fix the bids of all other bidders. Let $w^{\mathrm{ac}}(i)$ (i's*

[3]Unlike in [62], a bidder i's value is the same for any item (vertices as opposed to edges are weighted). My mechanism may then change the item i is currently assigned to at various stages in the algorithm.

Algorithm 4 $\texttt{M}_\alpha(\gamma)$: Allocation algorithm. A new bidder is accepted if he improves by at least a γ factor over his lowest-bidding indirect competitor. A bumped bidder receives a bump payment making up for his utility loss.

Let $A_0 := \emptyset$.

for each arriving bidder $i \geq 1$ bidding $w(i)$ **do**

  if $A_{i-1} \cup \{i\}$ can be matched **then**

    grant i a reservation: $A_i := A_{i-1} \cup \{i\}$.

  else let j^* be the lowest-bidding $j \in A_{i-1}$ such that $A_{i-1} \cup \{i\} \setminus \{j\}$ can be matched

    if $w(i) < (1+\gamma)w(j^*)$ **then** reject i: $A_i := A_{i-1}$.

    else cancel j^*'s reservation and pay him $\alpha w(j^*)$

      grant i a reservation: $A_i := A_{i-1} \cup \{i\} \setminus \{j^*\}$.

end for

Each bidder i in $S = A_n$ (i.e. survivors) is allocated an item from $L(i)$ and charged as in Eq. (4.2).

acceptance weight*) be the infimum of all bids that i can make such that i is accepted given i's arrival a_i and i's choice set N_i. Similarly, let $w^{\text{sv}}(i) \geq w^{\text{ac}}(i)$ (i's* survival weight*) be the infimum of all bids that i can make such that i is accepted* and survives until time T (the end) *given a_i and N_i.*

Note that $w^{\text{sv}}(i)$ always exists since it suffices to bid $(1+\gamma)\max_{j\neq i} w(j)$. Also, $w^{\text{ac}}(i)$ and $w^{\text{sv}}(i)$ are independent of i's actual bid, but may depend on the time i arrives and on the other bidders' bids or arrivals.

I will now introduce the prices charged by $\texttt{M}_\alpha(\gamma)$. If i is a survivor, i's price p_i is set as follows:

$$p_i = \begin{cases} w^{\mathrm{sv}}(i)(1-\alpha) & \text{if } w^{\mathrm{ac}}(i) < w^{\mathrm{sv}}(i). \\ w^{\mathrm{sv}}(i) & \text{if } w^{\mathrm{ac}}(i) = w^{\mathrm{sv}}(i). \end{cases} \tag{4.2}$$

These prices are designed with Theorem 11's conditions in mind, as will be clear in its proof.

I present now an example run of $\mathtt{M}_\alpha(\gamma)$. $w^{\mathrm{sv}}(i)$ at some time step is i's survival weight if $\mathtt{M}_\alpha(\gamma)$ stopped then.

Example 6 (a particular instance of $\mathtt{M}_\alpha(\gamma)$)**.** *Suppose* $\alpha < \frac{0.5}{0.5+1}$ *and let* $\gamma = 0.5$. *Consider two items* I_a, I_b *and bidders* $B_{1..4}$ *(*B_i *arrives at time* t_i; $T = t_4$*):* B_1 *bids* 6 *on* $L(1) = \{I_a, I_b\}$, B_2 *bids* 4.4 *on* $L(2) = \{I_b\}$, B_3 *bids* 10 *on* $L(3) = \{I_a\}$ *and* B_4 *bids* 7.5 *on* $L(4) = \{I_b\}$. $\mathtt{M}_\alpha(\gamma)$ *accepts* B_1 *at* t_1, *accepts* B_2 *at* t_2, *accepts* B_3 *and bumps* B_2 *at* t_3 *and then rejects* B_4 *at* t_4. *Both* $w^{\mathrm{ac}}(1)$ *and* $w^{\mathrm{ac}}(2)$ *are* 0; $w^{\mathrm{ac}}(3) = 1.5 \cdot 4.4 = 6.6$ *(to bump* B_2*) and* $w^{\mathrm{ac}}(4) = 1.5 \cdot 6 = 9$ *(to bump* B_1*);* $w^{\mathrm{sv}}(1) = \frac{7.5}{1.5} = 5$ *(to prevent being bumped by* B_4*),* $w^{\mathrm{sv}}(2) = 6$ *(to prevent being bumped by* B_3 *and* B_4*),* $w^{\mathrm{sv}}(3) = 6.6$, *and* $w^{\mathrm{sv}}(4) = w^{\mathrm{ac}}(4) = 9$. B_1 *and* B_3 *survive:* B_1 *pays* $(1-\alpha)w^{\mathrm{sv}}(1)$ *since* $w^{\mathrm{ac}}(1) < w^{\mathrm{sv}}(1)$ *and* B_3 *pays* $w^{\mathrm{sv}}(3)$ *since* $w^{\mathrm{ac}}(3) = w^{\mathrm{sv}}(3)$.

4.5 Algorithmic properties

For my incentive-aware bounds from Section 4.3, $\mathtt{M}_\alpha(\gamma)$ must find a good matching given the declared bids $\mathbf{w}$, regardless of those bidders' true values. This is a pure online-algorithms question (i.e., no game theory), which I treat in this section.

Recall that Opt[$\mathbf{w}$] denotes the optimal offline matching on the bids $\mathbf{w}$. For bids $\mathbf{w}$ and a set of bidders B, I let $w(B) = \sum_{i \in B} w(i)$ and $w^{\mathrm{sv}}(B) = \sum_{i \in B} w^{\mathrm{sv}}(i)$.

Theorem 14 shows a competitive ratio for efficiency (the difficult part of the proof is deferred to Section 4.7.3):

Theorem 14. *Mechanism $\mathtt{M}_\alpha(\gamma)$ is a $\frac{1}{1+\gamma}$-approximation to the optimal offline matching: $w(S) \geq \frac{1}{1+\gamma}\mathrm{Opt}[\mathbf{w}]$.*

Proof. My key technical lemma shows that if non-survivor bids are essentially given a late-arriving penalty, then the optimal offline matching coincides with $\mathtt{M}_\alpha(\gamma)$'s solution:

Lemma 2. $S = \mathrm{Opt}[\tilde{\mathbf{w}}]$ *for* $\tilde{\mathbf{w}}(i) = \begin{cases} w^{\mathrm{sv}}(i), & \text{if } i \in S \\ \frac{w(i)}{1+\gamma}, & \text{if } i \notin S \end{cases}$

I prove this Lemma in Section 4.7.3. To finish the theorem, let $\hat{\mathbf{w}}(i) = \max(w^{\mathrm{sv}}(i), \frac{w(i)}{1+\gamma})$ if $i \in S$, and $\frac{w(i)}{1+\gamma}$ otherwise. I have $w(S) \geq \hat{\mathbf{w}}(S) = \mathrm{Opt}[\hat{\mathbf{w}}] \geq \mathrm{Opt}[\mathbf{w}]/(1+\gamma)$: each inequality is implied by the fact that no bidder's contribution decreases when going from the left to the right hand side. Lemma 2 yields the equality: when going from $\tilde{\mathbf{w}}$ to $\hat{\mathbf{w}}$ only bids already in the optimum (i.e. S) can increase. □

The following bound assures us that not too much utility (of bumped bidders) is sacrificed for high efficiency:

Theorem 15. *Total bumped bidder weight* $w(R) \leq \frac{w^{\mathrm{sv}}(S)}{\gamma}$.

Proof. For an $r \in R$, let $\mathrm{s}^*(r) \in S$ be the survivor at the end of the sequence of bumps that starts from r. For an $s \in S$, let R_s be the refunded bidders in s's sequence of bumps: $R_s = \{r \in R : \mathrm{s}^*(r) = s\}$. As R is the disjoint union of R_s for all $s \in S$, the theorem follows by showing:

$$\text{For all } s \in S, \quad w(R_s) \leq w^{\mathrm{sv}}(s)/\gamma. \tag{4.3}$$

To show Eq. (4.3), fix $s \in S$, and, to simplify notation, assume that the elements in R_s are $1, 2, \ldots, s-1, s$: $j+1$ bumps j, $\forall\, 1 \leq j < s$. I now show that: $\sum_{j=1}^{s-1} w(j) \leq w^{\mathrm{sv}}(s)/\gamma$. As s bumps $s-1$, $w_{s-1} \leq \frac{w^{\mathrm{sv}}(s)}{1+\gamma}$. As $j+1$ bumps j, $\forall\, 1 \leq j \leq s-2$, $w_j \leq \frac{w_{j+1}}{1+\gamma}$. Thus by induction, $w_j \leq w^{\mathrm{sv}}(s)(1+\gamma)^{j-s}$, $\forall\, 1 \leq j < s$. I get $\sum_{j=1}^{s-1} w_j \leq w^{\mathrm{sv}}(s) \sum_{j=1}^{s-1}(1+\gamma)^{j-s} \leq w^{\mathrm{sv}}(s)/\gamma$. □

Let the *effective weight* of a solution be the weight of the matching minus a penalty amounting to the total utility loss by bumped bidders ($\alpha w(i)$ for each $i \in R$). Theorem 15 implies $w(S) - \alpha w(R) \geq w(S)(1 - \alpha/\gamma)$, implying (by Theorem 14) the following lower bound on effective weight:

Corollary 3. *The* $\mathtt{M}_\alpha(\gamma)$ *mechanism is a* $\frac{1-\alpha/\gamma}{1+\gamma}$*-approximation to the optimal offline matching in terms of effective weight:* $w(S) - \alpha w(R) \geq \frac{1-\alpha/\gamma}{1+\gamma}\mathrm{Opt}[\mathbf{w}]$.

Theorems 14 and 15 have analogs in [62]. My constants are tighter because in my model, a bidder's value for any item is the same, and all edges incident to a bidder arrive simultaneously. My bounds for revenue, efficiency and effective efficiency are almost tight:

Example 7 (tight bounds)**.** *Consider* $k+2$ *truthful bidders competing on one item; the* i*-th bidder to arrive has value* $(1+\gamma)^{i-1}$ *unless* $i = k+2$*, whose value is* $(1+\gamma)^{k+1} - \varepsilon$*. Bidder* $i+1$ *bumps* i*,* $\forall\, 1 \leq i \leq k$*. Only the* $k+1$*-st bidder survives. Bumped bidders' total weight is* $\sum_{i=0}^{k-1}(1+\gamma)^i = ((1+\gamma)^k - 1)/\gamma$*.* $\mathrm{Opt} = (1+\gamma)^{k+1} - \varepsilon$*.*

I also present an empirical[4] *upper* bound on how well *any* deterministic algorithm can approximate the effective weight:

Proposition 21. *Fix* n *(the number of bidders). No deterministic online algorithm can approximate the optimal offline matching in terms of effective weight with a factor better than* c_n*, where* c_n *is the lowest number (if any) in* $[0,1]$ *for which Eqs.* (4.12) *and* (4.13) *simultaneously hold (see Section 4.7.4). Based on computing* c_n *numerically, I conjecture that* c_n *approaches* $2\alpha + 1 - 2\alpha^{0.5}(\alpha+1)^{0.5}$ *as* $n \to \infty$*.*

For $\alpha < \frac{\sqrt{5}-1}{2} \simeq 0.618$ and the best γ given α, the approximation ratio in Corollary 3 for $\mathtt{M}_\alpha(\gamma)$ matches this upper bound (see Section 4.7.4 for further discussion).

Thus, despite $\mathtt{M}_\alpha(\gamma)$'s simplicity,

[4] Babaioff et al. [7] offer an insightful technical proof.

Theorem 16. *$\mathtt{M}_\alpha(\gamma)$ is optimal among deterministic algorithms with respect to the competitive ratio of effective weight.*

In an extension [8] of [7], the same authors show that there exist randomized algorithms that achieve a competitive ratio of $1+\alpha-2^{0.5}\alpha^{0.5}(1+\alpha)^{0.5}$ which is more than $2\alpha+1-2\alpha^{0.5}(\alpha+1)^{0.5}$ for $\alpha<0.5$, thus improving upon the optimal deterministic competitive ratio for such α.

The following Lemma is used in proving $\mathtt{M}_\alpha(\gamma)$'s constant competitiveness for (effective) efficiency.

Lemma 3. $v(S^T)+w^{\mathrm{sv}}(S^T)=\sum_{s\in S^T} v(s)+\sum_{s\in S^T} w^{\mathrm{sv}}(s)\geq \frac{\mathrm{Opt}[\mathbf{v}]}{1+\gamma}$

Proof. Let $w'(x) := \begin{cases} \max(v(x), w^{\mathrm{sv}}(x)), & \text{if } x\in S^T \\ w(x), & \text{if } x\notin S^T \end{cases}$. Clearly, $w(x)\geq w'(x), \forall\, x$ and $v(s)+w^{\mathrm{sv}}(s)\geq w'(s)\,\forall\, s\in S^T$.

$S^T(\mathbf{w})=S^T(\mathbf{w}')$ since only survivors in $S^T(\mathbf{w})$ change their bid, still bidding above their survival thresholds. $\sum_{s\in S^T} w'(s)\geq \mathrm{Opt}[\mathbf{w}']/(1+\gamma)$. The claim follows by noting that $\mathrm{Opt}[\mathbf{w}']\geq \mathrm{Opt}[\mathbf{v}]$ since $w'(x)=w(x)\geq v(x), \forall\, x\notin S^T$. □

Computational requirements. Each update step in $\mathtt{M}_\alpha(\gamma)$ amounts to finding an augmenting path with respect to the current matching, which can be done in $O(|E|)$ where E is the set of edges connecting each current and past bidder to each item in his choice set.

4.6 Game-theoretic properties

I now focus on the game-theoretic properties induced by my mechanism. I first offer some more intuition on survival and acceptance weights and then prove Theorem 11.

Recall from Section 4.4 that i is rejected if $w(i) < w^{\mathrm{ac}}(i)$, bumped if $w^{\mathrm{ac}}(i) \le w(i) < w^{\mathrm{sv}}(i)$ and a survivor if $w^{\mathrm{sv}}(i) \le w(i)$. If i bumps j^*, $w^{\mathrm{ac}}(i) = (1+\gamma)w(j^*)$ but $w^{\mathrm{sv}}(i)$ can either be $(1+\gamma)w(j^*)$, $w(k)$ for a (past or future) bumped bidder k or $\frac{w(k)}{1+\gamma}$ for a future rejected k. Thus, $w^{\mathrm{ac}}(i)$ can be computed when i arrives but $w^{\mathrm{sv}}(i)$ may depend on future bidders and can only be computed at $T+1$.

Let us focus now on a survivor i's ($w(i) \ge w^{\mathrm{sv}}(i)$) payment in Eq. (4.2). The common case is when $w^{\mathrm{ac}}(i) < w^{\mathrm{sv}}(i)$: i gets a discount amounting to the highest bump payment he could have otherwise obtained: $\alpha w^{\mathrm{sv}}(i)$. The special case of $w^{\mathrm{ac}}(i) = w^{\mathrm{sv}}(i)$ occurs when i's acceptance is enough for his survival (in particular if i is the last bidder). When $w^{\mathrm{ac}}(i) = w^{\mathrm{sv}}(i)$, from the bidder's point of view, $\mathtt{M}_\alpha(\gamma)$ posts a price of $w^{\mathrm{sv}}(i)$.

Say that in Example 6 a bidder B_5 were to arrive after B_4 bidding 10.5 on I_a. Only $w^{\mathrm{sv}}(3)$ would change to $\frac{10.5}{1.5} = 7$ and B_3's price would now be $(1-\alpha)\cdot 7$ which may be *lower* than 6.6. Unless a bidder i's $w^{\mathrm{sv}}(i)$ price coefficient goes from 1 to $1-\alpha$ (like in Example 6 for B_3 if B_5 arrives), i's price cannot go down if new bidders arrive.

4.6.1 Proof of Theorem 11.

If bidder i invests and bids his true value $v(i)$, then his utility is: $v(i) - p_i \ge v(i) - w^{\mathrm{sv}}(i) \ge 0$ if he survives, 0 if he is rejected, or $\alpha v(i) - \alpha v(i) = 0$ if he is bumped. This establishes (1).

Clearly, investing is preferred if and only if winning. If $w^{\mathrm{ac}}(i) < w^{\mathrm{sv}}(i)$, bidder i's highest possible bump payment is $\alpha w^{\mathrm{sv}}(i)$. The price of $(1-\alpha)w^{\mathrm{sv}}(i)$ has been chosen such that i prefers winning to being paid $\alpha w^{\mathrm{sv}}(i)$ if and only if $v(i) \ge w^{\mathrm{sv}}(i)$. That is, i's best response is to bid just below $w^{\mathrm{sv}}(i)$ if $v(i) < w^{\mathrm{sv}}(i)$ and to bid his true value otherwise. This establishes (2), (3) and (4) for this case.

If $w^{\mathrm{ac}}(i) = w^{\mathrm{sv}}(i)$, then i can never get a bump payment and i simply faces a take-it-or-leave-it offer of $w^{\mathrm{sv}}(i)$. Bidding truthfully is a best response in this case, and (2), (3) and (4) follow. $\square$

Theorem 11 establishes a separation property: for fixed bids by others, a bidder's best-response results in the same allocation for him as when truthful. Since this last statement is conditioned on others' bids, it is possible that the allocation at a pure Nash equilibrium (when each bidder is best-responding to others' bids) is different than the one resulting if all bidders were truthful.

Let us note a weakness of $\mathtt{M}_\alpha(\gamma)$'s best-response structure. Consider a bidder i that would not survive when truthful ($v_i < w^{\mathrm{sv}}(i)$) but for whom being bumped is possible ($w^{\mathrm{ac}}(i) < w^{\mathrm{sv}}(i)$). i can benefit by overbidding just below $w^{\mathrm{sv}}(i)$ instead of being truthful.

4.6.2 Speculators

In $\mathtt{M}_\alpha(\gamma)$, any participant that is bumped and did not invest obtains the bump payment for free. This motivates the following definition:

Definition 30. **[Speculators]** *A* speculator *is a bidder who does not invest. The auctioneer does not know if a bidder is a speculator or not.*

Speculators are thus bidders who *only* aim to get bump payments. Note that an investing bidder may also overbid for a higher bump payment if bumped (recall that, by Theorem 11, dishonest bidding cannot help a survivor), while having value for an item.

In this section I will focus on speculators' strategic behavior. At times, I may allow speculators more manipulations, such as colluding with each other or lying about their choice sets $L(i)$ or arrival times.

Prop. 22 bounds overbidders' (and thus speculators') total monetary profit:

Proposition 22. *Overbidders' profit is at most* $\alpha\mathrm{Opt}/\gamma$.

Proof. Denote overbidders' profit by $\Pi \leq -(1-\alpha)w^{\mathrm{sv}}(S \cap \overline{H}) + \alpha w(R)$. By Theorem 15, $w(R) \leq (w^{\mathrm{sv}}(S \cap \overline{H}) + w(S \cap H))/\gamma$. I get $\Pi \leq -(1-\alpha-\frac{\alpha}{\gamma})w^{\mathrm{sv}}(S \cap \overline{H}) + \frac{\alpha}{\gamma}w(S \cap H)$. The claim follows since $(1-\alpha-\frac{\alpha}{\gamma})w^{\mathrm{sv}}(S \cap \overline{H}) \geq 0$ and $w(S \cap H) \leq \mathrm{Opt}$. □

At first glance, it would seem that speculators' best strategy is to induce an assignment of actual bidders of weight as high as possible in the survivor set, since then overall bump payments would be maximized. This is true in some cases but not always, and indeed colluding speculators may even want some of them to survive:

Theorem 17. *Under the $\mathtt{M}_\alpha(\gamma)$ mechanism,*

1. *there exist input instances such that optimal speculator bidding induces optimal efficiency, and truthful bidding is a Nash equilibrium for all non-speculators;*
2. *there exist input instances where optimal speculator bidding induces sub-optimal efficiency,*
3. *there exist input instances where there is no pure-strategy Nash equilibrium,*
4. *there exist input instances where speculators achieve higher profit if they "sacrifice", i.e. some of them intentionally survive so that others obtain high refunds.*

Definition 31 will render the proof of Theorem 17 more concise. A speculator who is bumped with a bid of x could have obtained more bump payment by entering an earlier bid of at most $x/(1+\gamma)$; likewise, he could have obtained yet more by bidding earlier $x/(1+\gamma)^2$; and so on:

Definition 31. *Let $x > 0$. I say that the speculator σ is an x-*geometric speculator *with choice set N if σ places bids as follows on choice set N. Let ε be the minimum strictly positive bid that can be made and*

$$l = 1 + \left\lfloor \frac{\log(x/\varepsilon)}{\log(1+\gamma)} \right\rfloor \quad \textit{i.e.} \quad l \textit{ is integer } \& \frac{x}{(1+\gamma)^l} \geq \varepsilon > \frac{x}{(1+\gamma)^{l+1}}$$

Then σ places $l+1$ consecutive bids (each under a different identity) of $\frac{x}{(1+\gamma)^l}, \frac{x}{(1+\gamma)^{l-1}}, \ldots, \frac{x}{(1+\gamma)}, x$ on N.

Proof. 1. Consider two items and two speculators arriving before two bidders with values 1 and $C > 1+\gamma$. Assume that speculators cannot collude. I now show via a case analysis that there is no pure strategy Nash Equilibrium for speculators. Suppose that at a pure Nash equilibrium the two speculators bid $l < h$. If one of them is not refunded, then that is h and he can get a refund by underbidding l (unless $l = 0$, in which case l can bid between 0 and h and be refunded more). So both of them must be refunded. If h bids strictly less than $1-\varepsilon$ he can increase his bid to $1-\varepsilon$ and still be refunded. Since both bidders are refunded, l must bid at most $1/(1+\gamma)$. But he can do better by increasing his bid until right below $1-\varepsilon$ since he will be bumped by the C bidder (in particular he will prevent the actual 1 bidder from being accepted). But in this case h is no longer best-responding: he is not refunded anymore.

2. Fix a set of actual bids such that Opt[$\mathbf{v}$] bids arrive in increasing order. Suppose that speculators collude and want to maximize their joint revenue. Then optimal speculator bidding implies that:

   - no speculator survives, no investing bidder is bumped; all Opt bidders and only them are accepted.
   - speculators can achieve the highest payoff possible as given by Prop. 22.
   - truthful bidding is a Nash equilibrium for all investing bidders.

   Optimal speculator bidding in this case is as follows. For each bidder $i \in \text{Opt}$ there will be one $w(i)/(1+\gamma)$-geometric speculator σ_i with the same choice set.

   This result has an appealing interpretation. If very well informed, speculators can overcome the efficiency loss due to late bidders not being able to improve by a $1+\gamma$ factor over their earlier competitors.

In general however, speculators may prefer to induce a suboptimal perfect matching:

3. Consider two items i_1, i_2 and three bidders b_1, b_2, b_3 arriving in this order, each with choice set $\{i_1, i_2\}$. Note that any two bidders, but not all three, can be matched. Assume that $w(b_1) < w(b_3) < (1+\gamma)w(b_1)$ and $w(b_2) > 2w(b_3)$. The following analysis shows that speculators prefer the suboptimal set of actual bidders b_1 and b_2 to the optimal one with b_2 and b_3.

   - If both b_2 and b_3 survive, then speculators' profit is at most $2w(b_3)/\gamma$: the speculator bumped by b_2 must have a lower weight than the one bumped by b_3, which is at most $w(b_3)/(1+\gamma)$. Even if speculators are geometric, speculator profit can only go as high as $2w(b_3)/\gamma$.
   - If however b_1 and b_2 are alive when b_3 arrives, b_3 cannot bump b_1. By simply having one geometric $w(b_2)/(1+\gamma)$-speculator which is bumped by b_2, speculator profit is $w(b_2)/\gamma > 2w(b_3)/\gamma$.

4. Consider set I with k items, $k-1$ bidders bidding $C > 1$ all arriving before a bidder bidding 1; all k have choice set I. If speculators coordinate and participate with k identities as $C/(1+\gamma)$-geometric speculators on all the items then total speculator payoff is $(k-1)\alpha C/\gamma - (1-\alpha)C/\gamma = (k\alpha - 1)C/\gamma$, since $k-1$ will be bumped, but one will survive. If no speculator survives, the most money speculators can make is k/γ, by participating as k $1/(1+\gamma)$-geometric speculators. For any $\alpha > 1/k$, for a large enough C, speculators' profit is higher when one of them is sacrificed.

□

I may conclude from this theorem that it is unreasonable to expect stronger incentive properties than Theorem 11, such as truthfulness or a Nash Equilibrium. However, despite all the game-theoretic complexity that can arise from speculators, their effect on efficiency and revenue can still be bounded: implicitly via the results in Section 4.3 or explicitly in Prop. 22.

4.6.3 Relation to other game-theoretic concepts

I now position the mechanism's incentive properties with respect to several established game-theoretic concepts.

Algorithmic implementation in undominated strategies

In a game, an outcome is implemented by an equilibrium concept if the outcome is realized whenever players play strategies at such an equilibrium. Algorithmic implementation in undominated strategies has as desired outcome c-approximating an optimization problem if no dominated strategy is played.

Definition 32 (Babaioff et al. [9]). *A mechanism M is an* algorithmic implementation in undominated strategies *(AIUS) of a c-approximation algorithm if there exists a set D of* dominating *strategies such that*

- *M obtains a c-approximation whenever all players play strategies from D*
- *for any strategy $s \notin D$, there exists $s' \in D$ (computable in polynomial time from s) that dominates s.*

The following result effectively restates the Theorems in Section 4.3. It establishes the conditions under which $\mathtt{M}_\alpha(\gamma)$ is an algorithmic implementation.

Theorem 18. *Let D be the set of strategies representing truthful bids and overbids.*

- *Fix a* $c < \frac{1-\alpha-\frac{\alpha}{\gamma}}{1+\gamma}$. *Then* $\mathtt{M}_\alpha(\gamma)$ *is an AIUS with respect to the VCG revenue.*

- *Assume a c less than the appropriate constantin Theorem 12 and its total non-negative utility assumption. Then* $\mathtt{M}_\alpha(\gamma)$ *satisfies Def. 32 for (effective) efficiency.*

Being truthful is an equilibrium in undominated strategies of $\mathtt{M}_\alpha(\gamma)$ and I argue that it is a *focal point* (Schelling [82]), being highlighted by Theorem 11.

For the remainder of this section I will analyze the structure of best-responses, assuming that bidders' values and order of arrival are fixed and public knowledge.

CURB sets

For any player, a joint strategy in a CURB set [10] by the other players also has *all* best-responses to it in the CURB set.

Definition 33 (CURB set)**.** *Consider a game with n players. A set of pure strategies* $\Sigma_1 \times \cdots \times \Sigma_n$ *is* closed under rationalizable behavior *(CURB set) if, for all* i, Σ_i *contains* all *best responses to any* $\sigma_{-i} \in \Sigma_{-i}$.

A minimal CURB set is a CURB set such that none of its proper subsets is a CURB set.

Basu and Weibull [10] prove that any game with compact strategy sets and continuous payoff functions has at least one minimal CURB set C and C contains the support of at least one mixed Nash equilibrium. Additionally, the minimal CURB set is the least set-theoretic generalization of Nash equilibrium.

The following Proposition establishes the richness of any CURB set in a simple case of my algorithm with two players. It exhibits a minimal CURB set (which contains a pure strategy Nash equilibrium) even though the strategy sets are not compact (since a player can bid any positive number).

Proposition 23. *Suppose there is only one item and only two bidders, B_1 and B_2. Also suppose that $(1+\gamma)v_1 < v_2$ and that bidders' values are public knowledge. Then any pure Nash equilibrium is parametrized by a $w_2 \in [(1+\gamma)v_1, v_2]$, with B_1 bidding $\frac{w_2}{1+\gamma}$, B_2 bidding w_2 where only B_2 invests. Furthermore, any such pair of strategies is a Nash equilibrium.*

The collection of CURB sets is the collection of Cartesian products of sets $\mathbf{W}_1$ with the whole strategy space for B_2 ($\langle [0,\infty), \mathrm{I}_{0,1}\rangle$) where $\mathbf{W}_1$ contains $\langle [0,\infty), \mathrm{I}_1\rangle \cup \langle [v_1,\infty), \mathrm{I}_0\rangle$.

Proof. I will often refer to Theorem 11.

Note that B_2 cannot obtain any refund. So B_2's investment decision is uniquely determined by him being accepted. Therefore bidding truthfully (and investing) is *always* one of his best-responses.

At any pure Nash equilibrium, B_2 must win (and invest). Otherwise, B_1 would have bid at least $\frac{v_2}{1+\gamma}$ and paid at least $1-\alpha$ of that, which is worse than being refunded. Given that B_2 wins, B_1's best-response is to get the highest possible bump payment, which implies that B_2 pays his own bid. Winning is then a best-response for B_2 if and only if he affords the item: $w_2 \leq v_2$. Were w_2 less than $(1+\gamma)v_1$, B_1 would prefer to win and invest, since then B_1's utility would be at least αv_1.

B_2's set of best responses W_2' to B_1 bidding $\frac{v_2}{1+\gamma}$ is $\langle [v_2,\infty), \mathrm{I}_1\rangle$. Consider some $\langle w_2', \mathrm{I}_1\rangle \in W_2'$ with $w_2' > v_2$. B_1's best response w_1'' to w_2' is $\langle \frac{w_2'}{1+\gamma}, \mathrm{I}_0\rangle$. B_2's set of best responses W_2'' to w_1'' is bidding at most $(1+\gamma)w_1'' = w_2'$ and being bumped: $\langle [0, (1+\gamma)w_2'), \mathrm{I}_{0,1}\rangle$: the item becomes too expensive for B_2. Since w_2' can be arbitrarily large, the union of the sets W_2'' (one for each w_2') spans B_2's entire strategy space. If B_2 bids 0, any positive bid by B_1 coupled with investment is a best-response.

The only strategies that have not been noted as a best-response for B_1 so far are $\langle [0, \frac{v_2}{1+\gamma}), \mathrm{I}_0\rangle$. Consider now $x_2 \in ((1+\gamma)v_1, v_2)$, which is a possible best-response by B_2. B_1's best-response to x_2 is $\langle \frac{x_2}{1+\gamma}, \mathrm{I}_0\rangle$.

Let $x_1 < v_1$. I claim that $\langle x_1, \mathrm{I}_0\rangle$ cannot be a best-response. This is clear if B_2 bids higher than $(1+\gamma)x_1$. If B_2 bids lower than that, then x_1 strictly prefers investing to not investing and being bumped. $\square$

If $(1+\gamma)v_1 > v_2$ then the CURB sets are the same. The set of pure Nash equilibria, however, is the set of bids w_1 and $w_2 = (1+\gamma)w_1 - \varepsilon$, where $w_1 \in [\frac{v_2}{1+\gamma}, v_1]$ and only B_1 invests. Whether $(1+\gamma)v_1$ is higher than v_2 or not, the set of Nash equilibria has the truthful winner winning by bidding some w that is at most his value. Furthermore, if w was winner's true value, he would still win.

Interestingly, any CURB set also includes dominated (lower than true value) bids for B_2. Such bids are only weakly dominated and can thus be best-responses (even if B_1 bids truthfully). The existence of a continuum of best-responses when surviving or not being accepted is also the reason for the richness of the CURB sets.

Set-Nash equilibrium

Lavi and Nisan [56] introduce the following set-theoretic generalization of a Nash equilibrium.

Definition 34 (Set-Nash equilibrium)**.** *In a game with n players, a collection of n* sets *of strategies $\Sigma_1, \ldots, \Sigma_n$ is a Set-Nash equilibrium if for each $i = 1..n$, for each $\sigma_{-i} \in \Sigma_{-i}$, there exists a strategy $\sigma_i \in \Sigma_i$ that is a best-response to σ_{-i}.*

Note that the collection where each Σ_i is i's entire strategy space is trivially a Set-Nash equilibrium. At the other end of the spectrum, if $\Sigma_i = \{\sigma_i\}$ for all i then $(\sigma_1, \ldots, \sigma_n)$ is a Nash equilibrium. Thus the Set-Nash equilibrium concept is more powerful for "small" Σ sets.

For $\mathtt{M}_\alpha(\gamma)$, in the example of Proposition 23, if $\langle w_2, \mathrm{I}_{0,1}\rangle \in \Sigma_2$ with $w_2 \geq (1+\gamma)v_1$

then $\langle \frac{w_2}{1+\gamma}, \mathrm{I}_0 \rangle \in \Sigma_1$: B_1's best-response to a "high" bid by B_2 is to obtain a bump payment as high as possible.

4.6.4 Other game-theoretic considerations

I now ask a couple of "what if?" questions, whose answers further help motivate my model choices.

For stronger incentive properties, a standard modification to $\mathtt{M}_\alpha(\gamma)$ is to pay a bumped bidder i $\alpha w^{\mathrm{sv}}(i)$ (a bid-independent amount) instead of $\alpha w(i)$. Example 8 shows why this may result in a deficit:

Example 8 (Alternate bump payments). *Consider two bidders on one item: B_1 arrives first, bidding 1, followed by B_2 bidding $L > (1+\gamma)^2/\alpha$. Bidder B_2 survives and pays $1+\gamma$. If B_1's bump payment were $\alpha w^{\mathrm{sv}}(e) = \alpha L/(1+\gamma)$ then the choice of L ensures that B_1 is paid more than B_2 pays, i.e. the mechanism runs a deficit.*

Relation to outcome graph. Consider one item, a bidder i and bids by others w_{-i} such that $w^{\mathrm{ac}}(i) < w^{\mathrm{sv}}(i)$. Let $a = w^{\mathrm{ac}}(i)$ and $s = w^{\mathrm{sv}}(i)$.

Denote i's outcome $\phi(w_i)$ in $\mathtt{M}_\alpha(\gamma)$, dependent on his bid w_i, as follows: $\phi(w_i) = D$ if i is rejected ($w_i < a$), $\phi(w_i) = R$ if i is accepted and bumped ($a \leq w_i < s$) and $\phi(w_i) = S$ if i is a survivor ($s \leq w_i$).

Consider the following simplified version of $\mathtt{M}_\alpha(\gamma)$. Assume that, when bumped, i does not invest, not incurring any penalty at all (regardless of his bid). Despite this, i receives a bump payment of αs (as opposed to α times his bid as in $\mathtt{M}_\alpha(\gamma)$). Assume that i invests whenever winning.

For outcomes X and Y, recall from Section 2.2 $\delta_{XY} = \inf_{w_i:\phi(w_i)=X}\{w_i(X) - w_i(Y)\}$:

$$\delta_{DR} = 0 \qquad \delta_{RD} = 0$$

$$\delta_{DS} = -a \qquad \delta_{SD} = s$$

$$\delta_{RS} = -s \qquad \delta_{SR} = s$$

Note that ϕ satisfies WMON, and cycle DSD actually has positive weight $s - a$. Since the domain of bids is convex (the real line for each bidder), it is no surprise (recall Theorem 6 from Chapter 2) that both cycles with three outcomes have non-negative weight as well.

Payments in this simplified version of $\mathtt{M}_\alpha(\gamma)$ are $p_i(D) = 0, p_i(R) = -\alpha s$ and $p_i(S) = (1-\alpha)s$. They *cannot* be written as $p_i(X) = \delta^*_{XX_0}$ for a fixed outcome X_0 and any outcome X, which is a "universal" payment function for a truthful function (recall Eq. (2.8)).

Out of the $p_i(X) - p_i(Y) \leq \delta_{XY}, \forall X, Y \in \{D, R, S\}$ inequalities needed for truthfulness, $p_i(S) - p_i(D) < \delta_{SD}$, $p_i(S) - p_i(R) = \delta_{SR}$, $p_i(R) - p_i(S) = \delta_{RS}$ and $p_i(R) - p_i(D) < \delta_{RD}$. These four inequalities are a different formulation of best-responses in $\mathtt{M}_\alpha(\gamma)$: truthful bidding (and investing) for a survivor and obtaining the highest refund (amounting to an αs bump payment) without investing for a rejected or bumped bidder.

I assumed throughout that as soon as a bidder arrives, his choice set is known. If however that is private information as well, incentives become weaker: in Example 9, no bid by B^* on his true choice set $\{i_1, i_2\}$ is a best-response if bidding on different item(s) instead is allowed. This example also suggests why a naive generalization of $\mathtt{M}_\alpha(\gamma)$ to the setting where bidders have a different value for each of several items would not be able to incentivize bidders to bid at least their true value for each item.

Example 9 (**Private choice sets**). *Consider two items i_1, i_2 and the following three bidders, arriving in this order: $B^{-3/2}$ with value $(1+\gamma)^{-3/2}$, B^* with value $x < \alpha(1+\gamma)^{-3/2}$ and B^1 with value 1, all with choice set $\{i_1, i_2\}$. Assume $B^{-3/2}$ and B^1 bid truthfully. I will show that, whenever B^* bids on $\{i_1, i_2\}$, he can do* strictly *better by bidding on i_1 only.*

I claim that if B^* bids on $\{i_1, i_2\}$ then his utility is at most $\alpha(1+\gamma)^{-3/2}$. This is clear if he survives. If he is bumped by B^1, then his bid cannot be higher than $(1+\gamma)^{-3/2}$ ($B^{-3/2}$'s bid), since B^1 can replace any of $B^{-3/2}$ and B^*. But then B^*'s bump payment is at most $\alpha(1+\gamma)^{-3/2}$. Let $0 < \varepsilon < 1/2$. By bidding $(1+\gamma)^{-1-\varepsilon}$ on i_1 only and being bumped by B^1, B^* can get utility $\alpha(1+\gamma)^{-1-\varepsilon} > \alpha(1+\gamma)^{-3/2}$.

I have however the following conjecture: if a bidder prefers surviving to being refunded, he is better off bidding on his true choice set. Surviving yields the same value on any item in the choice set and restricting the choice set should not decrease the bidder's payment.

One can show that, if bidders myopically and simultaneously best-respond (over sequences of instances of $\mathtt{M}_\alpha(\gamma)$), they may reach bid vectors with a negative sum of utilities.

Given the instantaneous nature of bidders, temporal manipulations (such as changes in the order of arrival or multiple bids) become less motivated. Note however that in $\mathtt{M}_\alpha(\gamma)$ a certain item's price is almost never decreasing with a new bidder: see the continuation of Example 6 at the start of Section 4.6. If bidders could delay, but not hurry their arrival, only temporal manipulations that improve bump payments would be useful.

4.6.5 Tension of strategyproofness and competitiveness

Proposition 24 highlights the need for partial incentive compatibility constraints (i.e. not dominant-strategy) when aiming for constant competitiveness.

Call the *optimal stopping problem with preemptions and self-interest* the problem of selecting the maximum entry in a sequence, supplied in an online fashion by selfish players, when cancellations are allowed. This problem is this chapter's setting for one item. Note that, in the absence of game-theoretic considerations, this problem can be solved optimally by selecting the current entry in the sequence if and only if it is higher than the previously selected entry.

Proposition 24. *For the optimal stopping problem with preemptions and self-interest, no deterministic mechanism exists that is simultaneously*

- *dominant-strategy incentive-compatible*
- *constant-competitive with respect to efficiency*
- *individually-rational*

Proof. Suppose for a contradiction that there was a deterministic incentive-compatible individually-rational mechanism M with a (relative) competitive ratio of $c \in (0, 1)$: $\exists b \in \mathbb{R}$ such that

$$v_M(Z) \geq c \cdot \mathrm{Opt}(Z) + b \text{ for any input sequence } Z. \tag{4.4}$$

Let $X \in \mathbb{R}_{\geq 0}$ such that $0 < cX + b$.

Consider a sequence Z starting with a bid of X, possibly followed by a bid of X', where $X < X' \cdot c + b$. M must select X as otherwise it would fail to satisfy Eq. (4.4) for $Z = \{X\}$. But, similarly, M must select X' to satisfy Eq. (4.4) for $Z = \{X, X'\}$.

As M is individually-rational, the X bidder must be paid back at least αX when bumped. When the X bidder has true value 0 and X' arrives, he prefers overbidding such that he is bumped and receives the bump payment. □

4.6.6 Extensions

I now consider the robustness of $\mathtt{M}_\alpha(\gamma)$ to variations in the assumptions used.

Matroids

Matroids are abstract structures that capture the basic properties of numerous problems of interest in combinatorial optimization and economics. Indeed, problems such as matching in bipartite graphs (in particular allocating indivisible goods, or multiple units thereof) or forming spanning trees of a graph can be cast in matroid terms.

All my results, except the ones regarding $\mathtt{M}_\alpha(\gamma)$'s computational complexity, extend to a setting where the items for sale are elements of a matroid, defined as follows.

Instead of K items, the seller has a finite set E of *elements*. There is a *matroid structure* M on E defined by a collection $\mathcal{I}$ of *independent* subsets of E ($\mathcal{I} \subseteq 2^E$):

- The empty set is independent: $\emptyset \in \mathcal{I}$,
- If some $X \subseteq E$ is independent then any subset $X' \subset X$ is also independent.
- If X, Y are independent, $|X| < |Y|$ then there exists $y \in Y \setminus X$ such that $X \cup \{y\}$ is also independent. (*the exchange axiom*)

A bidder bids on *exactly* one element of the matroid, which is known ahead of time to the seller and may vary across bidders.

A bipartite matching problem generates a *transversal* matroid. For each subset S of items there is an element e_S (together with sufficiently many copies) in E. A multi-set of elements $\{e_S^1, \ldots, e_S^l\}$ is independent if it can be matched to l distinct items $I_1, \ldots, I_l$, i.e. $I_1 \in S_1, \ldots, I_l \in S_l$. A bidder with choice set S bids on exactly one copy of element e_S. A set of bidders (elements) is independent if the bidders can be matched to items such that each one receives an item from his subset.

Identical items

For the remainder of this chapter, I consider identical items, with each bidder having each copy of the item in his choice subset. I first prove that survival thresholds (and, thus, prices) for survivors are "close" in $\mathtt{M}_\alpha(\gamma)$ if items are identical.

Fix a survivor B. When a later bidder B^+ arrives, let $w^{\mathrm{sv}}_{B^+}(B)$ be the minimum of: 1) the lowest active bid w^* other than B's ($w^* = \min\{w_x : X \neq B \text{ is active before } B^+\}$) and 2) B^+'s adjusted bid $\frac{B^+}{1+\gamma}$.

As in Section 4.7.3, denote by $w^{\mathrm{sv}}_{\leq B^+}(B)$ the minimum bid that B could have made in order to survive until after the arrival of bidder B^+. At B^+, the survival threshold $w^{\mathrm{sv}}_{\leq \cdot}(B)$ of a bidder B is increased if and only if $w^{\mathrm{sv}}_{\leq B^+ -1}(B) < w^{\mathrm{sv}}_{B^+}(B)$. That is, if and only if, were B to bid just above $w^{\mathrm{sv}}_{\leq B^+ -1}(B)$, B^+ would bump B instead of bumping no one or the lowest other active bid at $B^+ - 1$, depending on whether $\frac{B^+}{1+\gamma} < w^*$ or not. Thus

$$w^{\mathrm{sv}}_{\leq B^+}(B) = \max\left\{w^{\mathrm{sv}}_{\leq B^+ -1}(B), w^{\mathrm{sv}}_{B^+}(B)\right\} \text{ where } w^{\mathrm{sv}}_{B^+}(B) = \min\left\{\frac{B^+}{1+\gamma}, w^*\right\} \tag{4.5}$$

Proposition 25 establishes an intuitive property for $\mathtt{M}_\alpha(\gamma)$: the survival threshold of a later unit-demand bidder is lower, but not by more than a $1+\gamma$ factor, than that of an earlier unit-demand bidder. An analogous relationship (with a factor of $\frac{1-\alpha}{1+\gamma}$) holds between their prices.

Proposition 25. *Suppose that items are identical. Let S_a and S_b be survivors in $\mathtt{M}_\alpha(\gamma)$ such that S_b arrives after S_a. Then the ratio of their survival thresholds and the ratio of their prices are bounded: $\frac{w^{\mathrm{sv}}(S_a)}{w^{\mathrm{sv}}(S_b)} \in [\frac{1}{1+\gamma}, 1]$ and $\frac{p_a}{p_b} \in [\frac{1-\alpha}{1+\gamma}, 1]$.*

Proof. I claim by induction that at any bidder B^+ after S_b,

$$\frac{w^{\mathrm{sv}}_{\leq B^+}(S_a)}{w^{\mathrm{sv}}_{\leq B^+}(S_b)} \in \left[\frac{1}{1+\gamma}, 1\right] \tag{4.6}$$

Suppose that S_b bumped a bidder B_1 who bids w_1. Then $w^{\text{sv}}_{\leq S_b}(S_b) = w^{\text{sv}}_{S_b}(S_b) = (1+\gamma)w_1$ and $w^{\text{sv}}_{S_b}(S_a) = w_1$. Thus $w^{\text{sv}}_{\leq S_b}(S_a) \geq \frac{1}{1+\gamma} w^{\text{sv}}_{\leq S_b}(S_b) = w_1$ by Eq. (4.5) for $B = S_a$ and $B^+ = S_b$. Since B_1 is active before S_b arrives, it must be that $w^{\text{sv}}_{\leq S_b}(S_a) \leq (1+\gamma)w_1$: otherwise, a bidder with value higher than B_1 would have been bumped before B_1, which is impossible. Therefore $w^{\text{sv}}_{\leq S_b}(S_a) \leq w^{\text{sv}}_{\leq S_b}(S_b)$. The bound in Eq. (4.6) for $B^+ = S_b$ is thus obtained.

For all B^+ after S_b at which the lowest active bid w^* is neither S_a's nor S_b's, $w^{\text{sv}}_{B^+}(S_a) = w^{\text{sv}}_{B^+}(S_a)$. Eq. (4.6) follows from Eq. (4.5) for each such B^+.

Suppose that, starting with some B_0, the lowest active bid is S_a's or S_b's. Since S_a and S_b survive, no bidder starting with B_0 can be accepted. In particular, for all B^+ after B_0, $\frac{w_{B^+}}{1+\gamma} < \min\{w_a, w_b\}$ and $w^{\text{sv}}_{B^+}(S_a) = w^{\text{sv}}_{B^+}(S_b) = \frac{w_{B^+}}{1+\gamma}$. Eq. (4.6) follows from Eq. (4.5) for each such B^+.

The bounds on the ratio of S_a's and S_b's survival thresholds ($\frac{w^{\text{sv}}(S_a)}{w^{\text{sv}}(S_b)} \in [\frac{1}{1+\gamma}, 1]$) are obtained by choosing B^+ in Eq. (4.6) as the last bidder to arrive.

Recall that a bidder's price is his survival threshold, possibly discounted by a $1-\alpha$ factor. A bidder B receives the α discount only when his survival threshold is increased beyond $w^{\text{ac}}(B)$. The inductive proof above also reveals that S_a and S_b never appear in each other's $w^{\text{sv}}_{B^+}(\cdot)$ and furthermore $w^{\text{sv}}_{B^+}(S_a) = w^{\text{sv}}_{B^+}(S_b)$ at any B^+ later than S_b. Since $w^{\text{sv}}_{\leq S_b}(S_b) \geq w^{\text{sv}}_{\leq S_b}(S_a)$, if S_b receives the discount then S_a must also receive the discount. $\square$

Note that $\frac{w^{\text{sv}}_{\leq S_b}(S_a)}{w^{\text{sv}}_{\leq S_b}(S_b)}$ may be strictly between $\frac{1}{1+\gamma}$ and 1. Consider a small modification to Example 10 with two bidders S_a and S_b bidding 5 and 7 instead of B and bumping bidders with bids of 2 and 3 respectively, where $\gamma = 1$. Then $w^{\text{sv}}_{\leq S_b}(S_a) = 4, w^{\text{sv}}_{\leq S_b}(S_b) = 6$. It is also possible that S_a receives the α discount, while S_b does not: continuing with this example, if a 9 bidder arrives after the 7 bidder S_b, then S_a's survival threshold increases to 4.5 but S_b's remains at 6.

Clearly, Proposition 25 does not hold if items are not identical: for example if there are two items such that half the bidders are interested in one item and the other half in the other item then there are practically two independent problems, one for each half-set of bidders and the corresponding item.

Multi-unit demand

One of the main limitations of my bidder model is that each bidder can only demand one item. I present now an informal discussion of the challenges in extending $\mathtt{M}_\alpha(\gamma)$ to a setting where a bidder can have value for more than one item.

If items are not identical, Example 9 suggests why a naive generalization of $\mathtt{M}_\alpha(\gamma)$ to multi-unit demand bidders would not be able to incentivize bidders to bid at least their true value for each item: a bidder may prefer bidding only on an item that yields a high bump payment instead of surviving on an item for which he has low value.

In the following I will consider, for simplicity, two identical items and a natural model where any bidder has a lower marginal value for the second item than for the first item received.

It seems difficult to generalize $\mathtt{M}_\alpha(\gamma)$ to incentivize winners to bid both their true values. Consider a bidder B with a value of 1 for each of two items, followed by a bidder with a value of 100 for exactly one item out of the two. Then B does best going for a high refund, which cannot be achieved unless he overbids on *both* items, since the *lower* bid is the one bumped.

There are challenges in generalizing $\mathtt{M}_\alpha(\gamma)$ to incentivize bidders to bid *at least* their true value for each item, a game-theoretic property of $\mathtt{M}_\alpha(\gamma)$. Fix, for simplicity, $\gamma = 1$: a new bid bumps an existing bid if and only if the new bid is at least twice the old bid.

Let us first note that such a generalization should consider the new bidder's bids in decreasing order and then essentially treat them as bids in $\mathtt{M}_\alpha(\gamma)$. Clearly, in this fashion, no bidder can bump one of his own bids.

Example 10. *Two unit-demand bidders with bids* 2 *and* 3 *arrive followed by* B *who bids* 7 *and* 5 *for the first and second item respectively. Even though* B *could bump both active bids if his bids were instead considered in increasing order, only* B*'s first bid (*7*) is accepted and bumps* 2*.* B*'s second (*5*) bid can no longer bump* 3 *and is rejected.*

If B *had both bids accepted then* B *should pay* 4 *and* 6*, achieving an utility of* $7 + 5 - 4 - 6 = 2$*. Following the descending order, only* B*'s* 7 *bid is accepted, for an utility of* $7 - 4 = 3$*. Intuitively, a new bidder's bids are accepted as long as that it is in his interest given his bids: the price for the second item is* 6*, higher than* B*'s value for it,* 5*.*

Proposition 25 establishes that any two survival thresholds in $\mathtt{M}_\alpha(\gamma)$ are within a $1 + \gamma$ factor of each other. In particular, when the last bidder B bumps some bidder B' instead of a survivor B^a (with bids $w_a < w_{B'}$), survival thresholds may be: $(1 + \gamma)w_{B'}$ for B, but $w_{B'}$ for B^a. Essentially because of this distinction, a multi-unit demand bidder (B^a) can prefer *underbidding* so that he competes against $w_{B'}$ without a $1 + \gamma$ multiplicative handicap, as shown in the following example.

Example 11. *Consider bidders* $B^\varepsilon, B^2, B^a, B^{40}$ *arriving in this order. Let* B^a *have equal marginal value* 5 *for the first and second item and bid an equal amount* w_a *(*3 *or* 5*) on each. All other bidders have zero value for the second item and bid truthfully* $\varepsilon, 2$ *and* 40 *respectively. Note that, regardless of whether* w_a *is* 3 *or* 5*,* B^{40} *and a bid by* B^a *survive.*

When B^a *bids truthfully (*$w_a = 5$*), he has a* 2 *handicap when bumping* B^2 *(he needs a value of at least* **4***). When* B^a *underbids (*$w_a = 3$*), he only needs to have a value of at least* **2** *to have* B^2 *bumped by* B^{40}*.*

There is, however, a more subtle difficulty. Example 12 shows that, by underbidding on his second bid, bidder 3 can have a bid survive by influencing the acceptance decisions for future bidders. Note that, similar to Example 11, a bidder (0.90) has the $1+\gamma$ handicap only for one of 3's true and lower second bids.

Example 12. *Bidder 3 has values* 1.30 *and* 0.50 *for first and second item respectively. All other bidders have zero value for the second item: ε, 0.20 (both before 3) and* 0.90, 1.50 *and* 6 *(after 3).*

The set of active bids evolves as follows, depending on whether 3's second bid is 0.50 *(truthful) or* 0.30 *(an underbid).*

	3	*0.90*	*1.50*	*6*
$\{1.30, 0.30\}$	{*1.30, 0.20*}	{*1.30, 0.90*}	{*1.30, 0.90*}	{***1.30***, *6*}
$\{1.30, 0.50\}$	{*1.30, 0.50*}	{*1.30, 0.50*}	{*1.30, 1.50*}	{*1.50, 6*}

Bidder 3 prefers underbidding since it leads to winning an item.

Surprisingly, before 1.50 arrives, the lowest active bid is higher *when 3 underbids than when truthful.*

Note that if bidder 3 only bids 1.30 and 0.30 is another bidder, then bidder 3's price in $\mathtt{M}_\alpha(\gamma)$ is $(1-\alpha)0.90$. Proposition 25 suggests the following conjecture for two item settings where a bidder with value for one or two items can win one item by underbidding: the utility from surviving (by underbidding and investing) is better than the utility for being bumped (by bidding truthfully and not investing) by no more than a constant fraction of one's value.

Other extensions

Here are some potential additional features for an auction with seller cancellations.

reserve prices By starting out $\mathtt{M}_\alpha(\gamma)$ with dummy bidders with non-zero values, the seller can set *reserve prices* on items. Reserve prices are particularly useful if the seller has information about future bidders, for example a distribution on their values as in Chapter 3.

bidder cancellations Suppose that a *bidder i* wants to cancel his allocation. Such a cancellation could be offered even for free when the seller receives another bid with a high value compared to bidders other than i. Otherwise, the seller could charge a fee for i's cancellation, which may have to be close to i's potential payment if the mechanism closed when i requests his cancellation.

discounting Suppose that bidders discount future utility. The game-theoretic properties of $\mathtt{M}_\alpha(\gamma)$ hold if bump payments are made at the same time as allocations, despite the fact that the amount and occurrence of the bump payment are known when the bidder is bumped.

departures If, unlike in $\mathtt{M}_\alpha(\gamma)$, a bidder's type includes his patience, then a bidder will aim to pretend a minimal patience in order to give future bidders a $1+\gamma$ handicap.

insurance Babaioff et al. [8] present an extension of $\mathtt{M}_\alpha(\gamma)$, with the same competitive ratio for effective efficiency, for the case where a bidder requests for him a higher α (a higher bump payment) or a higher γ (making it more difficult for future bidders to bump him). This algorithm, effectively providing insurance to bidders, accepts such a request provided the bid is high enough compared to the improved parameter requested.

optimality given model If the number of bidders and a distribution on bidders' values are known, then one can compute the optimal allocation rule with cancellations via a dynamic program analogous to [38]. For example, if there are two bidders with $Exp(1)$ values, then the efficiency-maximizing threshold for the first bidder is a value

R_1^α such that $e^{-(1+\alpha)R_1^\alpha} = 1 - R_1^\alpha$. Note that this equation has a positive solution if and only if $\alpha > 0$. As expected, $R_1^\alpha = 0$ (first bidder is accepted for any bid) when $\alpha = 0$ (cancellations are free) and $R_1^\alpha \to 1$ (first bidder is accepted with any bid higher than the mean of the $Exp(1)$ distribution) when $\alpha \to \infty$ (cancellations are essentially not allowed).

Summary

Mechanisms for reservations reduce the uncertainty for both buyers and sellers and are well established in many markets today. Sellers often have a large inventory of items and seek automatic, online pricing and allocation of reservations.

I present a simple model for auctioning items in advance, which allows canceling allocations, imposing a proportional utility loss to the buyer, compensated by the seller via a bump payment. I present an efficiently implementable online mechanism to derive prices and bump payments that has many desirable properties of incentives, revenue and efficiency. This mechanism is essentially a form of price discovery (like an auction), but its online nature induces a constant factor error. These properties hold even in the presence of speculators, who are in the game for earning bump payments only, and require no assumptions about bids' order of arrival or their value distribution.

If bidders have value for more than one item then achieving even partial incentive compatibility with an extension of the unit-demand mechanism appears challenging.

4.7 Missing proofs

4.7.1 Proof of Theorem 12

Proof. If in $\mathtt{M}_\alpha(\gamma)$ the sum of all bidders' utility is positive then

$$v(S^T) - (1-\alpha)w^{\mathrm{sv}}(S^T) + \alpha w(R^T) - \alpha v(R^T) \geq 0 \tag{4.7}$$

For efficiency, notice that:

$$v(S^T) \geq (1-\alpha)w^{\mathrm{sv}}(S^T) + \alpha v(R^T) - \alpha w(R^T) \tag{4.8}$$

$$\geq (1-\alpha-\frac{\alpha}{\gamma})w^{\mathrm{sv}}(S^T) \tag{4.9}$$

where Eq. (4.8) follows from Eq. (4.7) and Eq. (4.9) follows from Theorem 15. Lemma 3 then implies $v(S^T) \geq \frac{1-\alpha-\frac{\alpha}{\gamma}}{2-\alpha-\frac{\alpha}{\gamma}}\mathrm{Opt}[\mathbf{v}]$.

For effective efficiency, letting $y = \frac{1-\alpha-\frac{\alpha}{\gamma}}{1+\frac{\alpha}{\gamma}}$ yields

$$(1+y)(v(S^T) - \alpha v(R^T)) \geq (1-\alpha)w^{\mathrm{sv}}(S^T) - \alpha w(R^T) + y(v(S^T) - \alpha v(R^T)) \tag{4.10}$$

$$\geq (1-\alpha-(1+y)\frac{\alpha}{\gamma})w^{\mathrm{sv}}(S^T) + yv(S^T) \tag{4.11}$$

$$= y(v(S^T) + w^{\mathrm{sv}}(S^T))$$

where Eq. (4.10) follows from Eq. (4.7) and Eq. (4.11) follows from Theorem 15. Finally, Lemma 3 implies $v(S^T) - \alpha v(R^T) \geq \frac{1-\alpha-\frac{\alpha}{\gamma}}{(2-\alpha)(1+\gamma)}\mathrm{Opt}[\mathbf{v}]$. □

I conjecture that $v(S)-\alpha v(R) \geq \frac{1-\alpha-\frac{\alpha}{\gamma}}{1+\gamma}\mathrm{Opt}[\mathbf{v}]$: one can show that $v(S)-\alpha v(R) \geq \frac{1-\alpha-\frac{\alpha}{\gamma}}{1+\gamma}w^{\mathrm{sv}}(S)$. Note that $\frac{1-\alpha-\frac{\alpha}{\gamma}}{(2-\alpha)(1+\gamma)} \leq \frac{1-\alpha-\frac{\alpha}{\gamma}}{(2-\alpha-\frac{\alpha}{\gamma})(1+\gamma)} \leq \frac{1-\alpha-\frac{\alpha}{\gamma}}{1+\gamma}$.

4.7.2 Revenue - proof of Theorem 13.

Lemma 4. *A winning bidder's VCG payment is a losing bid. The VCG revenue can only increase if some bids in* $\mathbf{w}$ *are increased.*

Proof. An optimal matching can be found by adding bidders greedily to the matching in decreasing order of their values. This implies the following well-known (see e.g. [16], Fact 3.2) combinatorial property of my setting: $\forall\, i \neq x$, if $x \in \mathrm{Opt}[\mathbf{w}]$ then $x \in \mathrm{Opt}[\mathbf{w}_{-i}]$.

This fact implies that there exists a bidder k such that $\mathrm{Opt}[\mathbf{w}_{-i}] = \{k\} \cup (\mathrm{Opt}[\mathbf{w}] \setminus \{i\})$. But then i's VCG price must be $w'(k)$, i.e. a losing bid.

If Opt changes when bidder i's bid is increased, then i must displace a single lower bid by another bidder j since Opt is constructed greedily in decreasing order of bids. $\square$

Proof of Theorem 13. The payments received by $\mathtt{M}_\alpha(\gamma)$ are at least $w^{\mathrm{sv}}(S)(1-\alpha)$ and Theorem 15 implies that bump payments sum to at most $w^{\mathrm{sv}}(S)\alpha/\gamma$. Thus the theorem follows from showing $w^{\mathrm{sv}}(S) \geq \mathtt{REV}_{vcg}[\mathbf{w}]/(1+\gamma)$. I argue this in three steps below.

Let $\hat{\mathbf{w}}(i) = \max(w^{\mathrm{sv}}(i), w(i)/(1+\gamma))$ if $i \in S$, and $w(i)/(1+\gamma)$ otherwise.

1. I have $w^{\mathrm{sv}}(S) = \tilde{\mathbf{w}}(S) = \mathrm{Opt}[\tilde{\mathbf{w}}] \geq \mathtt{REV}_{vcg}[\tilde{\mathbf{w}}]$, where the second equality follows from Lemma 2, and the final inequality follows from the fact that VCG payments cannot be higher than VCG efficiency.

2. I claim that $\mathtt{REV}_{vcg}[\tilde{\mathbf{w}}] = \mathtt{REV}_{vcg}[\hat{\mathbf{w}}]$. To see this note that when going from $\tilde{\mathbf{w}}$ to $\hat{\mathbf{w}}$, only VCG winners may increase their bid. Increasing the bid of a winner has no effect on the allocation, and no effect on that winner's price. Furthermore it has no effect on any other price, since any price is a losing bid.

3. Finally, Lemma 4 implies $\mathtt{REV}_{vcg}[\hat{\mathbf{w}}] \geq \mathtt{REV}_{vcg}[\mathbf{w}/(1+\gamma)] = \mathtt{REV}_{vcg}[\mathbf{w}]/(1+\gamma)$ since VCG payments scale linearly if all the bids are multiplied by a scalar. $\square$

4.7.3 Proof of Lemma 2.

In this section it will be simpler to have as initial matching A_0 for $\mathtt{M}_\alpha(\gamma)$ an arbitrary perfect matching on dummy bidders instead of the empty matching. I introduce dummy bidders (one per item, each bidding 0) whose choice set is the whole set of items, arriving before all actual bidders. This will ensure that a perfect matching is maintained by $\mathtt{M}_\alpha(\gamma)$, but will not affect other arguments below[5]

At time t, I call currently accepted bidders *alive*, and denote the set of alive bidders as A_t. Let $X_t = \{b \in A_{t-1} : A_{t-1} \cup \{t\} \setminus \{b\}$ can be matched$\}$; X_t is the set of alive bidders at $t-1$ that can be swapped for t and $j^* = \operatorname{argmin}_{j \in X_i} w(j)$ (see Algorithm 4).

Assume wlog that bidder i arrives at time i. I denote by $w^{\mathrm{sv}}_{\leq t}(b)$ the minimum bid bidder b must make in order to survive up to and including time t. Then $w^{\mathrm{ac}}(b) = w^{\mathrm{sv}}_{\leq b}(b)$ and $w^{\mathrm{sv}}(b) = w^{\mathrm{sv}}_{\leq T}(b)$. It is clear that $w^{\mathrm{sv}}_{\leq t}(b) \leq w^{\mathrm{sv}}_{\leq t+1}(b)$.

Definition 35. *Let B be a set of bidders. I say that B is* tight *for a bidder i at time t if all bidders in B are alive at t, B can be matched but $B \cup \{i\}$ cannot be matched. I say that B* γ-dominates *a bidder i at time t if B is tight for i at t and for all $b \in B$, $w^{\mathrm{sv}}_{\leq t}(b) \geq w(i)/(1+\gamma)$.*

Lemma 5. *X_t is tight for t at t.*

Proof. X_t can be matched since $X_t \subseteq A_{t-1}$. Suppose for a contradiction that $X_t \cup \{t\}$ can be matched. Then $X_t \neq A_{t-1}$ since A_{t-1} is a perfect matching by assumption. Therefore there exists $X \subset A_{t-1} \setminus X_t, |X| = |A_{t-1}| - |X_t| - 1$ such that $X_t \cup \{t\} \cup X$ can be matched. There exists exactly one bidder $\{y\} = A_{t-1} \setminus (X_t \cup X)$ and $X_t \cup \{t\} \cup X = A_{t-1} \cup \{t\} \setminus \{y\}$ is a perfect matching, implying $y \in X_t$, contradiction. □

Let i^* be the time step when i ceases to be alive (i.e. $i^* = i$ if i is not accepted or the time i is bumped if i was accepted). I inductively construct a sequence $\{B_t\}_{i^* \leq t \leq n}$ as follows: if i is not accepted, $B_i = X_i$; if i is bumped by i^* then $B_{i^*} = X_{i^*} \cup \{i^*\} \setminus \{i\}$. At time $t \geq i^*+1$,

[5] When bidder t arrives, assume $A_{t-1} = A \cup D$ where D only contains dummy bidders and there exists a matching I_t of $A \cup \{t\}$ which matches t to some item i_t. By reassigning dummy bidders, I can assume that actual bidders are matched according to I_t. Then bidder t can bump at least the dummy bidder $d \in D$ that is matched to i_t in A_{t-1}.

- if no bidder in B_{t-1} is bumped, then I let $B_t = B_{t-1}$.
- if t bumps some $b \in B_{t-1}$ then I let $B_t = (B_{t-1} \cup X_t \cup \{t\}) \setminus \{b\}$

I will prove inductively on t that

Lemma 6. *B_t γ-dominates i at time t.*

Proof. By definition, all bidders in B_t are alive at t. I proceed by induction starting with the base case $t = i^*$. If i is not accepted ($i^* = i$), i cannot bump any bidder in X_i: therefore $\forall\, b \in X_i, w^{\mathrm{sv}}_{\leq i}(b) \geq w(i)/(1+\gamma)$. X_i is tight for i at i by Lemma 5. If i is bumped, then $w(i) \leq w^{\mathrm{sv}}_{\leq i^*}(r)$, $\forall\, r \in X_{i^*}$. $B_{i^*} = X_{i^*} \cup \{i^*\} \setminus \{i\}$ can be matched since they are all alive at i^*. $X_{i^*} \cup \{i^*\}$ cannot be matched: otherwise i^* would not bump $i \in X_{i^*}$.

In the inductive step, I assume that B_{t-1} γ-dominates i at $t-1$. If at time t, no bidder in B_{t-1} is bumped, then the claim obviously holds by the induction hypothesis. Otherwise, let $b \in B_{t-1}$ be the bidder that is bumped by t. Clearly, $(B_{t-1} \cup X_t \cup \{t\}) \setminus \{b\}$ can be matched since they are alive at t. Suppose for a contradiction that $B_t \cup \{i\} = (B_{t-1} \cup X_t \cup \{t\}) \cup \{i\} \setminus \{b\}$ could be matched. $i \notin B_t$ since i is no longer alive. $B_{t-1} \cup X_t$ can be matched since they are all alive at $t-1$. As $|B_{t-1} \cup X_t| = |B_t \cup \{i\}| - 1$, either $B_{t-1} \cup X_t \cup \{i\}$ or $B_{t-1} \cup X_t \cup \{t\}$ can be matched. The first case is not possible since a subset, $B_{t-1} \cup \{i\}$, cannot be matched (by the induction hypothesis); the second case is not possible since $X_t \cup \{t\}$ cannot be matched (Lemma 5). I have reached a contradiction, so B_t must be tight for i.

By the induction hypothesis, $\forall\, b' \in B_{t-1}$, $w^{\mathrm{sv}}(b')_{\leq t-1} \geq w(i)/(1+\gamma)$. As noted before, survival thresholds can only increase from $t-1$ to t and $w(t) \geq (1+\gamma)w(b)$. □

Proof of Lemma 2. Let V be the $\mathrm{Opt}[\tilde{\mathbf{w}}]$ assignment (where ties are broken in favor of bidders in S). Suppose for a contradiction that there exists a non-survivor $i \in V$. By

Lemma 6 for time n, i is dominated by a set $B_n \subseteq S$ at time n. Since $i \notin S$, but $B_n \subseteq S$, in $\tilde{\mathbf{w}}$ any bidder in B_n has a higher weight than i.

Since V is a perfect matching and B_n can be matched there must exist $V' \subset V \setminus B_n, |V'| = |V| - |B_n|$ ($V' = \emptyset$ if B_n is a perfect matching) such that $B_n \cup V'$ is a (perfect) matching. I know that $B_n \cup \{i\}$ cannot be matched, therefore $i \notin V'$. However, $i \in V$ therefore $i \in V \setminus V'$. $V \setminus \{i\}$ can be matched and has size $|V| - 1$. Therefore there $\exists b \in B_n \cup V'$, $b \notin V \setminus \{i\}$ such that $V \cup \{b\} \setminus \{i\}$ can be matched. That implies $b \in B_n \subseteq S$, i.e. $\tilde{\mathbf{w}}(b) \geq \tilde{\mathbf{w}}(i)$. But then $V \cup \{b\} \setminus \{i\}$ is a perfect matching of higher weight than V, contradiction. That is, $V \setminus S = \emptyset$, i.e. $V = S$ since both are perfect matchings. □

4.7.4 Proof of Proposition 21

For $c \in \mathbb{R}_+$, consider one item and a sequence of bids $\{a_k(c)\}_{1 \leq k \leq n}$ on it such that $a_1 = 1, a_2 = \frac{1}{c} > 1$ and $ca_{k+1}(c) = a_k(c) - \alpha \sum_{j=1}^{k-1} a_j(c) \,\forall\, k \geq 2$, implying

$$ca_{k+1} = (1+c)a_k - (1+\alpha)a_{k-1} \,\forall\, k \geq 2 \tag{4.12}$$

For a fixed $n \geq 1$, I will look for a $c = c_n$ such that

$$a_n(c) - \alpha \sum_{j=1}^{n-1} a_j(c) = ca_n(c) \iff \frac{a_n}{1+\alpha} = a_{n-1} \tag{4.13}$$

E.g. $c_2 = \frac{1}{1+\alpha} > c_3 = \frac{1}{1+2\alpha} > c_4 = \frac{2}{1+3\alpha+\sqrt{(1+5\alpha)(1+\alpha)}}$. Unfortunately, c_n does not have a nice closed form for $n \geq 4$ (in addition, c_n may be not be unique - the smallest $c_n \in [0,1]$ is then of interest). Furthermore, for certain c and n no such sequence may exist.

Proof of Theorem 21. Suppose towards a contradiction that there was a deterministic algorithm A with a competitive ratio $c' > c_n$. Assume that the bids that arrive are $a_1, \ldots, a_{k_0}$ for some $1 \leq k_0 \leq n$. Then at each k, the algorithm A must accept a_k, or its competitive ratio will be smaller than c_n when $k = k_0$. This is clear for $k = 1$. Fix $k \in [2, n-1]$. Let

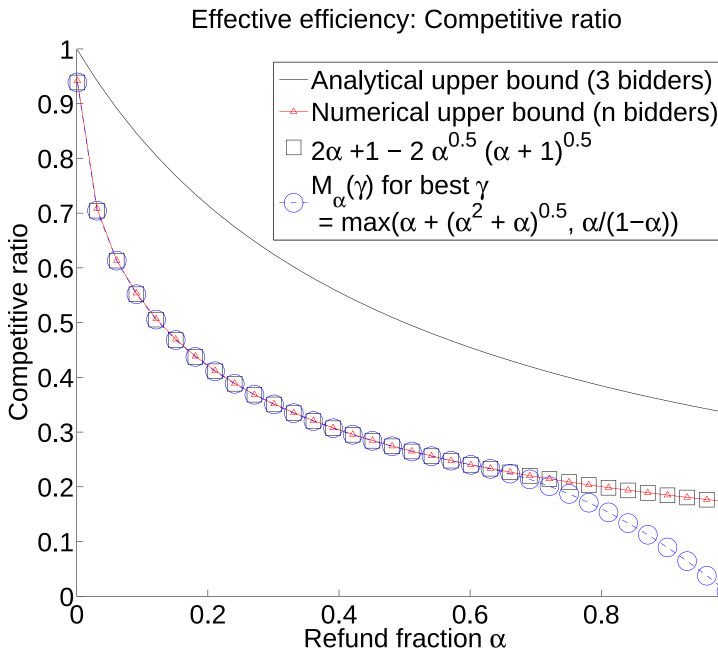

Figure 4.1: Effective efficiency competitive ratio (EECR) bounds as a function of α. The top curve is $c_3 = 1/(1+2\alpha)$. The middle curve is a numerical upper bound ($c = c_n$ of Eq. (4.13)) on any deterministic algorithm's EECR. The bottom curve shows (a *lower* bound on) my algorithm's EECR for the best γ_α: it matches the upper bound for $\alpha < 0.618$. γ is constrained by $\alpha < \gamma/(\gamma+1)$; if it were not, the bounds would match for all α.

M_k be the highest (i.e. the offline optimum) of $a_1, \ldots, a_k$. If A does not accept k then the competitive ratio on input $a_1, \ldots, a_k$ will be at most

$$\frac{a_{k-1}(c_n) - \alpha \sum_{j=1}^{k-2} a_j(c_n)}{M_k(c_n)} = \frac{c_n a_k(c_n)}{M_k(c_n)} \leq c_n$$

where the equality follows from Eq. (4.12). Now I claim that *whether or not* A accepts a_n, the competitive ratio will be at most c_n, which contradicts my assumption. If a_n is accepted, $\alpha \sum_{j=1}^{n-1} a_j$ has been lost due to bumping bidders $1, \ldots, n-1$; if a_n is rejected the effective efficiency is $a_{n-1} - \alpha \sum_{j=1}^{n-2} a_j$. By Eqs. (4.12) and (4.13), both quantities are a c_n fraction of a_n, which in turn is at most M_n, the optimal (effective) efficiency. $\square$

Figure 4.1 strongly suggests that the competitive ratio of any algorithm cannot be higher than $2\alpha + 1 - 2\alpha^{0.5}(\alpha + 1)^{0.5}$, shown as squares in the figure. Note that for this c the characteristic equation of Eq. (4.12) has a double root.

The triangles plot the minimum c found for the corresponding α for different values of n (I used Fibonacci values up to rank 12, i.e. largest n was 144). The c values were found via binary search. It was true in general, although not always, that the higher n, the lower c_n. I suspect that there exists an increasing sequence of integers $\{n_i\}_{i\geq 1}$ such that a solution c_{n_i} to Eqs. (4.12) and (4.13) converges from above to $2\alpha + 1 - 2\alpha^{0.5}(\alpha + 1)^{0.5}$ as $i \to \infty$.

Let $\underline{u}(\gamma) = \frac{1}{1+\gamma}\left(1 - \frac{\alpha}{\gamma}\right)$, the competitive ratio from Corollary 3. Subject to the constraint $\alpha \leq \frac{\gamma}{\gamma+1}$, $\underline{u}(\gamma)$ is maximized for $\gamma_0 = \max\{\alpha + \sqrt{\alpha^2 + \alpha}, \frac{\alpha}{1-\alpha}\}$. $\underline{u}(\gamma_0)$ is displayed in Figure 4.1 by circles. The value 0.618 (the golden ratio) is where $\frac{\alpha}{1-\alpha}$ becomes higher than $\alpha + \sqrt{\alpha^2 + \alpha}$. If $\alpha < 0.618$, $\underline{u}(\gamma_0) = 2\alpha + 1 - 2\alpha^{0.5}(\alpha + 1)^{0.5}$, which matches the numerical upper bound. The top curve plots $c_3 = 1/(1 + 2\alpha)$.

Chapter 5

Interdependent Values

Abstract. I initiate the study of incentive compatible dynamic auctions for bidders with interdependent values. These auctions are appropriate for dynamic environments with informational externalities (i.e. environments where other bidders' information on a good's value leads to a better estimate of one's own value for the good). I show that if bidders can misreport their departure times as well as their private signals, no reasonable auction satisfying a weak version of consumer sovereignty can be incentive compatible. For all other subsets of misreports I present conceptually simple incentive compatible auctions that respect consumer sovereignty.

I adopt a computational approach to design single-item revenue-optimal dynamic auctions with known arrivals and departures but (private) signals that arrive online. In leveraging the characterization of truthful auctions, I present a mixed-integer programming formulation of the design problem. I highlight general properties of revenue-optimal dynamic auctions in a parametrized example and study the sensitivity of prices and revenue to model parameters.

5.1 Introduction

In a dynamic market, being able to update one's value based on information available to *other* bidders currently in the market can be critical to having profitable transactions. In most dynamic auctions, however, the interaction between bidder valuations is limited and only indirect, through the competition between bids. If bidders have only partial information about the value of the item(s) to be sold then they may want a framework that allows them to "listen" to the market. Such a framework, well-established in the *static* auction theory literature [52, 63], is the model of *interdependent values* (IDV): each bidder's value is defined as an aggregation (expressed via a *valuation function*) of all the information in the market, but he only knows his own private information (his *signal*).

Suppose that every day a web content provider, such as the New York Times (NYT), auctions the right to have a banner ad appear next to the lead sports story the following day. The NYT must complete the auction by 4AM each morning, whereupon the auction for the subsequent day commences. One may imagine that advertisers will have a hard time knowing the value of having their own banner ad appear because this depends, in part, on the breaking news stories that day and on the associated user demographics. For this reason, an advertiser may naturally be interested in other bidders' information regarding the value of the banner ad. In fact, auctions are usually a means of obtaining information (the market price) for the seller. Continuing with the IDV model, each bidder's private information is encoded in his signal, but his value depends on his own signal as well as the signals of other bidders. For example, a bidder may express his value as \$1 (his signal) + one fifth the maximum (or the average) signal among all the other bidders. Each bidder becomes interested in buying the ad at a different time (construed as his arrival), and each bidder has a deadline (or departure) by which he must send a request to his ad design team.

For a second example consider a task allocation problem in which multiple agents are competing for the right to perform tasks in an dynamic, uncertain, environment. Together with an agent's private information about own capabilities and goals, each agent may have useful information about the domain; e.g., the reward or difficulty of performing a particular task. An interdependent value auction in this environment would allow each agent to condition the bid on this information, as reported by other agents, and enable each agent to ultimately report more accurately the value for a task.

I introduce a model of interdependent valuations in dynamic settings and identify a tension between interdependent values (IV), elementary misreports (EM), and dynamic dependencies (DD). A general negative result is established: no ex post IC and reasonable auction can satisfy all three of these properties. Elementary misreports model bidders as being able to report any signal, any departure but an arrival no earlier than their true one. An auction is reasonable if, loosely, there is some time beyond which any bidder can win for a high enough bid while the item remains unsold, but there is no bidder that can win for an arbitrarily low signal.

On the positive side, I design reasonable auctions when bidders can only misreport either their signal or departure (IV, ¬EM and DD). Previous work has identified reasonable auctions for bidders with interdependent values (IV, EM and ¬DD) in static environments [52, 27], as well as dynamic, private value environments (¬IV, EM, DD) [44]. Thus, any two of these three properties are possible but not all three together.

I leverage the characterization of incentive compatibility to design a revenue-optimal dynamic auction for IDV bidders. I propose a mixed integer program formulation and study its sensitivity to parameters in a simple example.

5.1.1 Related work

This work is the first that aims to bridge two distinct bodies of work: the economics literature on auctions for interdependent valuations and the computer science literature on online mechanisms.

In the economics literature there are two main models of interdependent valuations.

In the first model, each bidder's private information is his signal, but his valuation (aggregating his and others' signals) is publicly known. Dash et al. [28] present a multi-unit extension of the VCG mechanism for interdependent values [52]. Jehiel and Moldovanu [49] show that for incentive compatibility it is essentially necessary that bidders' private information has low dimensionality. Milgrom and Weber [63] study *affiliation* in bidders' signals, a form of correlation in which a higher signal for one bidder implies a priori higher signals for other bidders as well.

Dasgupta and Maskin [27] provide an IC (static) auction that is efficient when bidders' signals are one dimensional and introduce *contingent bids*. In this second model of interdependent values, a bidder does not report his signal, but rather a function describing how his value depends on the other bidders' *values* instead of signals. Ito and Parkes [48] instantiate this model to linear contingent bids and also extend it to single-minded combinatorial auctions.

Online (dynamic) mechanisms have received recent attention in computer science and operations research, motivated by the fast growing number of electronic commerce applications. Hajiaghayi et al. [44]'s setting is very similar to the one in this chapter, except that bidders' values are private. See also Lavi and Nisan [56], who establish an interesting negative result in private value environment with expiring goods and propose a relaxed notion of incentive compatibility as a work-around. The IC characterization in this chapter generalizes similar characterizations that are provided in earlier work for private value environments [56, 43].

5.2 Preliminaries

Assume that there is only one item for sale.[1] The item must be sold within T time units of the auction's start, which I will denote by time 1. Each bidder has an arrival and a departure time, and I call the interval between a bidder's arrival and departure the bidder's *activity interval*. I assume that any bidder i has a (single-dimensional) signal s_i and a valuation function v_i that aggregates all signals available in the market and determines i's actual value for the item. The signal s_i represents i's private information and v_i is a formula for i to compute his value from all the information in the market if i had this information. I will assume that the functional form of v_i for every bidder i is known.[2]

Signals are independent random variables, not necessarily from the same distribution, with values $0 \leq s_i < \infty$ unless otherwise specified. For time period $t \in [1, T]$ denote by $A^{\leq t}$ (respectively $\theta^{\leq t}$) the signals (respectively types) of bidders that arrive at or before t.

Similarly, if $a_i \leq t$, $A_{-i}^{\leq t}$ (respectively $\theta_{-i}^{\leq t}$) denotes the signals (respectively types) of all bidders except i that arrive at or before t.

Bidder i's value for the item is $v_i(s_i, A^{\leq d_i})$ at his departure and zero at any other time. Before his departure, the signal information from other bidders may not yet be revealed and thus i's value may be undefined. This is a significant change from the standard model for online, private-value auctions in which a bidder's valuation is known to him

[1] The characterization results generalize immediately to settings with known supply of multiple units of an item and unit-demand bidders.

[2] This assumption can be dropped, gaining practicality, if working in the contingent bids model instead, in which bidders state their values as a conditional value of the values of other bidders and a fixed point is determined [48].

throughout his time in the auction. A bidder's value is modeled as zero after departure to indicate that a bidder is uninterested in the item after his departure.

For example, $v_i(s_i, A^{\leq d_i}) = 0.8s_i + 0.2\max\{s_j : a_j \leq d_i\}$ if bidder i, whose signal is s_i, estimates the item's value to be the weighted average of s_i and the maximum signal of another bidder, where i's signal weighs four times more than the external signal. In a private values setting, $v_i(s_i, A^{\leq d_i}) = s_i$.

I consider a model of *elementary misreports* in which the misreports available are arbitrary signal misreports coupled with *late* arrival misreports and arbitrary departure misreports. I justify the assumption of late misreports of arrival by modeling the arrival time as the period at which a bidder first learns of the existence of the auction, or first realizes his demand for the item. Thus it is nonsensical to consider early reports. However, one may ask "why do bidders have arrivals at all if their value is only defined at departure?". Arriving in the system will grant a bidder the right to *compete* against other bidders. Moreover, since the arrival time is the time at which a bidder's signal is realized, it is desirable for the auction designer to provide incentives for a bidder to share that signal with other bidders.

I will call a bidder's *type* his private information: (arrival, departure; signal). Let bidder i's true type be $(a_i, d_i; s_i)$, i's reported type be $(a'_i, d'_i; s'_i)$ and let θ_{-i} denote a vector of types (not necessarily the true ones) from bidders other than i. An auction defines an allocation rule $\pi_i(a'_i, d'_i; s'_i, \theta_{-i}) \in \{0, 1\}$ (I will only consider deterministic auctions) to indicate whether or not bidder i is allocated the item, and a payment rule $p_i(a'_i, d'_i; s'_i, \theta_{-i}) \geq 0$ to define the payment made by bidder i. In a dynamic environment these must be online computable, i.e. $\pi_i(a'_i, d'_i; s'_i, \theta_{-i}) = \pi_i(a'_i, d'_i; s'_i, \theta_{-i}^{\leq d'_i})$ for all i, all θ, and similarly for payments. Moreover, payments must be collected by departure.

Bidders are modeled with quasilinear utilities: the utility of bidder i with type

$(a_i, d_i; s_i)$ when reporting $(a'_i, d'_i; s'_i)$ is

$$\pi_i(a'_i, d'_i; s'_i, \theta_{-i}) v_i(s_i, A_{-i}^{\leq d_i}) - p_i(a'_i, d'_i; s'_i, \theta_{-i}).$$

That is, utility is value minus price, where i's value for winning the item is aggregated by v_i from all true signals until i's departure. Note that i's true value for the item is $v_i(s_i, A_{-i}^{\leq d_i})$, whatever his report, where $A_{-i}^{\leq d_i}$ contains others' true signals.

Given this model of self-interest I restrict attention to incentive-compatible, online auctions such that there is an equilibrium in which every bidder chooses to report his true type immediately upon arrival into the auction:

Definition 36. *An auction is* incentive compatible *(IC) if, when the other bidders report their true types, the ex post utility of any bidder i is maximized if he reports his true type as well (i.e. truthful reporting is an* ex post Nash equilibrium*).*

This is *ex post* IC, meaning that when the other bidders are truthful (i.e. the set $\theta_{-i}^{\leq d_i}$ contains the *true* types of the other bidders arriving no later than d_i in the auction),

$$\begin{aligned} \pi_i(a_i, d_i; s_i, \theta_{-i}) v_i(s_i, A_{-i}^{\leq d_i}) \; - \; p_i(a_i, d_i; s_i, \theta_{-i}^{\leq d_i}) \quad \geq \\ \pi_i(a'_i, d'_i; s'_i, \theta_{-i}) v_i(s_i, A_{-i}^{\leq d_i}) \; - \; p_i(a'_i, d'_i; s'_i, \theta_{-i}^{\leq d'_i}), \end{aligned} \tag{5.1}$$

for all types $(a'_i, d'_i; s'_i)$ of bidder i such that $a'_i \geq a_i$.

Definition 37. *An auction is* individually rational *(IR) if the payment by the bidder winning the item is at most his true value for the item when all bidders (including the winner) report truthfully, and losing bidders pay zero.*

Note that the definitions of IC and IR do not specify anything about out-of-equilibrium behavior, thus being subject to the usual critiques of Nash equilibria, notably the question of how will bidders get to the Nash equilibrium and how to deal with multiple equilibria. On the other hand, the same critique serves to strengthen the negative result.

Let i be a bidder and $t \in [a_i, d_i]$. Let $s = A^{\leq d_i}$ for notational convenience. These following two assumptions are necessary for IC in static interdependent-value auctions (see [27]) and therefore I adopt them as well. The first assumption is that the bidders' valuations satisfy *v-monotonicity*: $v_i(s_i^+, s_{-i}) \geq v_i(s_i, s_{-i}) \, \forall\, i, \forall\, s_{-i}, \forall\, s_i^+ \geq s_i$. That is, a higher private signal cannot result in a lower value for the item. The second assumption is the *single crossing condition* (SCC): $\forall\, s$

$$\frac{\partial v_i(s)}{\partial s_i} > \frac{\partial v_j(s)}{\partial s_i}, \forall\, i, j : v_i(s) = v_j(s) = \max_{k \neq i,j} \{v_k(s)\}.$$

SCC requires that an infinitesimal change in bidder i's private signal influences i's value *more* than it influences the value of j if i's value is equal to j's and at least as high as the values of the other bidders.

5.3 Characterization of incentive compatibility

Let $\mathbb{A}$ be a dynamic auction for bidders with interdependent values. Consider bidder i and let $(\hat{a}_i, \hat{d}_i, \hat{s}_i)$ denote his reported type and fix the reports of other bidders $\theta_{-i}^{\leq \hat{d}_i}$. When referring to the types of other bidders $\theta_{-i}^{\leq \hat{d}_i}$, the superscript on i's reported departure $\hat{d}_i$ will be dropped if it can be inferred from the context.

I will show that the following conditions (to be denoted as CAD) are *necessary and sufficient* for $\mathbb{A}$ to be ex post IC and IR:

Critical signal: Let $s_i^c[\hat{a}_i, \hat{d}_i, \theta_{-i}] = \inf\{s_i : i \text{ wins in } \mathbb{A} \text{ reporting } (\hat{a}_i, \hat{d}_i, s_i)\}$

and ∞ if no such s_i exists (e.g. if the item has already been sold). Then when $\hat{s}_i > s_i^c[\hat{a}_i, \hat{d}_i, \theta_{-i}]$, bidder i must win in $\mathbb{A}$ at price $v_i\left(s_i^c[\hat{a}_i, \hat{d}_i, \theta_{-i}], A_{-i}^{\leq \hat{d}_i}\right)$.

Arrival monotonicity: $s_i^c[a_i^+, \hat{d}_i, \theta_{-i}] \geq s_i^c[\hat{a}_i, \hat{d}_i, \theta_{-i}]$,

$\forall\, a_i^+ \in (\hat{a}_i, \hat{d}_i]$, where $s_i^c[a_i, \hat{d}_i, \theta_{-i}]$ denotes as above the critical signal given i's arrival

a_i, departure $\hat{d}_i$ and the types of other bidders θ_{-i}.

Departure monotonicity: for $\forall\, d_i'$:

$$v_i\left(s_i^c[\hat{a}_i, d_i', \theta_{-i}^{\leq d_i'}], A_{-i}^{\leq d_i'}\right) \geq v_i\left(s_i^c[\hat{a}_i, \hat{d}_i, \theta_{-i}^{\leq \hat{d}_i}], A_{-i}^{\leq \hat{d}_i}\right)$$

Note that if d_i is fixed, i.e. bidder i cannot lie about his departure, then DEPARTURE MONOTONICITY trivially holds. Also note that s_i^c is allowed to fall with an early departure if the later signals are "bad news" for i.

Proposition 26. *The conditions* CAD *are necessary for IC and IR in an online, interdependent value environment.*

Proof. Let $\mathbb{A}$ be an auction having truthful reporting as ex post Nash equilibrium and consider bidder i.

Condition CRITICAL SIGNAL *is necessary.* Consider bidder i with true type $\theta_i = (a_i, d_i; s_i)$, fix the other bidders' types $\theta_{-i}^{\leq d_i}$ and assume they are truthful. To ease notation, I use s_i^c instead of $s_i^c[a_i, d_i, \theta_{-i}^{\leq d_i}]$. I prove that if the auction satisfies IR and IC, i must win whenever $s_i > s_i^c$. By s_i^c's definition, bidder i must lose when $s_i < s_i^c$.

Recall that v_i satisfies v-monotonicity i.e. $v_i(s_i^+, s_{-i}) \geq v_i(s_i, s_{-i})$ if $s_i^+ > s_i$. Note that i's price cannot be different for signals $s_i^1 \neq s_i^2$, with s_i^1 and $s_i^2 > s_i^c$: otherwise when i has the high-price signal he will be better off reporting the low-price one. By IR, since the other bidders are truthful, the price i pays must be at most $v_i(s_i^c, A_{-i}^{\leq d_i})$. But if i's price is strictly less than this, then when i has a signal barely under s_i^c, i will want to misreport a signal higher than s_i^c since the price is less than his value at d_i. Also, if bidder i loses for some $s_i' > s_i^c$ then i can misreport some $s_i^- \in [s_i^c, s_i']$ for which he wins. The price that i pays is $v_i(s_i^c, A_{-i}^{\leq d_i})$, which is less than $v_i(s_i', A_{-i}^{\leq d_i})$. Thus i would have an incentive to report s_i^-, contradicting IC.

Condition ARRIVAL MONOTONICITY *is necessary.* Assume the CRITICAL SIGNAL condition. Now, if ARRIVAL MONOTONICITY did not hold then bidder i could lower his price by reporting a late arrival, contradicting IC.

Condition DEPARTURE MONOTONICITY *is necessary.* Assume the CRITICAL SIGNAL condition so that the price for a report of $(\hat{a}_i, \hat{d}_i)$ has to be $v_i(s_i^c[\hat{a}_i, \hat{d}_i, \theta_{-i}^{\leq \hat{d}_i}], A_{-i}^{\leq \hat{d}_i})$. But if $v_i(s_i^c[\hat{a}_i, d'_i, \theta_{-i}^{\leq d'_i}], A_{-i}^{\leq d'_i}) < v_i(s_i^c[\hat{a}_i, \hat{d}_i, \theta_{-i}^{\leq \hat{d}_i}], A_{-i}^{\leq \hat{d}_i})$ for some $d'_i \neq \hat{d}_i$ then i will want to report d'_i instead of $\hat{d}_i$, contradicting IC. □

Say that an auction "allocates late" if the winning bidder is never allocated the item until his reported departure.

Proposition 27. *The conditions* CAD *are sufficient for IC and IR in an online, interdependent value environment when the v-monotonicity and SCC properties also hold and when the auction allocates late.*

Proof. Let i be a bidder, (a_i, d_i, s_i) his true type and assume that the other bidders are truthful. I will prove that if conditions CAD hold, then it is in the best interest of bidder i to be truthful as well. Note that IR trivially holds because of the way s_i^c is defined.

Whatever the signal report and for all misreports of departure, for i to be strictly better off by reporting a late arrival a_i^+, v-monotonicity implies that $s_i^c[a_i^+, \cdot] < s_i^c[a_i, \cdot]$. But this is specifically what ARRIVAL MONOTONICITY precludes. I can henceforth assume that i reports his true arrival. I now show that bidder i cannot be better off reporting some $d_i^- < d_i$, whatever signal i reports. If it were, then the price when reporting d_i^- must be less than the price for d_i (which may be ∞):

$$v_i\left(s_i^c[a_i, d_i^-, \theta_{-i}^{\leq d_i^-}], A_{-i}^{\leq d_i^-}\right) < v_i\left(s_i^c[a_i, d_i, \theta_{-i}^{\leq d_i}], A_{-i}^{\leq d_i}\right).$$

But this contradicts the DEPARTURE MONOTONICITY condition. In addition, it is never useful to report a late departure because the auction allocates late and this would have zero value. Thus I can assume that i reports his true departure and true arrival, whatever his signal misreport.

Bidder i's price does not depend on his signal and he can maximize his chances of winning by reporting his true signal. If i does not win reporting his true signal then $s_i^c > s_i$ i.e. his value is less than his price so i does not *want* to win. □

Theorem 19. *The conditions* CAD *are necessary and sufficient for IC and IR in an online, interdependent value environment, and when the auction allocates late.*

In private values settings $v_i(s_1, \ldots, s_n) = s_i$ and the conditions CAD amount to the existence of a critical value function $v_i^c[a_i, d_i, \theta_{-i}]$ such that $v_i^c[a_i^+, d_i^-, \theta_{-i}] \geq v_i^c[a_i, d_i, \theta_{-i}]$ if $[a_i^+, d_i^-] \subset [a_i, d_i]$. That recovers Theorem 5 from Hajiaghayi et al. [43].

The ability (due to interdependence) of changing the values of other bidders by temporal misreports makes the CRITICAL SIGNAL property harder to achieve than it may first seem and definitely harder than the critical value property from private values settings.

5.3.1 Obvious winner acceptance

In this section I use the IC characterization to show that any IC auction satisfying a weak version of consumer sovereignty (CS) must in fact transition at some time into a state in which it will subsequently sell (at a low price) to the first bidder to arrive, whatever his signal. Thus, no IC auction in the interdependent value, dynamic auction environment with elementary misreports can be *reasonable*, when reasonable is construed as providing some variant on CS[3] coupled with an absence of this "oblivious selling" property.

In private value settings, an auction satisfies consumer sovereignty if, with arbi-

[3]With sufficient uncertainty about the future, not selling to some bidder when it has a high signal (and therefore a high value as well) can significantly hurt the revenue or the efficiency of an auction. I leave it to the reader to judge the performance of an auction that sells to a bidder whatever he reports.

trary fixed values of the other bidders, any bidder can win provided he reports a high enough value. The following definition applies to an online setting, where one also conditions on the item still being available:

Definition 38. *An auction satisfies* obvious winner acceptance *(OWA) if there is some time* $\boldsymbol{T}$ *(the OWA* cut*) with the following property: whenever some bidder* w*'s (with* $a_w \geq \boldsymbol{T}$*) activity interval is* disjoint *from any other bidder's activity interval there is some finite value* $\boldsymbol{S}_w$ *(that can depend on the other bidders' signals) such that* w *wins the item with any signal at least as high as* $\boldsymbol{S}_w$ *(see Figure 5.1).*

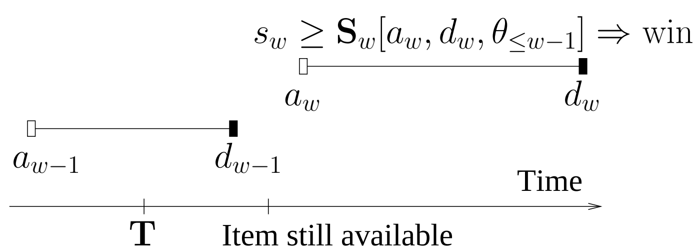

Figure 5.1: If *obvious winner acceptance* holds then w (whose interval is disjoint from the other bidders' intervals) must win for a high enough signal (assuming the item was not sold earlier).

The OWA condition requires that there is some time, past which if w faces no active competition then for some (high enough) signal w must win. Bidder w is in this case the "obvious winner". Note that if the OWA cut is before bidder w's $[a_w, d_w]$ interval and the auction is IC then the threshold $\mathbf{S}_w[a_w, d_w, \theta_{\leq w-1}]$ required by OWA must be $s^c_w[a_w, d_w, \theta_{\leq w-1}]$.

The following claim shows that any IC auction satisfying OWA must unconditionally sell, after the OWA cut, to a bidder whose arrival makes him the only active bidder (i.e. bidder 3 in the setting of Figure 5.2), whether or not some other bidder later arrives during his activity interval.

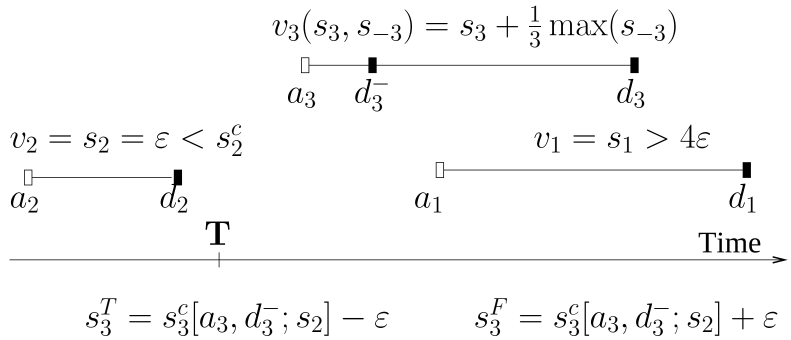

Figure 5.2: Illustrating Claim 28: any IC auction satisfying OWA must sell to agent 3 irrespective of his signal when 3 reports departure d_3^-.

Definition 39 (sensitive)**.** *Bidder i is said to be sensitive to the signal of bidder j if, for any s_i, s_i', s_{-ij}, there exists s_j such that $v_i(s_i; s_{-ij} \cup \{s_j\}) \geq v_i(s_i'; s_{-ij})$.*

That is, bidder i is sensitive to bidder j's signal if, whatever two values of bidder i that do not use j's signal, any one of them can be made higher than the other by allowing it only to also use some (perhaps sufficiently high) signal of j. This property holds for example if $v_i = \alpha s_j + f(s_{-j})$, where $\alpha > 0$ and v_i uses some fixed value for s_j if it is not (yet) available.

Proposition 28. *Let 2, 3 and 1 denote three bidders in order of arrival, and with $d_2 < a_3$ but $a_1 < d_3$. Suppose the OWA cut occurs in $\boldsymbol{T} \in (d_2, a_3)$, that the item is still available at $\boldsymbol{T}$, and that bidder 3 is sensitive to bidder 1's signal. Any IC auction with both signal and early departure misreports must allow bidder 3 to win for any reported signal and his price is that defined by the value for his minimal signal.*

Proof. Let $\mathbb{A}$ be an auction satisfying OWA and IC. Since $\mathbb{A}$ is IC, $\mathbb{A}$ must satisfy the CRITICAL SIGNAL condition. Let $s_3^c[d, s_{-3}]$ be shorthand for $s_3^c[a_3, d, s_{-3}]$. Note that $s_3^c[d_3^-, s_2] < \infty$ by OWA. Suppose $s_3^c[d_3^-, s_2] > 0$. Let $s_3^T < s_3^c[d_3^-, s_2]$: bidder 3 would not win if reporting truthfully. consider what happens for an s_1 such that $v_3(s_3, s_1, s_2) > v_3(s_3^c[d_3^-, s_2], s_2)$ (such an s_1 must exist since 3 is sensitive to 1's signal). Bidder 3 can simply report a signal higher than $s_3^c[d_3^-, s_2]$ and win: bidder 3 will be happy since the price he pays is less than his true value at d_1. Therefore $s_3^c[d_3^-, s_2] = 0$. This holds for all $d_3^- \in (a_3, a_1)$. But then, if 3 reports a departure after a_1, an IC auction must give the item to him as otherwise 3 can do better by misreporting a departure before a_1. In conclusion, the auction must sell unconditionally to bidder 3. □

This negative result immediately generalizes in the following sense: if at some (sufficiently late) time in the auction, there is no active bidder then the item must be allocated to the first arriving bidder f whatever his signal if f is sensitive to signals of future bidders.

Thus OWA, a reasonable requirement, is shown to imply this non-competitive property for a large set of scenarios. The requirement of incentive compatibility transforms the requirement of selling to the obvious *winner* into obliviously selling to what can be the obvious *loser*, e.g. when bidder 3 has a very low signal.

Theorem 20. *There is no reasonable, IC and IR auction in the interdependent values, online environment when bidders can misreport arrival (later only), departure and signal.*

5.3.2 Limited classes of misreports

As already seen, reasonable IC auctions are not possible even if bidders only have elementary misreports at their disposal. I now present reasonable auctions for environments in which bidders cannot report both a different signal and an early departure. In any auction in the sequel, I define the critical signal for any bidder to be ∞ after the item has been sold.

In the auctions of Figure 5.3 and 5.5 when I consider different arrivals a_i' for bidder i I will just consider the counter-factual world where i arrives at a_i' instead of a_i. Thus, when considering arrival a_i', i's signal will be available to bidders active at a_i' or later.

I will assume an upper bound Δ on the patience of any bidder ($d_i - a_i \leq \Delta, \forall\, i$). This assumption could be dropped, but then one would have to look back as far as the start of the auction for counter-factual arrivals of agent i.

Late arrivals and signal misreports

In this subsection I assume that bidders cannot lie about their departure. This can happen if, for example, the seller, sets a (possibly different for each bidder) deadline for selling the item. Thus, condition DEPARTURE MONOTONICITY is trivially satisfied.

The auction in Figure 5.3 sells to the highest bidder i, provided that no different arrival of i would have made another bidder j win early. Why worry about this? In retrospect (i.e. *ex post*), if bidder i loses to j and j's value depends significantly on i's signal and great news arrive for i after j's departure, i may be tempted to stop j from winning by hiding his signal from i via a late arrival misreport.

Proposition 29. *The auction in Figure 5.3 is IC and reasonable in interdependent, online environments, when v-monotonicity and SCC hold and when bidders can only misreport signals and arrival times but not departure times.*

Proof. Departure is fixed, so DEPARTURE MONOTONICITY is trivially satisfied.

The only way i's reported arrival can matter is if i has an arrival $a_i < d_j$ for some bidder j that would win (have the highest value) only if i were to report his true arrival a_i. i can instead report a false arrival of $a_i^+ > d_j$ and thus j would not win anymore. But then $s_i < s_i^{thr}[d_j] < s_i^*$, so, when misreporting, i would pay more than his true value. The auction is reasonable since the winner must outbid his competitors. □

The auction satisfies CS (a stronger property than OWA) if v_i's satisfy the following

Decision times for allocation: Bidders' departures.

Pricing rule: Fix bidder i. For all j such that $d_j \in [\hat{d}_i - \Delta, \hat{d}_i)$ let $s_i^{thr}[d_j]$ be the infimum of all signals s with the following property: whenever i reports some signal $s_i' \geq s$ along with an arrival of d_j, at d_j there is some active or departed bidder (that may be i) with a value strictly higher than j's value. Let

$$s_i^{\mathrm{PA}}[\hat{d}_i] = \max_{\hat{d}_i - \Delta \leq d_j < \hat{d}_i} s_i^{thr}[d_j]$$

Let $s_i^{hi}[\hat{d}_i]$ be the minimum signal i could have reported such that his value is the highest across all active and departed bidders at $\hat{d}_i$. Let

$$s_i^*[\hat{d}_i] = \max\{s_i^{\mathrm{PA}}[\hat{d}_i], s_i^{hi}[\hat{d}_i]\}$$

Sell to i only if $s_i > s_i^*[\hat{d}_i]$ ($s_i^*[\hat{d}_i]$ is i's critical signal)

Figure 5.3: IC and reasonable auction for the case of late arrival and signal misreports.

natural condition:

$$\forall\, i, \forall\, s_{-i}, \exists s_i^0 \text{ s.t. } v_i(s_i^0, s_{-i}) > \max_{j \neq i}\{v_j(s_i^0, s_{-i})\} \tag{5.2}$$

Indeed, by SCC and v-monotonicity, $\forall\, s_i \geq s_i^0$, i's value is the highest one, i.e. i wins ($s_i^*[d_i]$ is at most s_i^0).

When bidders cannot misreport arrivals either, this auction becomes the interdependent (generalized) second-price auction in [52].

Example for the truthful auction in Figure 5.3

In the example of Figure 5.4: $v_1(s_1, s_{-1}) = \frac{3}{5}s_1 + \frac{2}{5}\max(s_{-1})$, $v_2(s_2, s_{-2}) = \frac{13}{15}s_2 + \frac{2}{15}\text{avg}(s_{-2})$, $v_3(s_3, s_{-3}) = \frac{3}{5}s_3 + \frac{2}{5}400$ and signals $s_1 = 600, s_2 = 690, s_3 = 900$. If 3 reports

a_3

- at d_1: $v_1 = \frac{3}{5}600 + \frac{2}{5}690 = 636$; $v_2 = \frac{13}{15}690 + \frac{2}{15}600 = 678$. Thus, 1 does not win since 2 has a higher value.

- at d_2: $v_1 = \frac{3}{5}600 + \frac{2}{5}690 = 636$; $v_2 = \frac{13}{15}690 + \frac{2}{15}(\frac{600+900}{2}) = 698$; $v_3 = 700$. Thus, 2 does not win since 3 has a higher value. The values do not change at d_3 and therefore 3 wins the item at d_3.

Contrast this with what would happen if 3 reported $a'_3 = d_3 - \Delta$: $v_1 = \frac{3}{5}600 + \frac{2}{5}900 = 720$; $v_2 = \frac{13}{15}690 + \frac{2}{15}(\frac{600+900}{2}) = 698$, $v_3 = 700$. Thus, 1 wins since he has the highest value at d_1.

This is why the auction has the extra wrinkle of checking all possible arrivals for the candidate winner. With no check and simply selling to the bidder with the highest value, 3 would win when reporting a_3. But then, when 3 has true arrival $a'_3 = d_3 - \Delta$, 3 can do better by reporting a late arrival of a_3, hiding his signal from 1. Bidder 3 can win by reporting $s_3 > 1000$: then $v_3 > v_1$ for any arrival of 3 ($v_3 > v_2$ as well). Then one can set s_3^c to 1000 (or higher).

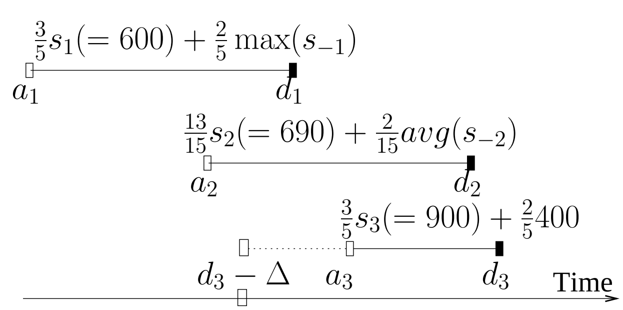

Figure 5.4: Example setting for the auction in Figure 5.3. Without the check for all possible arrivals, bidder 3 can benefit by reporting a_3 (thus hiding his signal from 1) when his true arrival is $d_3 - \Delta$.

Decision times for allocation: Bidders' arrivals

Pricing rule: Fix bidder i. For all $a_j \in [\hat{d}_i - \Delta, \hat{a}_i]$ let $s_i^{hi}[a_j]$ be the minimum signal i could have reported along with an arrival of a_j making i's value the highest across all active and departed bidders at a_j. Let

$$s_i^* = \max_{\hat{d}_i - \Delta \leq a_j \leq \hat{a}_i} s_i^{hi}[a_j]$$

If $s_i \geq s_i^*$, reserve the item for agent i but only give him the item at d_i. Charge i

$$\min_{\hat{a}_i \leq t \leq \hat{d}_i} v_i(s_i^*, A_{-i}^{\leq t}).$$

Figure 5.5: IC, reasonable auction for the case of late arrival and early departure misreports.

Departures and late arrivals misreports

In this subsection I assume that bidders cannot lie about their signals, but can misreport a late arrival and an early departure. This may be of relevance in a dynamic planning problem where the goal is to take the best decision given partial (local) information from each bidder, that may have limited (e.g. battery-powered) life.

The auction in Figure 5.5 sells to the highest bidder i, provided i would still be the highest bidder for any possible true arrival given his reported interval $[\hat{a}_i, \hat{d}_i]$. Since $d_i - a_i \leq \Delta$ and only early departure misreports, $a_i \in [\hat{d}_i - \Delta, \hat{a}_i]$. i's price is taken as the minimum value given his critical signal s_i^* across all his reported activity interval. This motivates i to report an interval as large as possible, i.e. his true one. Furthermore, reporting a tighter interval only has the effect of raising s_i^*, which is the max over more a_j's than when being truthful. Thus,

Proposition 30. *The auction in Figure 5.5 is IC and reasonable in interdependent, online environments, when v-monotonicity and SCC hold and when bidders can only misreport late* *arrival and early departure times but not signals.*

As before, the auction satisfies CS if v_i's satisfy condition (5.2).

5.4 Revenue maximization with known intervals

Recall the example from the introduction, in which the New York Times auctioned a banner ad slot. Assume as well that the NYT has contracted with various advertisers, where the contract defines a fixed bidding interval wherein the advertiser can choose to bid for the right at the same time each day (the "arrival" time of the bid) and is guaranteed a response by some subsequent time (the "departure" time of the bid.) Only the bid of the advertiser, and indeed whether or not the advertiser will choose to bid, is uncertain.

I first consider a naive generalization of the optimal *static* and IC auction for interdependent bidders [15] and point out that the solution obtained fails to satisfy IC constraints. The reason is that, in retrospect, a bidder i will regret reporting truthfully in the following scenario: i can misreport his signal changing the value or price of an earlier bidder h such that h is now precluded from winning and i is going to win (maybe because the future signals turn out to be "favorable" for i to misreport).

To determine the optimal, revenue-maximizing auction in this dynamic IDV environment I adopt a mixed-integer programming (MIP) formulation and follow the framework of automated mechanism design [23], building on the heritage of the Myerson [64] approach. For practical scalability I require that the interdependencies between bidders are of bounded degree, that the designer is able to constrain the number of signals that must be propagated from earlier periods into defining the price of bidders in later periods, and that a coarse discretization of signals can be tolerated. The formulation is illustrated in a simple, three bidder scenario. Based on this formulation I can compare the revenue from this, IC formulation with the one obtained using the naive generalization and the one obtained using a clairvoyant approach.

Related work

Branco [15] studies *static* revenue-optimal auctions for IDV bidders and shows that under a certain regularity condition, an asymmetric critical signal-based auction is optimal (I review this auction in Subsection 5.4.2). The regularity condition is satisfied if bidders' valuations are increasing and concave in their own signal and the signals' distribution has a non-decreasing hazard rate. Segal and Toikka [74] also extend Myerson [64]'s revenue equivalence theorem and optimal auction design to dynamic settings, but their formulation appears to be restricted to private values.

Aoyagi [1] investigates optimal pricing schemes in dynamic settings with IDV bidders, without considering incentive compatibility. A bidder infers his value indirectly, from the decisions of other bidders. If a previous bidder j accepted (respectively rejected) the price offered by the seller, then a current bidder i's estimation of j's signal will increase (respectively decrease), leading to a corresponding change in i's *value*. Aoyagi shows that for any simultaneous selling scheme, there exists a sequential one with at least as high a revenue.

Hajiaghayi et al. [44, 43] provide competitive mechanisms for selling one or more goods in a dynamic environment, but they model bidders' values as private. See Parkes [69] for a recent survey on mechanism design in dynamic environments. The strategy of finding optimal-revenue mechanisms through search is in the spirit of *automated mechanism design* [23]. However, rather than impose IC constraints directly, this formulation amounts to an *informed search*, since one searches only for critical signals that support a truthful allocation policy.

5.4.1 Preliminaries

The auctioneer has uncertainty about bidders' signals s and models each s_i as an independent draw from a distribution on non-negative values with cumulative distribution function (cdf) F_i and probability density function (pdf) f_i. I consider only deterministic auction rules and assume for simplicity that all bidders have disjoint departures.[4]

Without loss of generality I can focus on auction protocols that sell to a bidder upon his departure. Not only does this ensure that a bidder's own value is known at the time of his allocation but this allows the auctioneer to gain maximal information about other demand in the market. I assume that v_i is differentiable with respect to any bidder j's signal s_j and: (1) *v-monotonicity*: $v_i(s_i^+, s_{-i}) \geq v_i(s_i, s_{-i}) \, \forall \, i, \, \forall \, s_{-i}, \, \forall \, s_i^+ \geq s_i$. That is, a higher private signal cannot result in a lower value for the item; (2) the *single crossing condition* (SCC): an infinitesimal change in bidder i's private signal influences i's value *more* than it influences the value of j if i's value is equal to j's and at least as high as the values of the other bidders. Any non-trivial incentive compatible auction in static IDV environments must satisfy (1) and (2) (see [52]).

Incentive compatibility characterization

Consider a dynamic auction for IDV bidders that can only misreport their signal. In Section 5.3 I have established three conditions that are necessary and sufficient for IC in dynamic, IDV auctions. Two of the conditions require that a bidder's price does not go down if he misstates his interval – they are trivially satisfied in the known-interval setting considered. The third condition, adapted to a no-interval-misreports domain, requires:

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

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

This implies that the allocation rule is monotonic in the bidder's signal. The existence of an unconditional critical signal generalizes the "critical-value" concept in private-value settings, where a bidder wins iff his value is higher than the critical value, which is also the price he pays [69].

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

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

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

- i loses (before departing) to a competitor h when reporting signal s_i, but
- i can misreport some signal s_i' causing h to lose (e.g. if his critical signal goes from below to above s_h when i's signal changes from s_i to s_i'), and resulting in i now winning at a price less than his true value for the item.

To address this one must combine conditional critical signals $c_i[s_{-i}]$ with additional inter-temporal constraints:

Theorem 21. *A dynamic auction in the known-interval, IDV model is IC if and only if it has conditional critical signals with the property that there are no signals s_{-i} (of bidders arriving before i), s_i and s'_i such that all of the following hold: (a) $c_{j_0}[s_i, s_{-ij_0}] \leq s_{j_0}$ for some $j_0 < i$; (b) $c_j[s'_i, s_{-ij}] > s_j$ for all $j < i$; and (c) $c_i[s_{-i}] \leq \min\{s_i, s'_i\}$.*

The constraints in Theorem 21 are referred to as *inter-temporal IC constraints* (ITIC). Given this, there is never an instance for which some bidder i loses when reporting true signal s_i (a), could have prevented all earlier bidders from winning for some $s'_i \neq s_i$ (b), and wins for report s'_i and with a critical signal less than his true signal and thus a payment less than his true value (by v-monotonicity). Conditional critical signals that satisfy ITIC become *unconditional*: i wins if and only if his signal is at least $c_i[s_{-i}]$.

5.4.2 Special cases

I now present two simple dynamic IDV environments for which revenue optimal auctions can be easily constructed. I first review Branco's [15] solution for non-dynamic IDV environments and then provide a multi-step optimization formulation for the case of disjoint intervals.

Revenue-optimal static auctions

Definition 40. *If a bidder i's valuation is $v_i(s_i, s_{-i})$ then i's* virtual valuation *is*

$$\tilde{w}_i(s_i, s_{-i}) = v_i(s_i, s_{-i}) - \frac{\partial v_i}{\partial s_i}(s_i, s_{-i})\frac{1 - F_i(s_i)}{f_i(s_i)} \tag{5.3}$$

For example, if $i = 3$, signals are distributed uniformly on $[0, 1]$ and $v_3(s_1, s_2, s_3) = s_3 + \frac{s_1}{4} + \frac{s_2}{4} + \frac{1}{4}$ then $\tilde{w}_3(s_1, s_2, s_3) = s_3 + \frac{s_1}{4} + \frac{s_2}{4} + \frac{1}{4} - 1\frac{1-s_3}{1} = 2s_3 + \frac{s_1}{4} + \frac{s_2}{4} - \frac{3}{4}$. Branco [15] provides a revenue-optimal static auction for IDV bidders if i's *virtual* valuation is increasing in s_i, for all i. This holds in particular if bidder valuations are increasing and concave in their own signal and the distribution of each bidder's signal has a non-decreasing hazard rate ($\frac{f_i(s_i)}{1-F_i(s_i)}$). Branco's auction generalizes Myerson's [64] revenue-optimal private-value auction, and is thus based on the fact that in equilibrium, *virtual valuation = revenue.* This insight suggested the virtual valuation-based heuristic for the non-clairvoyant mechanism in Subsection 5.4.4. In Branco's solution, the bidder with the highest virtual valuation $\tilde{w}_i$ wins, but only if $\tilde{w}_i$ is non-negative. The winner pays his value computed at the lowest signal for which he still wins.

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

Disjoint intervals

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

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

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

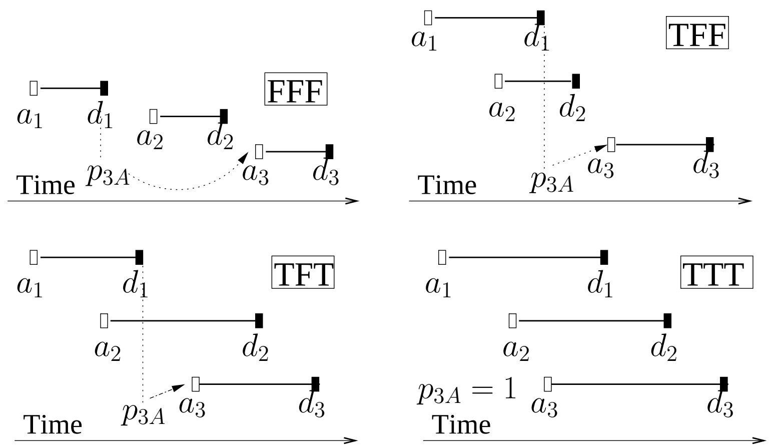

Figure 5.6: Bidder configuration as known a priori (at a_1) in each of the four scenarios. At d_1 or d_2, if 3's scheduled arrival a_3 is later, the probability that 3 will arrive at all is p_{3A}. If bidder 3 arrives, then his interval is *known* to be $[a_3, d_3]$.

selling to bidder j at a price defined by critical signal c and selling to future bidders under the *optimal* critical signal schedule. For the last bidder, $c_n \in \text{argmax}_c\{(1-F(c))\cdot v_n(c, s_{<n})\}$. If the n-th bidder's valuation is increasing and concave in his own signal then $\tilde{w}_n(c_n, x_{<n}) = 0$.

5.4.3 Working with a specific problem instance

For the remainder of the chapter I will work with a set of four specific 3-bidder scenarios as shown in Figure 5.6. In all scenarios bidders' arrivals and departures are known and fixed, with a slight variation for the scenarios in which bidder 3 is supposed to arrive later than time d_1. In those scenarios, if bidder 3 has not arrived yet, the auctioneer only knows the (correct) probability of bidder 3 arriving: $0 \leq p_{3A} \leq 1$.

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

For simplicity, I will model the signal of any bidder as being *uniformly* distributed on $[0, 1]$. In the continuous Branco formulation, the distribution-dependent part of i's virtual valuation (see Eq. (5.3)) will then be $1 - s_i$.

MIP formulation

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

The signal space of each bidder b is discretized such that s_b can take on values $s_{b1}, \ldots, s_{bm_b}$. I will hereafter assume that bidders' valuations are linear in the signals available: $v_i(s_i, s_{-i}) = s_i + \sum_{j \neq i} v_{ij} s_j$. In Section 5.4.4 specific numeric values will be used for the v_{ij} weights, i.e. for the weight that bidder i assigns to the signal of some bidder $j \neq i$. This discretized type space provides an approximation to a continuous one.

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

Decision variables As seen in Theorem 21, any IC auction can be defined by conditional critical signals when coupled with ITIC constraints. The MIP includes decision variables c_i, instantiated for each bidder i on the discretized signals of the other relevant bidders, and defining the conditional critical signals (just "critical signals" hereafter). A bidder gets the item if no earlier bidder won the item and the bidder's signal satisfies $c_i \geq c_i(s_{-i})$ upon departure. Given critical signals, binary variables λ, μ and ν encode whether each bidder

(1, 2 and 3 respectively) does *not* win the item:

$$\begin{array}{rr}
\lambda_{ij} = 1 \iff & s_{1i} < c_1(s_{2j}) \\
\mu_{ijk} = 1 \text{ iff } s_{2j} < c_2(s_{1i}, s_{3k}) & \mu_{ij}^{\neg 3} = 1 \text{ iff } s_{2j} < c_2^{\neg 3}(s_{1i}) \\
\nu_{ijk} = 1 \iff & s_{3k} < c_3(s_{1i}, s_{2j})
\end{array} \tag{5.4}$$

Within the MIP, one can capture logic such as $\lambda_{ij} = 1 \iff s_{1i} < c_1(s_{2j})$ via a linear constraint such as $-M\lambda_{ij} \leq s_{1i} - c_1(s_{2j}) < M(1 - \lambda_{ij})$ where a "big M" is adopted, and set to the smallest constant that can be proved to be larger than the maximal absolute value of $s_{1i} - c_1(s_{2j})$. Monotonicity constraints on indicator variables are $\lambda_{ij} \leq \lambda_{i-1,j}$, $\mu_{ijk} \leq \mu_{i,j-1,k}$, $\mu_{ij}^{\neg 3} \leq \mu_{i,j-1}^{\neg 3}$, and $\nu_{ijk} \leq \nu_{i,j,k-1}$.

These indicator variables are used both in the ITIC constraints and in the objective function. For example, λ_{ij} indicates whether bidder 1 does not win the item when his signal is s_{1i} and bidder 2's signal is s_{2j}. Note that in TFT bidder 1's critical signal depends on s_2, bidder 3's on (s_1, s_2) and bidder 2's depends on whether or not bidder 3 arrives. Variables $c_2(s_{1i}, s_{3k})$ and μ_{ijk} capture the behavior of the auction in respect to bidder 2 when bidder 3 arrives while variables $\mu_{ij}^{\neg 3}$ and $c_2^{\neg 3}(s_{1i})$ are for the case without bidder 3.[5]

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

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

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

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

which, using the variables in Eq. (5.4), can be written as

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

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

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

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

  which is easily expressed into linear constraints as above.

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

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

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

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

$$\begin{cases} [c_2(s_{1i}, s_{3k}, s_{4l}) \leq s_{2j} \textbf{ or } c_3(s_{1i}, s_{2j}, s_{4l}) \leq s_{3k}] \text{ and} \\ s_{2j} < c_2(s_{1i}, s_{3k}, s_{4l'}) \text{ and} \\ s_{3k} < c_3(s_{1i}, s_{2j}, s_{4l'}) \text{ and} \\ c_4(s_{1i}, s_{2j}, s_{3k}) \leq s_{4l} \text{ and } c_4(s_{1i}, s_{2j}, s_{3k}) \leq s_{4l'} \end{cases} \tag{5.5}$$

which can also be easily expressed in linear constraints.

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

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

where

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

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

Formulation size The following problem characteristics define the size of the MIP in a general scenario:

- D^+ ($\leq m$): the maximum number of earlier bidders' signals that can influence a bidder's value

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

By generalizing Eq. (5.4), one can see that for each bidder i, each discrete signal s_i and each possible signal tuple of bidders that influence i's value, there is a variable controlling whether s_i is smaller than the critical signal c_i computed at that particular tuple. Since in the tuple there can be at most D^+ bidders, *the number of variables is* $O(nS^{D^++1})$.

To analyze the number of constraints, one needs to extrapolate Eq. (5.5) since the bulk of constraints will be of this form. Each such ITIC constraint must ensure that no bidder i ($i = 4$ in Eq. (5.5)) can misreport his signal such that: all earlier bidders (2 and 3 in Eq. (5.5)) who should be winning when i is honest do not win anymore and furthermore, i wins. There are n bidders in total. Each bidder i can influence the value of at most D^- earlier bidders and each of those other bidder's ITIC constraint depends on his own signal, together with perhaps m other signals, where m is the maximal number of other signals that can influence his critical signal. Thus, *the number of constraints is* $O(nD^-S^{m+1})$. Note that there are no ITIC constraints when $D^- = 0$, which occurs either in the IDV but disjoint interval case or in the non-disjoint but private-values case. In both of these cases it is sufficient for IC to simply formulate the decision problem as one of setting critical signals.

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

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

5.4.4 Instantiation

I now initiate the simulation study, by instantiating particular scalars for the valuation model described above.

Let $s_1, s_2, s_3 \sim U[0,1]$ and assume the following valuations (whose choice will be motivated shortly):

$$v_3(s_1, s_2, s_3) = s_3 + \frac{s_1}{4} + \frac{s_2}{4} + \frac{1}{4}$$

$$v_2(s_1, s_2, s_3) = \begin{cases} s_2 + \frac{s_1}{4} + \frac{s_3}{4} + \frac{1}{4} & \text{,if } d_3 \leq a_2 \\ s_2 + \frac{s_1}{2} + \frac{1}{4} & \text{, otherwise} \end{cases}$$

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

The virtual valuations are

$$\tilde{w}_1(s_1, s_2, s_3) = \begin{cases} 3s_1 - \frac{5}{4}, & \text{, if 1 cannot see 2 or 3} \\ 2s_1 + \frac{s_2}{2} - \frac{3}{4} & \text{, if 1 can see 2, but not 3} \\ 2s_1 + \frac{s_2}{4} + \frac{s_3}{4} - \frac{3}{4} & \text{, if 1 can see both 2 and 3} \end{cases}$$

$$\tilde{w}_2^{\neg 3}(s_1, s_2, s_3) = 2s_2 + \frac{s_1}{2} - \frac{3}{4}$$

$$\tilde{w}_2(s_1, s_2, s_3) = 2s_2 + \frac{s_1}{4} + \frac{s_3}{4} - \frac{3}{4}$$

$$\tilde{w}_3(s_1, s_2, s_3) = 2s_3 + \frac{s_1}{4} + \frac{s_2}{4} - \frac{3}{4}$$

The valuations were chosen to be symmetric and such that:

- they depend linearly on the signals available,
- $\frac{\partial v_i}{\partial s_i} > \frac{\partial v_i}{\partial s_j} \, \forall \, i, j$, implying the v-monotonicity and SCC conditions
- $v_1 = v_2 = v_3$ and $\tilde{w}_1 = \tilde{w}_2 = \tilde{w}_3$ when $s_1 = s_2 = s_3$ in *all* scenarios, except $\tilde{w}_1 < \tilde{w}_2 = \tilde{w}_3$ in scenario FFF.

The final property ensures that there is no a priori bias between bidders and it is meant to facilitate the analysis of the interaction of interdependent values and uncertainty about the future. One cannot have, however, $v_1 = v_2 = v_3$ and $\tilde{w}_1 = \tilde{w}_2 = \tilde{w}_3$ when $s_1 = s_2 = s_3$ in *all* scenarios unless values are private. For the weights chosen, $\mathbb{E}[\tilde{w}_1(s_1)] = 0.25$ in scenario FFF, less than 0.5, the expected value of any bidder's virtual valuation in all other scenarios.

Experimental setup

I compare the performance of the MIP-based auction with two additional auctions as summarized in Table 5.1. The *continuous, clairvoyant* (ContCV) auction provides a

Table 5.1: Auctions used for evaluation (left), as characterized by whether or not the auctioneer is clairvoyant (CV), whether or not a discretization is imposed when designing the auction (Cont == "continuous," no discretization), and whether or not the auction is IC. The expected revenue is summarized by scenario (right) in an environment where bidder 3 always arrives.

Auction	CV	Cont	IC
ContCV	$\surd$	$\surd$	$\surd$
Cont NonCV	X	$\surd$	X
MIP policy	X	X	$\surd$

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

best-case revenue. Here, I allow the auctioneer to observe the signals of all bidders from the start, regardless of the scenario. Bidder i's critical signal is then computed as the least signal for which i has the highest non-negative virtual valuation $\tilde{w}_i$ (recall that $\tilde{w}_i$ depends on the scenario). In other words, ContCV implements Branco's auction as if the auctioneer could have perfect knowledge of all signals.

The *continuous, non-clairvoyant* (Cont NonCV) auction is a naive generalization of Branco's offline auction that is not IC in general. At each departure d_i, the item is sold to i iff i's virtual valuation $\tilde{w}_i$ is higher than the maximum of the expectations of the virtual valuations of bidders still to arrive and zero (as in Eq. (5.6)). If any bidder j has already departed, $\tilde{w}_i$ is *not* compared with $\tilde{w}_j$. For example, in scenario TFT, this solution requires selling to:

$$\begin{aligned}
&\text{Bidder 1,}\ \ \text{if } \tilde{w}_1 \geq (1-p_{3A})\max(\tilde{w}_2^{\neg 3},0) + p_{3A}\mathbb{E}[\max(\tilde{w}_2,\tilde{w}_3,0)] \\
&\text{Bidder 2,}\ \ \text{if } \tilde{w}_2^{\neg 3} \geq 0 \text{ when 3 does not arrive and haven't sold to 1} \\
&\text{Bidder 2,}\ \ \text{if } \tilde{w}_2 \geq \max(\tilde{w}_3,0) \text{ when 3 arrives and haven't sold to 1} \\
&\text{Bidder 3,}\ \ \text{if } \tilde{w}_3 \geq 0 \text{ when 3 arrives and haven't sold to 1 or 2}
\end{aligned} \tag{5.6}$$

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

For the *MIP-based policy*, recall that there are m_i discrete signals for each bidder

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

The MIP formulation was encoded using CPLEX and JOpt, a simplified Java Wrapper for mixed integer or linear programming[6]. It was run on each scenario and auction for three hours on a Pentium IV at 3GHZ, allowing 256MB of memory for CPLEX. For comparison, 256MB is not enough memory in the TTT scenario for $m_1 = m_2 = m_3 = 19$. Note that the CPLEX solution was stopped before reaching its tolerance level of 99.9%; any feasible solution to the MIP is however IC, because it must satisfy the ITIC constraints.

Unless stated otherwise, $p_{3A} = 1$, meaning that bidder 3 always arrives. For each auction, a critical signal set matrix was obtained by solving the MIP and then evaluated using sampling of signals: 500000 independent uniform samples of $(s_1, s_2, s_3) \in [0, 1]^3$ were taken for the clairvoyant tests and 10000 for the non-clairvoyant ones, due to these taking significantly longer.

Empirical results

The average revenue for the MIP formulation and the Clairvoyant and Non-Clairvoyant settings are shown in Table 5.1, right. Let us take a closer look at these numbers. In all three auctions, as the bidders intervals' overlap increases, so does the rev-

[6] `http://www.eecs.harvard.edu/econcs/jopt`

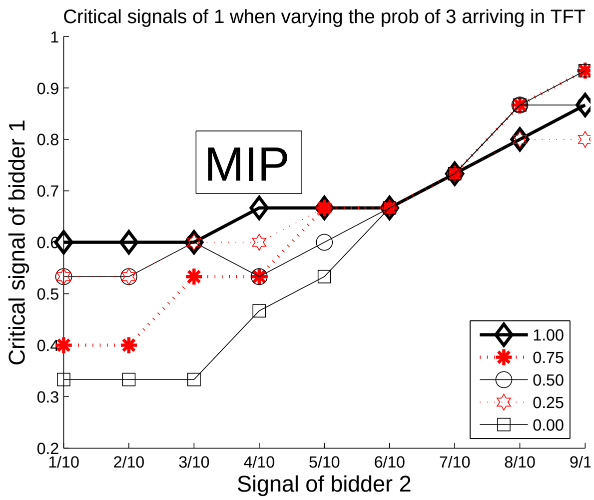

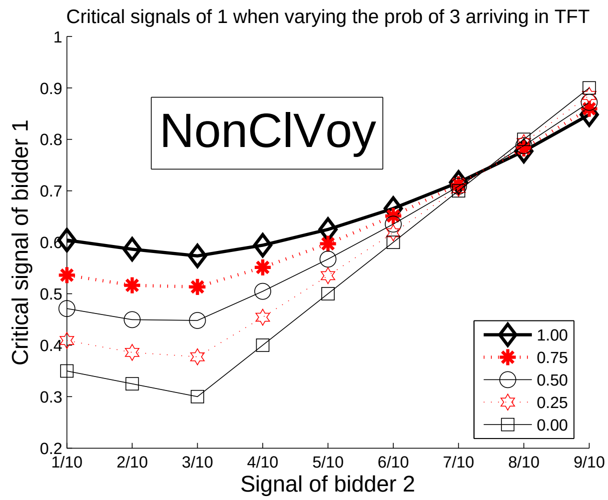

Figure 5.7: Critical signals of bidder 1 as a function of bidder 2's signals in scenario TFT for several values of p_{3A} (shown in legend). The critical signals obtained from the MIP policy are on the left, while the ones obtained from the (Branco) ContNonCV heuristic are on the right.

enue – this is expected since the uncertainty in the model decreases with the amount of overlap.

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

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

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

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

The critical signals for bidders 2 and 3 computed at time d_1 are basically identical for all values of p_{3A}. What differs are the critical signals of bidder 1, plotted in Figure 5.7, left. The signals have an increasing, convex shape and appear to become flatter as p_{3A} increases.[7] The similarity with the signals predicted by the Branco non-clairvoyant (on the right in the same Figure) is striking. Recall though that the Branco non-clairvoyant are *not IC*! It is an interesting question whether this phenomenon is in fact *output ironing* in order to achieve IC (in the sense of Myerson [64]) or whether it is due to the discretizations imposed in the MIP methodology together with suboptimal solutions. I conjecture this effect to be due to leveling of critical signals in regions of the signal space that may violate IC (in this case, pairs of low s_1 and s_2 in Figure 5.7). By way of comparison, I also used the above formulation on the TTT scenario, which is effectively an one-shot, non-dynamic scenario. The (approximate) optimal auction was almost identical to the optimal IC auction predicted by Branco's result, as described in Subsection 5.4.2.

Summary

In this chapter I extended the model of interdependent values to dynamic single-item auctions for bidders with arrivals and departures. I showed that an incentive compatible dynamic auction with interdependent values must assign to each bidder a critical signal that may depend on others' signals and intervals. This critical signal cannot decrease if the bidder reports a later arrival. I showed that no ex post incentive-compatible and reasonable auction exists in the general setting. I developed incentive-compatible, reasonable auctions

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

for settings where bidders cannot misreport both signal and departure.

I considered a simple single-item dynamic setting where bidders have interdependent values, but their arrivals and departures are public. I formulated the design problem of revenue-optimal auctions as a mixed-integer program defined on discretized signals. The formulation has reasonable size if the amount of interdependence between bidders' valuations is small, a coarse signal discretization can be tolerated, and a bound is imposed on the maximal dependence on earlier signals of any bidder's price. The formulation leverages my characterization in terms of critical signals for dynamic incentive-compatible auctions. The design problem thus becomes one of setting appropriate critical signals, which in turn induce an optimal decision policy. In the example I have considered, the optimal policy is close, but not identical, to a heuristic policy generalizing the static revenue-optimal auction.

Chapter 6

Conclusions and Perspectives

In this thesis I designed and analyzed expressive, resilient to manipulations, dynamic auctions with good performance under economic measures such as revenue or social welfare.

I proceed by reviewing the main contributions, discussing the tradeoffs encountered between the different desiderata and finally presenting promising directions for future work.

Summary: Thesis contributions

Incentive compatibility via self-correction

I extended to the domain of non-expiring goods the framework of self-correction, applied to optimization based on a distributional model of future demand.

Self-correction, introduced in [71] for expiring goods, is a local, run-time, rectification technique that establishes the global property of incentive compatibility. Self-correction eliminates opportunities for profitable manipulations in CONSENSUS, an online stochastic combinatorial optimization algorithm [84]. From an allocation perspective, self-correction amounts to discarding the items of bidders who might have benefited from a manipulation.

For the non-expiring goods model that I considered, simply applying self-correction to Consensus leads to frequent discards due to impossibility of verifying that bidders cannot profitably report earlier departures. To overcome this problem, I designed the `NowWait` heuristic adaptation of Consensus, that only allocates to a bidder when his reward is higher than his current estimated opportunity cost. To estimate this opportunity cost, `NowWait` uses the prior distribution on departures (and not any information on reported departures). `NowWait` has encouraging efficiency, despite more discards, and limited computational overhead when compared to a naive simplification of the method in [71]. My framework also allows targeting revenue; I showed how to incorporate the classical method of virtual valuations [64] towards this goal.

Costly cancellations of reservations for impatient buyers

I initiated the worst-case study of mechanisms for impatient buyers bidding on reservations that can be canceled by the seller. I provided an online mechanism $\mathtt{M}_\alpha(\gamma)$ for a model in which each buyer has value for a single item and makes an instantaneous offer for it, incurring a loss of a fraction of this value if his reservation is canceled later.

$\mathtt{M}_\alpha(\gamma)$ is always at most a multiplicative constant factor away from the optimum achievable if all offers arrived simultaneously, in terms of social welfare with or without accounting for the value losses of bidders whose reservations were canceled. A similar bound holds for the revenue of $\mathtt{M}_\alpha(\gamma)$ when compared to the standard VCG mechanism.

If bidders are moderately truthful (the sum of their utilities is non-negative) then $\mathtt{M}_\alpha(\gamma)$ is constant-competitive with respect to bidders' true values as well.

In $\mathtt{M}_\alpha(\gamma)$, no bidder can profit by underbidding, whereas overbidding is profitable only if the bidder would have his reservation canceled when truthful. Constant competitiveness is no longer achievable if one requires that a bidder be best off by bidding his true value under any circumstances.

Competitive and incentive-compatible extensions of $\mathtt{M}_\alpha(\gamma)$ to domains where bidders' information is no longer one-dimensional appear challenging.

Interdependent values

I extended the classical static model of interdependent values from the auction theory literature to single-item dynamic auctions. In my model, each bidder has a private estimate (*signal*) regarding the item's quality, a public mapping from all signals to his value, and an arrival and a departure specifying the interval in which he has value for the item.

I provided a characterization for incentive compatibility, which requires the existence of a critical signal for each bidder: a bidder wins if and only if his reported signal is higher than his critical signal. A bidder's critical signal cannot increase if the bidder reports a later arrival. Using this characterization, I showed that when a bidder can manipulate his departure, no reasonable auction can be incentive-compatible.

I formulated the problem of maximizing revenue in a simple dynamic setting with interdependent values and known arrivals and departures as a mixed-integer program. I showed the similarity of the optimal solution with the solution obtained via a heuristic that generalizes Myerson's [64] virtual valuation auction to dynamic interdependent values.

Tradeoffs between desiderata of dynamic auctions

In the mechanisms designed in this thesis, tradeoffs had to be made between

performance Typical economic performance measures are revenue and social welfare. A usual concern in computer science is the complexity of running a certain auction, for example in computing the optimal allocation or prices.

expressiveness Expressiveness measures the extent to which information provided by a buyer is taken into account by the seller towards making a decision. For instance, allowing each bidder to specify a quantity and perhaps an arrival-departure interval in addition to a bid and taking them into account when deciding an allocation increases an auction's expressiveness.

incentive properties Guarantees of a mechanism's performance are most powerful when they are consistent with predictions on player behavior, for example via equilibrium concepts. A central goal for the auctions in this thesis was achieving strong incentive compatibility, which requires that it is best, under a small set of assumptions, for a bidder to reveal his private information truthfully.

Here are some of the tradeoffs encountered:

Chapter 2 Any truthful function for unrestricted preferences must be an affine maximizer. In contrast, for the very restrictive case of single-minded preferences, there is a rich set, beyond affine maximizers, of social choice functions that are truthful. In dynamic environments, dynamic VCG maximizes expected social welfare and has good incentive properties but is not scalable computationally.

Chapter 3 When all bidders have unit-demand, no cancellations due to self-correction were encountered in any of the experiments. In particular, if all other (active and sampled) bidders have unit-demand, reducing one's quantity cannot result in not being selected anymore by `NowWait`.

Interestingly, `NowWait`'s smooth interfacing with self-correction is guaranteed only when bidders with high values and high patiences tend to be rare.

There exist analytic characterizations of revenue-optimal dynamic auctions only for simple, unit-demand, domains. In contrast, the complex interactions between bidders in my domain seem to require a computational approach.

Chapter 4 If bidders treat items identically and each bidder has value for exactly one item, then no bidder can be strictly better off by bidding below his true value in $\mathtt{M}_\alpha(\gamma)$. However, incentives properties become more involved when bidders have more refined preferences. First, if a bidder has value for exactly one item in a choice subset (and 0 for any other item) then this subset must be known by the seller to ensure that the bidder does not prefer pretending a different choice subset. Second, naively extending the unit-demand algorithm to a setting where items are identical and a bidder may have value for more than one item leads to new opportunities for manipulation.

Call a mechanism with cancellations for allocating a single item *competitive* if, on any sequence of bids, the bid it allocates to is at least a constant factor of the highest bid. It is established in Chapter 4 that no competitive mechanism can incentivize a bidder to always bid his true value. In contrast, if all bids are received at once, the second-price auction always allocates to the highest bid and has truthful bidding as dominant strategy.

Chapter 5 For interdependent values, the most powerful variant of incentive compatibility usually considered is *ex post*: a bidder is best off by truthfully reporting his private information given that others do the same, for any private information the other bidders may have.

There exists an ex post incentive-compatible, reasonable, dynamic auction for bidders with interdependent values if and only if no bidder can misreport his departure. In contrast, in static auctions with interdependent values, a generalization of the VCG mechanism chooses the socially-optimal allocation and is ex post incentive-compatible.

Opportunities for future work

Despite the common thread defining this thesis, in each chapter there are some assumptions that are not touched on in the other chapters. For example, the ability of cancelling previous allocations, correlation between the components of a bidder's type, correlation or interdependence of a bidder's value on other bidders' private information etc. It is of interest to consider the robustness of results in one chapter when assumptions are relaxed or enriched.

I first describe several appealing open questions related to each chapter and then I go on to discuss open-ended questions of more general interest.

Incentive compatibility via self-correction. Before discards due to self-correction, `NowWait` achieves higher social welfare on some distributions when compared to the CONSENSUS algorithm, that only allocates a bidder upon departure. This suggests that better performance may be attained using more sophisticated online stochastic combinatorial optimization methods, such as the EXPECTATION algorithm [84]. Better estimates of opportunity costs for `NowWait` may also lead to better performance. Either extension will, however, require care in coupling with self-correction. It is clearly of interest to study the "first best" solution, i.e. the value- or revenue-maximizing policy among monotonic policies. Obtaining this solution appears a significant technical and computational challenge because monotonicity constraints break the "principle of optimality" that underlies many computational approaches.

Self-correction holds promise in other settings where optimization has been studied mostly without considering incentives. Potential applications include both static traditional computer science domains and dynamic settings such as dynamic auctions with interdependent values (Chapter 5) or dynamic combinatorial auctions. Another research question, suggested by Chapter 4, is the applicability of self-correction to establishing less restrictive notions of incentive compatibility. Such a question may benefit from characterizations of such notions, for example via outcome graph considerations as in Section 4.6.4.

Reservations. Chapter 4 introduces a relaxed form of incentive compatibility: truthful winners are best-responding and underbidding is never better than bidding one's true value. Is (constant) competitiveness possible in other domains under this incentive compatibility notion where it is impossible under dominant strategy incentive compatibility? Such domains should probably involve more complex decisions than a binary decision for players with one-dimensional private information, for which the incentive compatibility notion defined in Chapter 4 appears identical to dominant strategy incentive compatibility. Lavi and Nisan [56] establish such a competitiveness gap in an expiring goods setting for strategies in a Set-Nash equilibrium, a relaxation of dominant strategy equilibrium that may be more or less general than the notion in Chapter 4.

Interdependent values. One appealing variation of the model in Chapter 5 is to allow correlation between bidders' signals, which I assumed to be independent. It is of interest to better understand the relationship between virtual valuation-based methods [15, 64] and the solutions generated through optimization subject to incentive compatibility constraints. A relevant extension towards overcoming the curse of dimensionality are graphical models of interdependent values with rich structure.

In all chapters in this thesis, a bidder's value dropped to 0 at departure. A more practical assumption is that a bidder's value is decreasing over time after departure. A model in which a bidder's value is discounted via a publicly known function is appealing because it does not introduce further opportunities for manipulations.

There is a significant gap between the economics approach of assuming distributions on bidder valuations and the theoretical computer science approach of worst-case analysis. In this thesis, I adopt the distribution-based approach in Chapters 3 and 5 and the distribution-free one in Chapter 4. Obtaining better bounds with significantly less information than a distributional model (such as bounded values [55] or uniform order of arrival, as in the secretary problem and its variants) is an important direction of research aimed at bridging the aforementioned gap.

Perhaps one of the most promising approaches to simultaneously aiming for incentive compatibility and optimization is reducing the former to the latter, as achieved by, for example, Awerbuch et al. [5] or Lavi and Swamy [57]. It seems likely that, regardless of the approach, tradeoffs of incentive compatibility and optimization will need to be made; the findings in this thesis support this statement. A unified metric for losses in performance and profitable manipulations is currently missing and of interest.

"Dynamic" and "online" are sometimes used interchangeably for specifying that an auction takes place over time. More often however[1], the first term is used for persistent populations with changing types whereas the second term is used for changing populations with persistent types. There is a clear need to unify the literatures for the two types of environments, but the unified treatment by Cavallo [18] is the only one that I am aware of.

Beyond dynamic auctions, dynamic mechanism design is very promising for collaborations between computer scientists and economists. While economists are traditionally the experts in mechanism design, the prescriptive approach and the increased relevance of computational considerations in dynamic, as opposed to static, mechanism design render it amenable to computer science expertise.

[1] I feel that "dynamic" is an appropriate name for both settings, whereas "online" auctions are often taken to mean "on the Internet".

Bibliography

[1] M. Aoyagi. Optimal sales schemes against interdependent buyers. *American Economic Journal: Microeconomics*, 2009. forthcoming.

[2] A Archer and E Tardos. Truthful mechanisms for one-parameter agents. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science*, 2001.

[3] Kenneth Arrow. The property-rights doctrine and demand revelation under incomplete information. In M. Boskin, editor, *Economics and Human Welfare*, pages 23–39. Academic Press, New York, 1979.

[4] Susan Athey and Ilya Segal. An efficient dynamic mechanism. `http://www.stanford.edu/~isegal/agv.pdf` (accessed April 8 2009), February 2007.

[5] Baruch Awerbuch, Yossi Azar, and Adam Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proc. ACM Symposium on Theory of Computing (STOC'03)*, pages 503–510, 2003.

[6] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems and online algorithms. In *SODA*, 2007.

[7] Moshe Babaioff, Jason Hartline, and Robert Kleinberg. Selling banner ads: Online algorithms with buyback. In *Fourth Workshop on Ad Auctions*, 2008.

[8] Moshe Babaioff, Jason Hartline, and Robert Kleinberg. Selling ad campaigns: Online algorithms with cancellations. In *ACM Conference on Electronic Commerce*, July 2009. To appear.

[9] Moshe Babaioff, Ron Lavi, and Elan Pavlov. Single-value combinatorial auctions and implementation in undominated strategies. In *SODA*, 2006.

[10] Kaushik Basu and J orgen Weibull. Strategy subsets closed under rationalizable behavior. *Economics Letters*, 36(2):141–146, June 1991.

[11] D. P. Bertsekas. *Dynamic Programming: Deterministic and Stochastic Models*. Prentice-Hall, Englewood Cliffs, NJ, 1987.

[12] S. Bikhchandani, S. de Vries, J. Schummer, and R. Vohra. An ascending Vickrey auction for selling bases of a matroid. Accessed April 2009, August 2008.

[13] Eyal Biyalogorsky and Eitan Gerstner. Contingent pricing to reduce price risks. *Marketing Science*, 23(1):146–155, 2004.

[14] Craig Boutilier, David C. Parkes, Tuomas Sandholm, and William E. Walsh. Expressive banner ad auctions and model-based online optimization for clearing. In *Proc. 23rd AAAI Conference on Artificial Intelligence*, pages 30–37, 2008.

[15] F. Branco. Multiple unit auctions of an indivisible good. *Economic Theory*, 8:77–101, 1996.

[16] Matthew C. Cary, Abraham D. Flaxman, Jason D. Hartline, and Anna R. Karlin. Auctions for structured procurement. In *SODA*, 2008.

[17] Ruggiero Cavallo. *Social Welfare Maximization in Dynamic Strategic Decision Problems*. PhD thesis, Computer Science, Harvard University, 2008.

[18] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Efficient online mechanisms for persistent, periodically inaccessible self interested agents. Tech. report, Harvard Univ., 2007.

[19] Xi Chen and Xiaotie Deng. Settling the complexity of 2-player nash-equilibrium. In *Electronic Colloquium on Computational Complexity*, 2005.

[20] Yiling Chen, Lance Fortnow, Nicolas Lambert, David Pennock, and Jennifer Wortman. Complexity of combinatorial market makers. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 190–199, New York, NY, USA, 2008. ACM.

[21] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[22] Richard Cole, Shahar Dobzinski, and Lisa Fleischer. Prompt mechanisms for online auctions. In *First Symposium on Algorithmic Game Theory*, 2008.

[23] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Uncertainty in Artificial Intelligence*, pages 103–111, 2002.

[24] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proc. 18th Conference on Uncertainty in Artificial Intelligence*, pages 103–110, 2002.

[25] Vincent Conitzer and Tuomas Sandholm. Incremental mechanism design. In *IJCAI*, 2007.

[26] Peter Cramton, Yoav Shoham, and Richard Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.

[27] P. Dasgupta and E. Maskin. Efficient auctions. *The Quarterly Journal of Economics*, CXV, 2000.

[28] Rajdeep Dash, Alex Rogers, and Nicholas Jennings. A mechanism for multiple goods and interdependent valuations. In *Agent-Mediated Electronic Commerce VI*, 2004.

[29] C. Daskalakis, P.W. Goldberg, and C.H. Papadimitriou. The complexity of computing a nash equilibrium. In *Electronic Colloquium on Computational Complexity*, 2005.

[30] C d'Aspremont and L Gerard-Varet. Incentives and incomplete information. *J. of Public Economics*, 11:25–45, 1979.

[31] Nedialko B. Dimitrov and C. Greg Plaxton. Competitive weighted matching in transversal matroids. UT Austin, Technical Report T-08-03, January 2008.

[32] U. Feige, N. Immorlica, V. Mirrokni, and H. Nazerzadeh. A combinatorial allocation mechanism for banner advertisement with penalties. In *World Wide Web Conference (WWW)*, 2008.

[33] J. Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006.

[34] J. Gallien and S. Gupta. Temporary and permanent buyout prices in online auctions. *Mgmt. Science*, 53:814–833, 2007.

[35] Alex Gershkov and Benny Moldovanu. Learning about the future and dynamic efficiency. Forthcoming in American Economic Review.

[36] Alex Gershkov and Benny Moldovanu. Dynamic revenue maximization with heterogeneous objects: a mechanism design approach. `www.econ2.uni-bonn.de/gershkov/pdf/dynamicrev.pdf`, February 2008.

[37] Alan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–602, 1973.

[38] John Gilbert and Frederick Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61(313):35–73, March 1966.

[39] Andrew V. Goldberg and Jason D. Hartline. Envy-free auctions for digital goods. In *In Proc. of 4th ACM Conference on Electronic Commerce*, pages 29–35. ACM Press, 2003.

[40] Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

[41] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[42] Hongwei Gui, Rudolf Muller, and Rakesh Vohra. Dominant strategy mechanisms with multidimensional types. 2004.

[43] Mohammad Hajiaghayi, Robert Kleinberg, Mohammad Mahdian, and David Parkes. Online auctions with re-usable goods. In *ACM Conf. on Electronic Commerce*, 2005.

[44] Mohammad Hajiaghayi, Robert Kleinberg, and David C. Parkes. Adaptive limited-supply online auctions. In *ACM Conf. on Electronic Commerce*, pages 71–80, 2004.

[45] Mohammad Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI 2007*, pages 58–65, 2007.

[46] Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1:2007, 2002.

[47] Leonid Hurwicz. On the existence of allocation systems whose manipulative Nash equilibria are Pareto optimal. Unpublished paper presented at the 3rd World Congress of the Economic Society, Toronto, 1975.

[48] T. Ito and D. Parkes. Instantiating the contingent bids model of truthful interdependent value auctions. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2007.

[49] P. Jehiel and B. Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69:1237–1259, 2001.

[50] Michael J. Kearns, Yishay Mansour, and Andrew Y. Ng. A sparse sampling algorithm for near-optimal planning in large Markov Decision Processes. In *IJCAI*, 1999.

[51] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems.* Springer, Berlin, Germany, 2004.

[52] Vijay Krishna. *Auction Theory.* Academic Press, 2002.

[53] Vijay Krishna and Motty Perry. Efficient mechanism design. Technical report, Pennsylvania State University, 2000. Available at: http://econ.la.psu.edu/˜vkrishna/vcg18.ps.

[54] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *Proceedings of FOCS '03*, 2003.

[55] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conf. on Electronic Commerce*, pages 233–241, 2000.

[56] Ron Lavi and Noam Nisan. Online ascending auctions for gradually expiring items. In *Symposium on Discrete Algorithms (SODA)*, 2005.

[57] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In *FOCS*, 2005.

[58] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with

decreasing marginal utilities. In *EC '01: Proceedings of the 3rd ACM conference on Electronic Commerce*, 2001.

[59] Daniel Lehmann, Rudolf M uller, and Tuomas Sandholm. The winner determination problem. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 12. MIT Press, 2006.

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

[61] Pinyan Lu, Shang-Hua Teng, and Changyuan Yu. Truthful auctions with optimal profit. In *WINE*, pages 27–36, 2006.

[62] A. McGregor. Finding graph matchings in data streams. In *Proceedings of APPROX-RANDOM*, pages 170–181, 2005.

[63] Paul Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.

[64] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, VI:58–73, 1981.

[65] Roger B Myerson. Multistage games with communication. *Econometrica*, 54(2):323–358, 1986.

[66] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

[67] John Nash. Equilibrium points in N-person games. In *Proceedings of the National Academy of Sciences*, volume 36, pages 48–49, 1950.

[68] Mallesh Pai and Rakesh Vohra. Optimal dynamic auctions. `http://www.kellogg.northwestern.edu/faculty/vohra/ftp/vohra91.pdf`. Accessed April 2009., March 2008.

[69] David C. Parkes. Online mechanisms. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 16. Cambridge University Press, 2007.

[70] David C. Parkes. When analysis fails: Heuristic mechanism design via self-correcting procedures. In *Proc. 35th Conf. on Current Trends in Theory and Practice of Computer Science (SOFSEM'09)*, 2009.

[71] David C. Parkes and Quang Duong. An ironing-based approach to adaptive online mechanism design in single-valued domains. In *Proc. 22nd National Conference on Artificial Intelligence (AAAI)*, 2007.

[72] David C. Parkes and Satinder Singh. An MDP-Based approach to Online Mechanism Design. In *Proc. NIPS*, 2003.

[73] David C. Parkes, Satinder Singh, and Dimah Yanovsky. Approximately efficient online mechanism design. In *Proc. 18th Annual Conf. on Neural Information Processing Systems (NIPS'04)*, 2004.

[74] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: Revenue equivalence, profit maximization and information disclosure. `www.stanford.edu/~isegal/dmd.pdf`, June 2008.

[75] David Pennock and Rahul Sami. Computational aspects of prediction markets. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 26. Cambridge University Press, 2007.

[76] Warren B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley-Interscience, 2007. Wiley Series in Probability and Statistics.

[77] Kevin Roberts. The characterization of implementable choice rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences*, pages 321–348. North-Holland, Amsterdam, 1979.

[78] Jean-Charles Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16:191–200, 1987. North-Holland.

[79] Michael Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In *ACM EC*, 2005.

[80] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. In *Artificial Intelligence*, pages 542–547, 2001.

[81] Mark A Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[82] T. C. Schelling. *The Strategy of Conflict*. Harvard University Press, 1960.

[83] Kalyan T. Talluri and Garrett John Van Ryzin. *The theory and practice of revenue management*. Boston, Mass. [u.a.], Kluwer Academic Publ, 2004.

[84] Pascal van Hentenryck and Russell Bent. *Online Stochastic Combinatorial Optimization*. MIT Press, 2006.

[85] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[86] Robert Wilson. Computing equilibria of N-person games. *SIAM Journal on Applied Mathematics*, 21(1):80–87, July 1971.