POLITECNICO DI MILANO
SCHOOL OF INDUSTRIAL AND INFORMATION ENGINEERING
Master Degree in Automation and Control Engineering

Posted Pricing with Time-Discounted Valuations and Poisson Arrivals

Supervisor: Prof. Nicola Gatti
Co-Supervisors: Giulia Romano, Alberto Marchesi

Master Thesis of:
Gianluca Tartaglia
ID 899867

Academic Year 2018-2019

Sommario

In questa tesi definiamo dei meccanismi economici detti di *posted-price* al fine di vendere un unico bene entro un certo periodo di tempo. Un esempio pratico è quello degli affitti a lungo termine di camere e/o appartamenti. Il nostro obiettivo è quello di progettare meccanismi che siano in grado di massimizzare il guadagno. Si assume che gli agenti arrivino nel mercato in modo sequenziale, secondo un *processo di Poisson* e venga offerto loro un *take-it-or-leave-it price*. Le loro valutazioni sono ignote al venditore e *tempo varianti*. In particolare sono assunte decrescere nel tempo secondo una funzione di sconto. Definiamo e analizziamo due scenari che differiscono nelle assunzioni sulle valutazioni. Questi scenari sono studiati sia con un approccio teorico, che tramite *programmazione matematica*.
Inizialmente, studiamo lo scenario *Identical Valuation*, dove tutti gli agenti hanno la stessa valutazione del bene, ignota al venditore. In questo scenario, proponiamo il meccanismo *ottimo*, secondo una *worst-case competitive analysis*, nel caso di generica funzione di sconto. Nel caso di sconto lineare del tempo, il meccanismo ottimo, che chiamiamo $\mathcal{M}_1$, ha una strategia di prezzo continua nel tempo. Nello stesso scenario, viene presentato un approccio tramite programmazione matematica che approssima arbitrariamente bene il meccanismo ottimo.
Successivamente, studiamo lo scenario *Random Valuation*, dove le valutazioni degli agenti sono estratte da un'ignota distribuzione di probabilità F che presenta *monotone non-decreasing hazard rate*. Proponiamo un nuovo meccanismo $\mathcal{M}_2$ con una strategia di prezzo costante a tratti, nel caso di sconto lineare nel tempo delle valutazioni. Studiamo le prestazioni di $\mathcal{M}_1$ e $\mathcal{M}_2$, calcolando due *lower bounds* per il fattore competitivo. Questi valori risultano costanti rispetto alla scelta di F. Lo scenario *Random Valuation* viene studiato anche da una punto di vista di programmazione matematica.

Abstract

In this thesis, we design economic mechanisms said of *posted-price* in order to sell a unique good in a finite period of time. A practical example is the long-term rental of rooms and/or apartments. The aim of the mechanisms is to maximize the seller's revenue. We assume agents sequentially arrive to the market according to a *Poisson process* and are offered a *take-it-or-leave-it* price. Their valuations are *unknown* to the seller and *time-variant.* In particular, they are assumed to be decreasing over time according to a discount function. We define and analyze two different scenarios which differ in the assumptions made on the valuations. The settings are studied both with theoretical and mathematical optimization approach.

We first study the *Identical Valuation* setting, where all agents have the same valuation for the item, that is unknown to the seller. We provide the *optimal* pricing mechanism according to a worst-case *competitive analysis*, in the case of a generic discount function. We specify the optimal mechanism $\mathcal{M}_1$ when the discount linearly depends on time. $\mathcal{M}_1$ has a pricing strategy that is continuous over time. In the same scenario, a mathematical programming approach is presented. This method turns out to approximate arbitrarily well the optimal mechanism.

Then, we study the *Random Valuation* setting, where agents' valuations are drawn from an unknown distribution F with monotone non-decreasing *hazard rate*. We provide a new mechanism $\mathcal{M}_2$ that propose a staircase pricing strategy, in the case of linear discount. We study the performances of $\mathcal{M}_1$ and $\mathcal{M}_2$, deriving two lower bounds for the competitive ratios. These values turn out to be constant with respect to the choice of F. The Random Valuation setting is also studied from an optimization perspective.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Research Area and Main Problem

In this thesis we design economic mechanisms said of *posted-price* for selling a unique good within a finite period of time. We study such mechanisms in an *online* setting, in the sense that agents' arrival time has a dynamic behavior that is modeled by a Poisson process and the valuations are unknown to the seller. The seller interacts with agents by sequentially offering a take-it-or-leave-it price. The mechanism terminates when an agent accepts the offer and buys the item or the deadline expires. If an agent rejects the offer, she leaves the market and never returns. In particular, agent i buys the item if the resulting utility $v_i - p_i$ is positive, where v_i is her personal valuation of the good and p_i is the price posted by the seller. If the utility is negative, the agent rejects the offer.

The aim of the mechanism is to maximize the seller's revenue that, in this setting, is equivalent to maximize the social welfare. Asking agents to report their true values as prescribed by classical mechanism design is unrealistic: they have no incentives to reveal more information than necessary. Therefore, we use the bids as a proxy for their unknown valuations in order to find the welfare-maximizing outcome. In this scenario the bid is the price paid by the agent accepting the offer. Hence, the revenue-maximizing mechanism also maximizes the social welfare. Posted-price auctions are also dominant-strategy incentive-compatible (*DSIC*), individually rational and weakly budget balanced. In this setting the DSIC property means that acting truthfully is a dominant strategy: buying or leaving reflects agents' true preferences. Even if valuations are not reported, we know that there is no strategic behavior.

Another advantage of the posted-price mechanisms is that they are robust

with respect to collusion, (*a.k.a. group strategyproof*). An agent could help another one only refusing an offer that could be convenient for her. But this should be in contrast with her own utility.
Our mechanism is *model-free* because it does not rely on the assumption of knowing the distribution of valuations, which would be highly improbable for our setting. Indeed, in a single-item single-unit scenario there is no possibility of learning the demand curve, therefore, the seller cannot have accurate beliefs about the distribution of valuations. Even in a repeated scenario it could be difficult collecting samples. This explains why a Bayesian approach cannot be used.
A further advantage of posted pricing is that agents do not need to wait the outcome of the mechanism. Similarly to non-parametric learning, in posting pricing the main *computational effort* is offline. Consequently, the online computational effort is practically negligible and the agents immediately know whether they buy the item or not.
We propose for the first time in literature a *time-variant* model. We consider a time-dependent behaviour of the agents' valuations. Specifically, a time decreasing *discount* is applied to the initial valuations of the agents. This meets two possible interpretations and corresponding microeconimic settings. In the first one the agents' willingness to pay an item decreases over time. In the second one the agents' valuations are constant; but it is crucial for a seller the time instant at which the item is sold. Specifically, the later the item is sold, the smaller the revenue of the seller. It is to the latter interpretation that we refer the *motivating example*. The *time-dependency* challenge also highlights the role of the *deadline*. The presence of a deadline clearly modifies the mechanism. The pricing strategy should be tuned with respect to how much time there is before reaching the deadline.
Following the literature, the present work deals with two models. The first one is the *Identical Valuation* scenario, in which all agents have the same unknown discounted valuation. The second one is the *Random Valuation* scenario, in which the undiscounted valuations are drawn from an unknown probability distribution F. We suppose the widely recognised assumption that F is a *Monotone Hazard Rate* - MHR - distribution.

1.1.1 Motivating Example

The long-term rental market in specific metropolitan areas is widespread and fast growing. Handling the market changes and the dynamic of the demand in an online manner, and simultaneously guaranteeing the revenue maximization could be hard in several real microeconomic scenarios. An economic mechanism able to automatically solve this challenge could be a good choice in terms of *efficiency* and *revenue guarantees*. A first possible instance of this problem is that of a seller which aims to rent a room or an apartment within a deadline. The potential buyers have their private valuations over that room. The number of these agents are not *a-priori* known to the seller who cannot even know the arrival time instants of these agents. The private valuation of an agent could be assumed constant with respect to a fixed amount of time, a *month* for example, or it could abruptly change depending on the *seasons*. If for a buyer the number of rental months is a mere question of *need*, for a seller it is a crucial point. A seller would largely prefer to rent a room at the minimum price for twelve months rather than for a month at the maximum price. It is clear that the *time* assumes an essential role. A possible solution could be considering the cumulative buyers' valuations decreasing with respect to time. This thesis tries to solve this problem when the seller have no information about the willingness-to-pay of the customers.

1.2 Original Contributions

In the Identical Valuation framework, several alternative settings are considered. We propose the *optimal mechanism* with general shapes of continuous decreasing private valuations. This mechanism represents the *Upper Bound* on the basis of *Competitive Analysis*. This means that no mechanism reaches a *competitive ratio* higher than the ours. In the case of linear discounting valuations, we call the optimal mechanism $\mathcal{M}_1$. We also consider seasonability of the market. We suppose the agents' valuations could change abruptly. This means that they have not a smooth behaviour in time. We propose an analysis of this scenario and a mechanism able to achieve a constant competitive ratio.

In the Random Valuation scenario we take into account the case of linearly dependent valuations. In this framework we evaluate mechanism $\mathcal{M}_1$ which exhibits a constant lower bound of the competitive ratio, independently from the unknown MHR F. We also present a new mechanism $\mathcal{M}_2$ able to reach a higher constant lower bound for the competitive ratio.

We also present an *automated mechanism design* method to both the settings. We provide a *flexible* tool which can be adapted to different variants of the model and/or additional constraints. In Identical Valuation, following a mathematical programming approach, we introduce an algorithm able to optimize the competitive ratio and reaching an estimate - both a lower and an upper bound - very close to the one of the optimal mechanism in a reasonable amount of time. We call this algorithm the *Maximum Valuation Algorithm*, or *MVA*. We also propose a mechanism handling a variant of the model, in which seller cannot continuously change the price over time. We adapt the *MVA* also in the Random Valuation scenario, proposing a new possible interpretation of the model in this framework.

1.3 Thesis Structure

The Thesis is structured in the following way:

- Chapter 2 lays the theoretical groundings on which this work is based. It presents some key concepts from Mechanism Design and Automated Mechanism Design, presenting some methods to handle pricing problems from the literature.

- In Chapter 3 the *identical Valuation* scenario is studied. We provide the optimal mechanism in three different scenario. Then we study the case in which the agents' valuation abruptly change.

- In Chapter 4 the same problem is addressed from a mathematical programming point of view. It is presented an algorithm able to optimize the competitive ratio, the *MVA*. Then some experimental results are shown. We consider also possible variations of our model, including one in which there is one more constrained, *i.e.* the price cannot be modified continuously over time.

- Chapter 5 deals with the *Random Valuation* scenario in the case of linear decreasing valuations. We discuss the role of the benchmark mechanism, propose a new mechanism $\mathcal{M}_2$ and evaluate both $\mathcal{M}_1$ and $\mathcal{M}_2$.

- In Chapter 6 the *Random Valuation* scenario is treated applying an optimization approach. We propose an interpretation of the model with a *Markov chain* and a new version of the *MVA*.

- In Chapter 7 we draw conclusions. We summarize the main results of our work and we propose some future developments.

Chapter 2

Preliminaries

Our work refers to a general and common *microeconomic* scenario, well-known in some engineering fields, such as, for example, *Economic Mechanism Design*, *Algorithmic Game Theory* and *Multi-Agent Systems*. In this scenario, an agent - the *seller* - aims to sell an item, when some other agents - the *buyers* - are interested in purchasing that item. Every agent is *self-interested*. In general, it does not necessarily mean that they want to cause harm to each other, or even that they care only about themselves. Instead, it means that each agent has his own description of which states of the world he likes—which can include good things happening to other agents—and that he acts in an attempt to bring about these states of the world. In order to represent the grade of happiness of an agent, the *utility theory* is the common practice. A utility function is a mapping from states of the world to real numbers: each agent has a *private valuation*, or *type* over the item, that represents her willingness-to-pay. The seller is a profit-maximizer.
In literature, many instances of this basis scenario have been studied. We underline the main features considered. We may have one seller (*monopoly*), multiple sellers (*oligopoly*) or many sellers (perfect *competitive market*). One could be interested in selling one item, multiple items, and/or multi units of a single item. The private valuations over these items should be considered in a *Bayesian scenario*, where the buyers' valuations are samples from some probability distributions, which may be *known*, or *unknown*.
The theoretical approaches studied to manage these kinds of scenarios are many. Obviously, they have similar characteristics, but at the same time, they can be very different. We underline four main approaches' techniques: *Mechanism design*, *Online* mechanisms, *Learning* approach and *Dynamic Programming*. It is worth mentioning the fact that strategies from different approaches could share the same characteristics.

2.1 Economic Mechanism Design

Mechanism Design is a subfield of economic theory that is rather unique within economics in having an engineering perspective, (Nisan et al., 2007). We consider a scenario with a finite set of agents which are asked to declare their type; but they may do *truthfully* or not. A simple example is that of *sealed-bid auction* for a single item. Mechanism design is the art of designing the rules of the game so that the agents are motivated to report their type preferences truthfully, and a desirable outcome is chosen, (Conitzer and Sandholm, 2004). This is the kind of question that a mechanism designer tries to answer. Mechanism design has to decide two things: *who wins what* and *who pays what*, that is basically the ultimate task of any pricing strategies.
Consider the following scenario:

- $N = \{1, \dots, n\}$ set of agents;
- O is a set of outcomes;
- $\Theta = \Theta_1 \times \cdots \times \Theta_n$ is a set of possible joint type vectors;
- f is a probability distribution over Θ;
- $u = u_1, \dots, u_n$, where $u_i : O \times \Theta \longmapsto \mathbb{R}$ is the utility function for each player i;
- $A = A_1 \times \cdots \times A_n$, where A_i is the set of actions available to player i;
- g: $A_1 \times \cdots \times A_n \longmapsto O$ is the outcome function.

The first fundamental concept is that one of *social choice function.* It is a function that maps the preferences of the players into an outcome.

Definition 1. ***(Social Chioce Function)*** *A social choice function over N and O is a function $C : \Theta \longmapsto O$.*

Basically, given the preferences of the agents, a social choice function investigates ways in which they can be aggregated, respecting some properties, such as efficiency or individually rationality. When studying the properties of social choice functions, one assume to have perfect information of the types of the agents. We start talking about strategic agents when we consider a mechanism.

Definition 2. ***(Economic Mechanism)*** *An economic mechanism is a tuple (A, O, g).*

Definition 3. ***(Bayesian Game)*** *A bayesian game is a tuple (A, O, g, Θ, f, u).*

A social choice function chooses the outcome only considering the types of the players, regardless the fact that a player should misreport her type. We say that a mechanism *implements* a social choice function only if *there exists* an equilibrium *such that* the outcome chosen by g is the same of the social choice function. The aim of mechanism design is to select a mechanism, given a particular Bayesian game setting, whose equilibria have desirable properties, (Shoham and Leyton-Brown, 2008).

Definition 4. ***(Direct economic mechanism)*** *Given a social choice function C, a direct (revelation) economic mechanism is a mechanism (Θ, O, C).*

Thanks to the concept of direct mechanisms, we can now define the following;

Definition 5. ***(Incentive compatibility)*** *A social choice function C is incentive compatible (or truthfully implementable) if the Bayesian game induced by the direct revelation economic mechanism (Θ, O, C) has a pure equilibrium (according to some solution concept, e.g. dominant-strategy) $s_1^*, \ldots, s_n^*$ such that $s_i^*(\theta_i) = \theta_i$ for every player i and type θ_i.*

For a mechanism to be *DSIC* - dominant strategy incentive compatible - means that for the players truthfully responding is a dominant strategy, no matter what other players do. Moreover, every truthfully players is guaranteed non-negative utility. A theorem by Gibbard and Satterthwaite, known as *impossibility theorem* - (Gibbard), (Satterthwaite, 1975) - forces the research to make more assumptions on the utility function (actually the theorems maintains that other assumptions are possible, but they should be too restrictive). We consider the case of quasi-linear environment, and specifically single-parameter linear environment; the latter is a special subclass of the former and corresponds to the single-item single-unit case. With this environment, Myerson achieves a fundamental result, (Myerson, 1981), providing the optimal auction. The Myerson's mechanism is the milestone of economic mechanism design, the foundation on which the most of the next mechanisms are built.

There are successful applications of economic mechanism design, for example the *Sponsored Search Auctions (SSAs)* - see (Farina and Gatti, 2017) or (Gatti et al., 2015). Similar approaches are based on *Game Theory*, for example we mention the *bargaining problem*, - see (Giunta and Gatti, 2006), or (An et al., 2016).

Myerson's mechanism and, more generally, Vickrey auctions[1] satisfy the sur-

[1] In a specific scenario the Myerson optimal auction is a Vickrey auction with an optimally chosen reserve price.

plus maximization property if the players report truthfully. Even though the bidder valuations were *a priori* unknown to the auctioneer, the auction nevertheless successfully identifies the bidder with the highest valuation. Unfortunately, the assumptions made are too strong to be implemented in several real economic scenarios[2]. This is the great weakness of Myerson's mechanism. See for example (Ausubel et al., 2006) with its very expressive and meaningful title *'The lovely but lonely Vickrey auction'*.
Vickrey auctions have several weaknesses. Varying the scenarios of interest, these weaknesses can be more or less relevant. For instance, it should be considered a multi-unit selling, or the presence of competitors in the market, or some of the assumptions may be unrealistic, *e.g* the knowledge on the probability distribution from which the buyers sample their private valuations. Moreover, they are vulnerable to collusion by a coalition of losing players and to the use of multiple bidding identities by a single bidder. It is quite interesting another objection made by (Kleinberg et al., 2016). Theory neglects the uncertain investment required to investigate purchases. Indeed, information acquisition costs play a very important role in many economic mechanisms. They propose an interesting connection with a real option.
In any case, the two main issues are the following: 1) very often it is not possible to consider that all the agents are in the market in a specific time instants, and that they wait until the end of the auction. Moreover, in some settings, taking an auction is considered unfair. This is the case of the rent, for example. 2) asking agents to report their valuations might be unfeasible. Very often, buyers are unwilling to report their valuations. This encourages alternative mechanisms.
To solve the last two issues a new kind of *online* mechanism has been studied: the *Posted-Pricing Mechanism*, or *Posted Price Auction*.

2.2 Online Mechanisms

Online mechanisms handle a dynamic environment, in the sense of online algorithm. Agents may arrive or departure, there is uncertainty about the types of the players and, more in general, there is uncertainty about the future. Information is revealed online and decisions must be made dynamically. There are two main frameworks in which to study the performance of online mechanisms. The first is *model-free* and is useful when a designer does not have good probabilistic information about future agent types. The

[2]In some cases they can be used and are currently implemented, for example in contextual advertising, by Google and Facebook.

second is *model-based* performing in a data-rich environment. It is reasonable to believe that the seller can build an accurate model to predict the distribution on types of buyers.
One of the most important type of mechanisms is the *Posted-Pricing*. In this mechanism the seller offers a *take-it-or-leave-it* price to each potential buyer facing the mechanism. The buyer either accept or rejects the offer. If she accepts the offer then she wins the item; otherwise, if she rejects, she leaves the market and the seller waits for the next player. In posted-price mechanisms with single item, bidders have a dominant strategy to accept any offer which is below their values, and reject it otherwise. Note that the bidders are not expected to reveal their exact valuation, but only to send a reject or accept message, (Babaioff et al., 2011).
From this point on, we will focus only on the posted-pricing, underlying some of the most interesting works in literature. Basically there are two possible ways to handle the dynamic of the demand. The first one, the simplest way, is consider a certain number of agents that sooner or later appear in the market. The dynamic behaviour of the demand is consider simply trough a sequence of agents. The second one - ours choice - considers the arrival times of the agents as a stochastic process, a common practice in literature is the Poisson process.

A first typology of posted-pricing is the *Sequential Posted pricing Mechanism*, or *SPM*. SPMs has been studied by (Sandholm and Gilpin, 2003) and (Blumrosen and Holenstein, 2008), but we refer our attention on *'Multi-parameter Mechanism Design and Sequential Posted Pricing'*, (Chawla et al., 2010). In SPM only the question *when the agents arrive* is uncertain. The questions *how many agents* or *whether they arrive or not* have certain answers. Note that in our work also the last two are uncertain. Moreover, the authors consider a data-rich environment, hence a *model-based* framework. SPM offers in sequence take-it-or-leave-it prices to the agents, defining both the the prices and the sequence. They reach a very good approximation with respect to the optimal auction. They also consider the so-called *Oreder-oblivious Posted pricing Mechanisms*, or *OPMs*, in which it is not possible to define the sequence of the agents facing the mechanism. However, the main results are reached by proposing an intuitive and quite simple way to switch from single-parameter to multi-parameter scenario, using the lens of approximation. The basic idea is that each multi-dimensional agent is represented by many independent single-dimensional agents. An interesting discussion regards the incentives of the agents. If in single-parameter truthfully responding is a dominant strategy, in multi-parameter this does not hold anymore. They propose a method to preserve this guarantee.

An exciting connection exists between the theory of SPM and the so-called *Economic Prophet Inequalities*, from the optimal stopping theory. (Hajiaghayi et al., 2007) notice a natural analogy between the single-item posted price problem and this theory proven in the $70s$. Actually the analogy is much larger, including also the multi-parameter scenario. We refer to (Lucier, 2017) for a good summary. Prophet inequalities theory basically considers the following setting, let us call it the *Treasure Game*: 'there are n locked treasure chests, each contain a treasure prize. Every chest has a non-negative distribution over the possible prize. The values of the chest are drawn independently. It's possible to open a chest at a time. If you accept the prize, the game ends; if you refuse the prize, it is lost to you forever.' The pricing analogy is evident.

A similar setting is studied by Kleinberg and Leighton in their *'The Value of Knowing the Demand Curve'*, (Kleinberg and Leighton, 2003). They consider n buyers potentially interested in a single-item with unlimited supply in a *model-free* environment. Notice a very different assumption of *unlimited* available units. They examine an additive regret $\mathbb{E}[\mathcal{S}] - \mathbb{E}[\mathcal{S}^{bench}]$, where $\mathcal{S}$ is the strategy of the studied mechanism, and $\mathcal{S}^{bench}$ is the strategy of the benchmark. The interpratation proposed is very interesting: in the case where buyers' valuations are *i.i.d.* from a fixed unknown probability distribution - usually specified by a demand curve - one can interpret the additive regret as *how much the seller should be willing to pay for knowledge of the demand curve*.

We want now to report other two fascinating works, which face issues not yet solved in posting pricing, a possible strategic behaviour of agents and the presence of competitors in the market. Let us consider the first issue, reporting a work by Mohri and Munoz, (Mohri and Munoz, 2014). They proposed a quite different single-unit setting, in which the seller and a buyer face each other in a real *trade*. They suggest that this scenario has three possible interpratations: 1. a *Second-Price Auction* with reserve with one bidder - even if there can be many potential buyers there's always a $1vs1$ situation; 2. *Two-player repeated non-zero sum game* with incomplete information and 3. a *Multi Armed-Bandit* problem, since only the reward for the price chosen is accessible by the seller. The assumption that the valuations of the agents are samples from an unknown distribution is broken. An agent should act strategically against the seller, because they are in a two-player repeated game. The proposed idea is to consider a tree, where every node is labelled with a price that is offered to the buyer. When a price is rejected, the seller offers the same price for a specific number of rounds. The second work is a recent study by Rong, Qin and An, (Rong et al., 2018). This study focus on a competitive market. They consider a dynamic demand which implies dynamic inventories for reusable resources. For each provider, they consider a *Birth Death* process - a special Markov process - modelling the dynamic supply. Since each provider aims to maximize her expected revenue, the optimal policy is supposed to be a *pure strategy Nash equilibrium.* Clearly, the providers have not full information to compute such an equilibrium, the idea proposed is an approximate Nash equilibrium that reaches very good results. The goodness of this model is inside its intuitive way of representation of the demand and the market competition, from the price sensitivity, to the providers' attractiveness.

Finally, we briefly discuss the two closest works with respect to our setting, (Babaioff et al., 2011) and (Zheng et al., 2016). Both of them consider a *model-free* single-unit scenario. They suppose the private valuations of the agents to be samples *i.i.d.* from an unknown probability distribution, from the family of *Monotone Hazard Rate*. Some of the results reached in these works will be generalised in the present thesis.

2.3 Learning approach

This section is dedicated to briefly report some methods that use *Machine Learning* techniques to solve pricing problems. It is clear that many learning techniques are actually *online* methods. Nevertheless, we want to distinguish the following approaches to that ones of the previous section because the former are more *data-oriented*.

First of all, note that a learning approach is possible only in a multi-unit scenario. In posted pricing, when a player is interested in selling a single unit of an item - or even few units - the learning ability is limited, because the mechanism ends with the first *accept* message. Moreover, different from the common mechanism design scenario, the players do not reveal their exact valuations. Hence, we can only learn from a *reject* message and consequently infer that the private valuation might be below that price. In other settings - like the ours - these *reject* messages are not available because when an agent faces a too high price, she simply leaves the market. Hence, the only way to learn is to turn upside down a well-known proverb into a new formulation:

silence gives dissent.
A further consideration about mechanism design and learning. In an auction mechanism, if an agent knows that a revenue optimization algorithm is used, she can decide to misreport her valuations. Indeed, consistent empirical evidence of strategic behaviour by advertisers has been found by (Edelman and Ostrovsky, 2007).

We now present a quite different posted pricing scenario with respect to the ours, in which an unlimited inventory is available. There is no deadline, and the assumption that there is a *best* price that can be chosen. Note that such a best price may change during time. A currently used technique in many real economic settings is the *A/B/n testing* which is an *offline* learning technique. A huge amount of data is collected, then analysed and finally the best candidate is chosen. A/B/n testing has several weaknesses. A much more efficient method is *Multi Armed Bandit.*
One of the main works dealing with bandits from pricing actually has been already mentioned, (Kleinberg and Leighton, 2003). We want now to refer to two more recent works, (Trovo et al., 2015) and (Trovo et al.).
In a bandit problem, there is a finite number of prices candidates - the *arms* - and at each time instant an algorithm should chose an arm. The goal is the minimization of the regret, *i.e.* the loss the seller incurs in the choice of a suboptimal arm. The *Clairovoyant* algorithm is defined as the ideal seller which *a-priori* knows the expected value of each candidate and every time selects the best one. Consider the following notation:

- A is a set of arms;
- a_t is the arm played at time t;
- a^* is the optimal arm;
- μ_a is the expected reward of arm a;

Given an algorithm $\mathcal{U}$, the cumulative regret after T rounds $\mathcal{R}_T(\mathcal{U})$ is defined as follows:

$$\mathcal{R}_T(\mathcal{U}) = T\mu_a - \sum_{t=1}^{T} \mu_{a_t} \tag{2.1}$$

The assumptions made is the monotonicity of the *demand curve* and the presence of a-priori information about the order of magnitude of the *conversion rate*, *i.e.* the probability that a buyer purchases the good. *MAB* techniques like *UCB* or *Thompson sampling* can be used, ensuring the minimization of the regret.

Finally, we briefly present some machine learning techniques in *second-price auctions*, (Medina and Mohri, 2014) and (Cesa-Bianchi et al., 2014). The problem is to set the *reserve price* in a sequence of auctions. The proposed idea is quite intuitive: by setting the reserve price low, the second highest bid is observable. Hence it is possible to estimate the second highest bid distributions. This allows to set a better reserve price. In order to facilitate future *exploration*, the new reserve price has to be set at the lowest potentially optimal value.
Another very interesting work in auction is (He et al., 2013). In the framework of *sponsored search*, a game theoretic machine learning approach is presented. The authors propose to learn a Markov model from historical data to describe how advertisers change their bids in response to an auction mechanism. They basically want to learn the *strategic behaviour* of the players, avoiding the so called *second order effect*. Then, use the predicted future bids to learn the auction mechanism.

2.4 Dynamic Programming approach

The last possible approach is the so called *automated mechanism design*. A classical pricing setting can be seen as an *optimization* problem. The parameters of mechanism design become the input, while the output is a nonmanipulable mechanism that is optimal with respect to some objects. Basically, there are two kinds of objects. A *benevolent* designer aims to maximize the *social walfare*; a *self-interested* designer has an utility function that depends only on the final outcome. There are two main advantages with respect to the previous approaches: *flexibility*, *i.e.* the applicability to a broader set of problems and *efficiency*, shifting the burden of a classical method from humans to a machine.
We refer to (Conitzer and Sandholm, 2004). The authors present a very intuitive representation of the optimization problem, identifying the *Individual rationality* and the *Incentive compatibility* constraints. They study the case of a self-interested designer, showing that the payment-maximizing AMD problem is closely related to an interesting variant of the optimal combinatorial auction design problem, where the bidders have "best-only" preferences.

2.5 Competitive Analysis

We conclude this chapter exposing the evaluation criteria commonly used in literature. If the economists have traditionally assumed that buyers' valuations are *i.i.d.* samples from a known probability distribution, computer scientists adopted the worst-case model for buyers' valuations. Our approach, following (Kleinberg and Leighton, 2003) and (Babaioff et al., 2011), stays in the middle. The *i.i.d.* hypothesis is preserved but the probability distribution is unknown to the seller. The evaluation criterion follows is that one of *worst-case*, and, more in general, of the *Competitive Analysis*.
In online problems the input is only partially available, some relevant input data will be accessible only in the future. Competitive analysis is a powerful tool to analyze the performance of an online mechanism. The idea is to compare such an *online* mechanism to an *offline* one, which knows in advance the input. This comparison is made through the concept of *competitive ratio*.

Definition 6. ***(Competitive Ratio)*** *For a specific set of instances $\mathcal{I}$, possibly infinite, let $\mathbb{E}[\mathcal{R}(\mathcal{M})]$ be the expected value of a given mechanism $\mathcal{M}$; let bench be the expected revenue of the offline benchmark mechanism; we say that $\mathcal{M}$ has competitive ratio $c_{\mathcal{M}}$ if:*

$$c_{\mathcal{M}} = \min_{\mathcal{I}} \left\{ \frac{\mathbb{E}[\mathcal{R}(\mathcal{M})]}{bench} \right\} \tag{2.2}$$

Chapter 3

Identical Valuation

3.1 The Model

We consider the following scenario, both for Identical and Random Valuation. A seller is interested to sell a single item within a finite period D, and the item is single unit. At any $t \in [0, D]$ the seller offers a *take-it-or-leave-it price*. We use $p(t) : [0, D] \to \mathbb{R}^+$ to denote the *pricing strategy* of the seller. Agents arrive online according to a Poisson process defined by parameter λ. Since λ is a measure of the expected arrival rate of the agents, it can be easily estimated and it's considered known. It is worth underlining the fact that we don't know both the exact arriving times of the agents and the number of total arrivals, but we have probability distributions over these events, modeled with a Poisson process.
Each agent has a private valuation $v(t)$ of the item. If she arrives at time t, she buys the item *if and only if* $v(t) \geq p(t)$. It is important to stress the time dependency of the private valuations. We assume a dependency of this type: $v(t) = \tilde{v} \cdot \eta(t)$ where $\eta(t) : [0, T] \to [0, 1]$ is a non increasing function *s.t.* $\eta(0) = 1, \eta(T) = 0$. We call $\eta(t)$ *discounting valuation rate* or *discounting rate* and it is constant for all agents. $\tilde{v}$ is the *undiscounted private valuation.* The posted price mechanism in a single parameter scenario is clearly incentive compatible. This is no longer true when we are in multi-parameter scenario. Here an agent could face a strategic dilemma of whether to accept an offer early on or wait for a later offer. The seller doesn't know the private valuations and she can't learn or infer any information about them. There are different reasons why. First of all, we are in a Single-Item Single-Unit scenario, hence there's no way to infer something on an agent, given the information collected on the previous ones. Secondly, in this model, we don't ask the agents to reveal their private valuations, not even when one of them

buys the item. In that case, one can only say that the private valuation is above the price. Indeed asking agents to report their private valuations might be unrealistic. Moreover, an agent has no incentive to reveal it, as she may plan to participate in similar markets in the future.
The undiscounted private valuations of the agents $\tilde{V}$ belong to the range $\left[\tilde{V}_{\min}, \tilde{V}_{\max}\right]$, which is conventionally normalized in $[1, h]$, where $h = \frac{\tilde{V}_{\max}}{\tilde{V}_{\min}}$. From this point on we call $\tilde{v}$ the *undiscounted normalized private valuation* belonging to $[1, h]$. We consider $\tilde{v}$ drawn from a unique distribution F, which is the cumulative distribution function. F belongs to the family $\mathcal{F}$. Such a distribution is unknown to the seller.
The definition of the family $\mathcal{F}$ splits our problem into the two sub-problems: in the *Identical Valuation* scenario, $\mathcal{F}$ is the family of all possible Dirac delta function distributions over $[1, h]$, in the *Random Valuation* scenario, $\mathcal{F}$ is the family of *Monotone Hazard Rate* distributions.

The goal of the seller is to define the pricing strategy that maximizes the profit via competitive analysis with respect to an omniscient seller, which has knowledge of the distribution F and the discounting rate. From this point on we refer to her pricing strategy as the *benchmark*. In this chapter we focus on the *Identical Valuation* scenario.

3.2 Three Different Scenarios

In the *Identical Valuation* scenario, all buyers' undiscounted valuations $\tilde{v}$ are equal. This settings is already studied in literature, see (Kleinberg and Leighton, 2003) or (Babaioff et al., 2011). This value is unknown to the seller. Notice that this doesn't mean that the valuations of the buyers are equal. Indeed, even if both $\tilde{v}$ and $\eta(t)$ are equal to all agents, the arriving times are different[1]. The *benchmark* knows F, hence the undiscounted valuation $\tilde{v}$, but it does not know the arriving times.
In this chapter we analyse three different scenarios - Figure 3.1, one for each choice of the *discounting rate* $\eta(t)$:

(a) $\eta(t) = 1$: is the undiscounted valuation case in which the valuations are always equal to the undiscounted private valuations. This case has already been discussed in (Zheng et al., 2016);

[1] The Poisson distribution is also called *Law of Rare Events*, where the rare event is not the arrival of an agent, but the event that a certain person arrives. Hence in the Poisson model, the probability that two events happen at the same time instant is considered null. In our setting, this means that it is impossible that two buyers have the same valuation.

(b) $\eta(t) = 1 - \frac{t}{T}$: the linear discounting rate case;

(c) generic $\eta(t)$

Without loss of generality, we consider $D = T$. From this point on, where not specified, we use the parameter T and we call it *deadline*. Hence T is both the time within we have to sell the item and the time at which the valuations of all players vanish - except for the undiscounted valuation case.

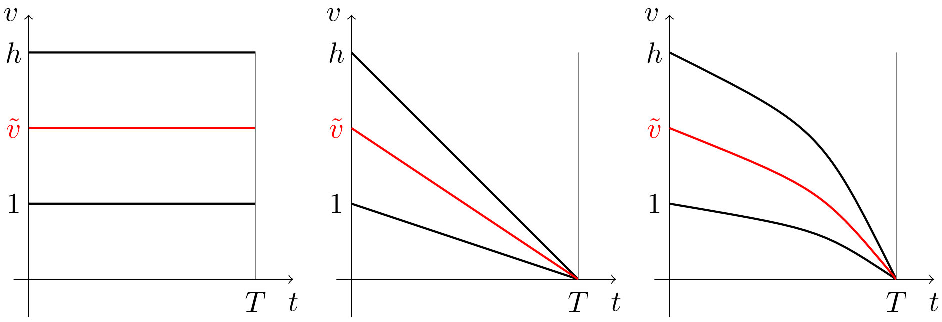

(a) $\eta(t) = 1$ (b) $\eta(t) = 1 - \frac{t}{T}$ (c) generic $\eta(t)$

Figure 3.1: The three different scenarios

3.3 The Upper Bound

We provide three different results in the following theorems, one for each scenario. The third one is a generalization of the first two.

Theorem 1. *In the Identical Valuation scenario with discounting rate $\eta(t) = 1$, the deterministic posted-price mechanism that achieves the best competitive ratio is the one that posts:*

$$p(t) = \begin{cases} h^{1-\frac{t}{t_0}} & t \in [0, t_0) \\ 1 & t \in [t_0, T] \end{cases} \tag{3.1}$$

where $t_0 \leq T$ is the time s.t. $1 - e^{-\lambda(T-t_0)} = k$ and $k \leq 1$ is a number depending on λ, h and t_0. Such a mechanism achieves a costant competitive ratio of $\frac{k}{1-e^{-\lambda T}}$.

Theorem 2. *In the Identical Valuation scenario with discounting rate $\eta(t) = 1 - \frac{t}{T}$, the deterministic posted-price mechanism that achieves the best competitive ratio is the one that posts:*

$$p(t) = \begin{cases} h \cdot \left(1 - \frac{t}{T}\right) \cdot e^{\lambda\left(1-\frac{1}{k}\right)t + \frac{\lambda}{2kT}t^2} & t \in [0, t_0) \\ 1 - \frac{t}{T} & t \in [t_0, T] \end{cases} \tag{3.2}$$

where $t_0 \leq T$ *is the time s.t.* $1 - \frac{1}{\lambda T}\left(1 + \lambda t_0 - e^{-\lambda(T-t_0)}\right) = k$ *and* $k \leq 1$ *is a number depending on* λ, h, T *and* t_0*. Such a mechanism achieves a competitive ratio of* $\frac{k}{1-\frac{1}{\lambda T}\left(1-e^{-\lambda T}\right)}$*. Let us call this mechanism* $\mathcal{M}_1$*.*

Theorem 3. *In the Identical Valuation scenario with discounting rate* $\eta(t)$*, the deterministic posted-price mechanism that achieves the best competitive ratio is the one that posts:*

$$p(t) = \begin{cases} A \cdot e^{\int f(t)dt} & t \in [0, t_0) \\ \eta(t) & t \in [t_0, T] \end{cases} \tag{3.3}$$

where A *is a constant,* $f(t) = \lambda - \frac{\lambda}{k\zeta(t)} - \frac{\zeta'(t)}{\zeta(t)}$ *and* $\zeta(t) = \frac{1}{\eta(t)}$*.* $k \leq 1$ *is a number depending on* λ, h, T *and* t_0*. Such a mechanism achieves a competitive ratio of* $\frac{k}{k_{bench}}$*.* k_{bench} *is the expected revenue of the benchmark when* $\tilde{v} = 1$*.*

Before reporting the proofs of these theorems, we first refer to the following (Babaioff et al., 2011):

Observation 1. *A posted-price mechanism should set the minimum price* $p(t) = \eta(t)$ *for a non null time interval. Otherwise the competitive ratio is zero.*

Indeed, consider a mechanism that never sets the minimum price and a set of agents having undiscounted valuation $\tilde{v} = 1$. Here, the necessary and sufficient condition $v(t) \geq p(t)$, *s.t.* an agent buys the item, never holds. Hence the expected revenue is zero, and so the competitive ratio. In our setting, using the notation of (Zheng et al., 2016), the time interval in which we set the minimum price is $[t_0, T]$.
We now present two lemmas that highlight two crucial properties which characterize optimal posted-price mechanisms in the IV setting.

Lemma 4 implies that the pricing strategy of an optimal mechanism must be such that the undiscounted price defined as $\frac{p(t)}{\eta(t)}$ is non-increasing in t, whereas Lemma 5 shows that any mechanism which always provides a constant fraction of the expected revenue of the benchmark, independently of the agents' undiscounted valuation $\tilde{v}$, is an optimal mechanism.

Lemma 4. *In the IV setting, given any deterministic posted-price mechanism* $\mathcal{M}$*, there always exists a deterministic posted-price mechanism* $\mathcal{M}'$ *with an undiscounted price* $\frac{p_{\mathcal{M}'}(t)}{\eta(t)}$ *non-increasing in* t *such that* $\mathbb{E}_v[\mathcal{R}(\mathcal{M})] \leq \mathbb{E}_v[\mathcal{R}(\mathcal{M}')]$ *for every possible agents' undiscounted valuation* $\tilde{v} \in [1, h]$*.*

Proof (Lemma 4). We only need to prove the result for mechanisms $\mathcal{M}$ whose undiscounted price $\frac{p_{\mathcal{M}}(t)}{\eta(t)}$ is *not* non-increasing in t, otherwise the statement of the lemma is trivially true.

The main idea of the proof is to let the time period $[0, T]$ be evenly partitioned into time intervals of length τ such that the undiscounted price function of $\mathcal{M}$ is constant in each interval. This is w.lo.g. if we take $\tau \to 0$. Then, there must be two consecutive time intervals, namely $I_1 := I_{s,\tau}$ and $I_2 := I_{s+\tau,\tau}$ for some starting time $s \in [0, T-\tau]$, such that there exist $p_1 < p_2 \in [1, h]$ with $\frac{p_{\mathcal{M}}(t)}{\eta(t)} = p_1$ and $\frac{p_{\mathcal{M}}(t)}{\eta(t)} = p_2$ during I_1 and I_2, respectively (otherwise the undiscounted price would be non-increasing). Now, let us define a mechanism $\mathcal{M}'$ whose undiscounted price function is the same as that of $\mathcal{M}$, except for the fact that $\frac{p_{\mathcal{M}'}(t)}{\eta(t)} = p_2$ during I_1 and $\frac{p_{\mathcal{M}'}(t)}{\eta(t)} = p_1$ during I_2 (*i.e.*, intuitively, we exchange the values in the two intervals so as to make the undiscounted price non-increasing in that window of time, see Figure 3.2 in case of linear discounting rate).

We show that the expected revenue provided by $\mathcal{M}'$ is always greater than or equal to that achieved by $\mathcal{M}$, as long as $\tau \to 0$. In order to make compare the expected revenues of the two mechanisms, it is sufficient to focus on the window of time $I_1 \cup I_2$, where their price functions differ. Given p_1 and p_2, we can partition the agents' undiscounted valuations $\tilde{v} \in [1, h]$ into three different subsets, as follows:

- $\tilde{v} < p_1$, implying that $\tilde{v}\,\eta(t) < p_{\mathcal{M}}(t)$ and $\tilde{v}\,\eta(t) < p_{\mathcal{M}'}(t)$ for every time instant $t \in I_1 \cup I_2$;
- $p_1 \leq \tilde{v} \leq p_2$, implying that $p_{\mathcal{M}}(t) \leq \tilde{v}\,\eta(t) \leq p_{\mathcal{M}'}(t)$ for every time instant $t \in I_1$ and $p_{\mathcal{M}'}(t) \leq \tilde{v}\,\eta(t) \leq p_{\mathcal{M}}(t)$ for every time instant $t \in I_2$;
- $\tilde{v} > p_2$, implying that $\tilde{v}\,\eta(t) > p_{\mathcal{M}}(t)$ and $\tilde{v}\,\eta(t) > p_{\mathcal{M}'}(t)$ for every time instant $t \in I_1 \cup I_2$.

In the first case, $\mathbb{E}[\mathcal{R}(\mathcal{M}(\tilde{v}))] - E[\mathcal{R}(\mathcal{M}'(\tilde{v}))] = 0$, since both $\mathcal{M}$ and $\mathcal{M}'$ achieve an expected revenue equal to 0 during the time window $I_1 \cup I_2$, given that the item is never sold in that window (as both $p_{\mathcal{M}}(t)$ and $p_{\mathcal{M}'}(t)$ are always higher than the agents' valuation $\tilde{v}\,\eta(t)$). As for the second case, let us assume $p_1 < \tilde{v} < p_2$ (since the cases $\tilde{v} = p_1$ and $\tilde{v} = p_2$ are analogous). Then, $\mathcal{M}$ can sell the item only during the interval I_1, while $\mathcal{M}'$ can sell the item only during the other interval I_2. Thus, the difference $\mathbb{E}[\mathcal{R}(\mathcal{M}(\tilde{v}))] - E[\mathcal{R}(\mathcal{M}'(\tilde{v}))]$ is equal to:

$$\int_{s}^{s+\tau} p_1 \eta(t)\, \lambda e^{-\lambda t} dt - \int_{s+\tau}^{s+2\tau} p_1 \eta(t)\, \lambda e^{-\lambda t} dt$$

which goes to 0 as long as $\tau \to 0$, given that η is continuous. Finally, in the third case, we can compute the difference between the expected revenues of the two mechanisms $\mathbb{E}[\mathcal{R}(\mathcal{M}(\tilde{v}))] - E[\mathcal{R}(\mathcal{M}'(\tilde{v}))]$ as follows:

$$\begin{aligned}
&\int_{s}^{s+\tau} p_1\eta(t)\,\lambda e^{-\lambda t}dt + \int_{s+\tau}^{s+2\tau} p_2\eta(t)\,\lambda e^{-\lambda t}dt + \\
&\qquad - \int_{s}^{s+\tau} p_2\eta(t)\,\lambda e^{-\lambda t}dt - \int_{s+\tau}^{s+2\tau} p_1\eta(t)\,\lambda e^{-\lambda t}dt = \\
&= (p_1 - p_2)\int_{s}^{s+\tau} \eta(t)\,\lambda e^{-\lambda t}dt - (p_1 - p_2)\int_{s+\tau}^{s+2\tau} \eta(t)\,\lambda e^{-\lambda t}dt = \\
&= (p_1 - p_2)\left[\int_{s}^{s+\tau} \eta(t)\,\lambda e^{-\lambda t}dt - \int_{s+\tau}^{s+2\tau} \eta(t)\,\lambda e^{-\lambda t}dt\right]
\end{aligned}$$

which is less than or equal to 0 as $\tau \to 0$, by continuity of η.

By re-iterating the procedure on all the pairs of consecutive infinitesimal intervals defined as I_1 and I_2 (each time using the last mechanism $\mathcal{M}'$ as the new $\mathcal{M}$), we can render the undiscounted price function non-increasing, obtaining a final mechanism $\mathcal{M}'$ such that $\mathbb{E}[\mathcal{R}(\mathcal{M}(\tilde{v}))] \leq E[\mathcal{R}(\mathcal{M}'(\tilde{v}))]$ for every possible agents' undiscounted valuation $\tilde{v} \in [1, h]$. □

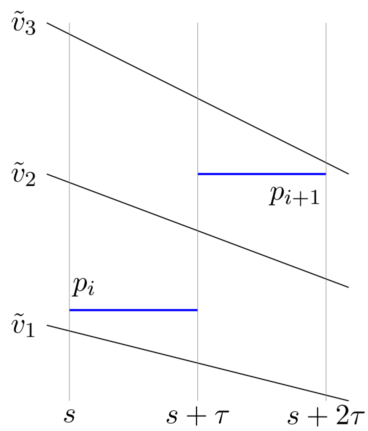

(a) The first mechanism

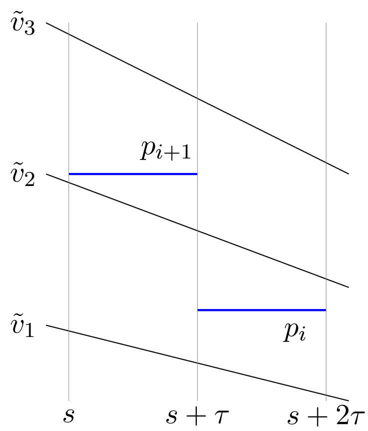

(b) The second mechanism

Figure 3.2: Non-increasing mechanism: linear $\eta(t)$

Lemma 5. *In the IV setting, let $\mathcal{M}$ be a deterministic posted-price mechanism whose pricing strategy $p_{\mathcal{M}}$ satisfies $p_{\mathcal{M}}(0) = h$ and $p_{\mathcal{M}}(t) = \eta(t)$ for $t \in [t_0, T] \subseteq [0, T]$. If the ratio $\frac{\mathbb{E}[R(\mathcal{M}(\tilde{v}))]}{\text{bench}(\tilde{v})}$ does not depend on the agents' undiscounted valuation $\tilde{v}$, then $\mathcal{M}$ is an optimal mechanism.*

Proof (Lemma 5). By contradiction, suppose that $\mathcal{M}$ is *not* optimal, *i.e.*, there exists another deterministic posted-price mechanism $\mathcal{M}'$ such that $c_{\mathcal{M}'} > c_{\mathcal{M}}$. According to Observation 1 and Lemma 4, $\mathcal{M}'$ must be defined by a pricing strategy $p_{\mathcal{M}'}$ such that the undiscounted price $\frac{p_{\mathcal{M}'}(t)}{\eta(t)}$ is non-increasing in t and the minimum price is selected for a time interval $[t_0', T] \subseteq [0, T]$ having non-zero length (recall that $c_{\mathcal{M}}(\tilde{v}) > 0$ does not depend on $\tilde{v}$ and $c_{\mathcal{M}} = \min_{v \in [1,h]} c\mathcal{M}(\tilde{v})$).

Case $t_0' \geq t_0$. Let us consider the undiscounted valuation $\tilde{v} = 1$. Then, we have that the expected revenue of mechanism $\mathcal{M}$ is $\mathbb{E}[\mathcal{R}(\mathcal{M}(\tilde{v}))] = \int_{t_0}^{T} \eta(t)\lambda e^{-\lambda t} dt$ (accounting for the case in which an agent arrives at $t \geq t_0$ and buys the item at price $\eta(t)$), which is greater than or equal to the expected revenue of mechanism $\mathcal{M}'$, defined as $\mathbb{E}[\mathcal{R}(\mathcal{M}'(\tilde{v}))] = \int_{t_0'}^{T} \eta(t)\lambda e^{-\lambda t} dt$. Intuitively, $\mathbb{E}[\mathcal{R}(\mathcal{M}(\tilde{v}))] \geq \mathbb{E}[\mathcal{R}(\mathcal{M}'(\tilde{v}))]$ since $\mathcal{M}'$ posts the minimum price for a period of time shorter than that of $\mathcal{M}$. Therefore, it holds $c_{\mathcal{M}'} \leq c_{\mathcal{M}'}(\tilde{v}) \leq c_{\mathcal{M}}(\tilde{v}) \leq c_{\mathcal{M}}$, which is a contradiction.

Case $t_0' < t_0$. First, suppose that there exists a time instant $t' \in [0, t_0']$ defined as $t' \coloneqq \sup\{t \in [0, t_0'] \mid p_{\mathcal{M}}(t) < p_{\mathcal{M}'}(t)\}$, *i.e.*, the last time instant in which $p_{\mathcal{M}}(t)$ changes from being less than $p_{\mathcal{M}'}(t)$ to being larger than or equal to $p_{\mathcal{M}'}(t)$. Clearly, it holds $p_{\mathcal{M}}(t) \geq p_{\mathcal{M}'}(t)$ for every $t \in [0, T] : t > t'$. Moreover, let us consider the agents' undiscounted valuation $\tilde{v}' \in [1, h]$ such that $\tilde{v}'\eta(t') = p_{\mathcal{M}}(t')$ and focus on the case in which $p_{\mathcal{M}}(t) = p_{\mathcal{M}'}(t)$ (as the other cases are analogous). Notice that, for every time instant $t \leq t'$, mechanism $\mathcal{M}'$ cannot sell the item, since, by using Lemma 4, we get:

$$\tilde{v}\eta(t) < \tilde{v}p_{\mathcal{M}'}(t)\frac{\eta(t')}{p_{\mathcal{M}'}(t')} = \tilde{v}p_{\mathcal{M}'}(t)\frac{\eta(t')}{p_{\mathcal{M}}(t')} \leq \tilde{v}p_{\mathcal{M}'}(t)\frac{\eta(t')}{\tilde{v}\eta(t')} \leq p_{\mathcal{M}'}(t).$$

Additionally, with an analogous reasoning we can shows that, for all the times $t \in [0, T] : t > t'$, both mechanisms may sell the item, but the price posted by $\mathcal{M}'$ is always less than or equal to that chosen by $\mathcal{M}$, with a non-empty time interval in which the former is strictly less than the latter (as $t_0' < t_0$). Thus, in this case, it holds $c_{\mathcal{M}}(\tilde{v}) > c_{\mathcal{M}'}(\tilde{v})$, which implies that $c_{\mathcal{M}'} < c_{\mathcal{M}}$, a contradiction. Finally, it remains to analyze the case in which a time instants t' defined above does *not* exist. Since the undiscounted price functions are non-increasing by Lemma 4 and $t_0' < t_0$, it must be the case that there is no intersection point between the two functions. Hence, it must be $p_{\mathcal{M}}(t) > p_{\mathcal{M}'}(t)$ for all $t \in [0, t_0]$, which implies that $c_{\mathcal{M}'} < c_{\mathcal{M}}$ by taking $\tilde{v} = h$. This leads to a contradiction. $\square$

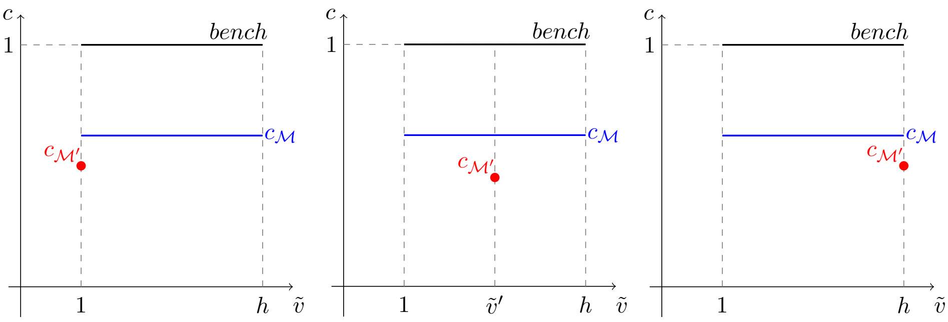

(a) if $t'_0 > t_0$ (b) if $\exists\, t'$ (c) if $t'_0 < t_0$, $\nexists\, t'$

Figure 3.3: Parts of the Proof

3.3.1 The Benchmark

The last step before the proofs is a discussion about the omniscient *benchmark*. In the *Identical Valuation* scenario, it knows both the discounting rate $\eta(t)$ and the point mass distribution F, and so the undiscounted private valuation $\tilde{v}$ of the agents. Nevertheless, it doesn't know the arriving times of the customers. The strategy of the *benchmark* is to post the price $p_{bench}(t) = \tilde{v}\eta(t)$ for $t \in [0, T]$. We give the following:

Observation 2. *The expected revenue of the benchmark is linearly dependent on the undiscounted private valuation $\tilde{v}$.*

It's clear that the revenue of the benchmark is not a fixed value, but it's a random variable depending on the Poisson process. Let's compute the Expected Value of the Revenue of the *benchmark*. We refer to it as *bench*:

$$bench = \int_0^T p_{bench}(t)\lambda e^{-\lambda t} dt = \int_0^T \tilde{v}\eta(t)\lambda e^{-\lambda t} dt = \tilde{v}\int_0^T \eta(t)\lambda e^{-\lambda t} dt = \tilde{v}\cdot k_{bench} \tag{3.4}$$

It is evident the linear behaviour. The slope is called k_{bench} and it represents the expected revenue of the *benchmark* when the agents have the minimum undiscounted valuation $\tilde{v}$.

3.3.2 Proofs

We now present the proofs of the Theorem 1, Theorem 2, Theorem 3.
The basic idea is the following: *since*

i. the best deterministic posted-price mechanism has constant competitive ratio for all the undiscounted private valuation $\tilde{v}$ (Lemma 5) *and*

ii. the *benchmark* has a linear dependency in the variable $\tilde{v}$ (Observation 2) *then*

We search for a deterministic posted price mechanism whose expected revenue is linearly dependent on the variable $\tilde{v}$.

We give the following definition, next we proceed to the proofs.

Definition 7. *Given a pricing strategy $p(t)$ and a valuation $v(t)$ we call t^* the time instant s.t. $t^* \in [0, t_0]$ and $p(t^*) = v(t^*)$.*

Given an undiscounted valuation $\tilde{v}$, t^* is the time instant after which the condition $v(t) \geq p(t)$ is satisfied.

Proof (Theorem 1). Our aim is to define a pricing strategy $p(t) : [0, T] \to \mathbb{R}^+$ *s.t.* the Expected Revenue can be written as $\mathbb{E}[R(\tilde{v})] = k \cdot \tilde{v}$ for some k. We already know that we should offer the minimum price for $[t_0, T]$. Hence, we can restrict the problem to the pricing strategy $p(t) : [0, t_0] \to \mathbb{R}^+$. Since $\eta(t) = 1$ we can write $\mathbb{E}[R(\tilde{v})] = \int_{t^*}^{t_0} p(t)\lambda e^{-\lambda(t-t^*)}dt + \int_{t_0}^{T} \lambda e^{-\lambda(t-t^*)}dt$. While $\tilde{v}$ can be written as $p(t^*)$. We have the following equation to solve, where the function $p(t)$ is unknown.

$$e^{\lambda t^*} \int_{t^*}^{t_0} p(t) \cdot \lambda e^{-\lambda t} dt + e^{\lambda t^*} \int_{t_0}^{T} \lambda e^{-\lambda t} dt = k \cdot p(t^*) \tag{3.5}$$

Let us derive with respect to t^*. The left side becomes:

$$\begin{aligned}
\frac{\partial \mathbb{E}[R(\tilde{v})]}{\partial t^*} &= \frac{\partial}{\partial t^*} \left(e^{\lambda t^*} \int_{t^*}^{t_0} p(t) \cdot \lambda e^{-\lambda t} dt + e^{\lambda t^*} \int_{t_0}^{T} \lambda e^{-\lambda t} dt \right) = \\
&= e^{\lambda t^*} \frac{\partial}{\partial t^*} \left(\int_{t^*}^{t_0} p(t)\lambda e^{-\lambda t} dt \right) + \lambda e^{\lambda t^*} \int_{t^*}^{t_0} p(t)\lambda e^{-\lambda t} dt + \\
&\quad + \lambda e^{\lambda t^*} \int_{t_0}^{T} \lambda e^{-\lambda t} dt
\end{aligned}$$

Rewrite the term $\frac{\partial}{\partial t^*} \left(\int_{t^*}^{t_0} p(t)\lambda e^{-\lambda t} dt \right)$ as $\frac{\partial}{\partial t^*} G(t^*)$ where

$$G(t^*) = \int_{t^*}^{t_0} p(t)\lambda e^{-\lambda t} dt = -\int_{t_0}^{t^*} p(t)\lambda e^{-\lambda t} dt = \int_{t_0}^{t^*} g(t) dt$$

Applying the fundamental theorem of calculus:

$$\frac{\partial}{\partial t^*} G(t^*) = g(t^*) = -p(t)\lambda e^{-\lambda t}$$

Hence the derivative of the expected revenue becomes:

$$\begin{aligned}\frac{\partial \mathbb{E}[R(\tilde{v})]}{\partial t^*} &= e^{\lambda t^*}\frac{\partial G(t^*)}{\partial t^*} + \lambda e^{\lambda t^*}\int_{t^*}^{t_0} p(t)\lambda e^{-\lambda t}dt + \lambda e^{\lambda t^*}\int_{t_0}^{T}\lambda e^{-\lambda t}dt \\ &= e^{\lambda t^*} g(t^*) + \lambda\left[e^{\lambda t^*}\int_{t^*}^{t_0} p(t)\lambda e^{-\lambda t}dt + \lambda e^{\lambda t^*}\int_{t_0}^{T}\lambda e^{-\lambda t}dt\right]\end{aligned}$$

Notice the term in square brackets is exactly the left side of the Equation 3.5. We can now write the derivative of (3.5):

$$-e^{\lambda t^*} p(t^*)\lambda e^{-\lambda t^*} + \lambda\left[k \cdot p(t^*)\right] = kp'(t^*)$$

$$\Rightarrow p'(t^*) = \lambda\left(1 - \frac{1}{k}\right)p(t^*) \tag{3.6}$$

The solution of the differential equation (3.6) is:

$$p(t) = A \cdot e^{\int \lambda\left(1-\frac{1}{k}\right)dt} = A \cdot e^{\lambda\left(1-\frac{1}{k}\right)t} \tag{3.7}$$

Adding the boundary conditions *s.t.* $p(0) = h$ and $p(t_0) = 1$, we can derive $A = h$ and $k = \frac{\lambda t_0}{\lambda t_0 + \ln(h)}$. Finally we have to define t_0. It must be the time instant *s.t.* $\mathbb{E}[R(1)] = k$. Hence:

$$\int_0^{T-t_0} \lambda e^{-\lambda t}dt = 1 - e^{-\lambda(T-t_0)} = k$$

We can finally write the deterministic posted-price mechanism that achieves the best competitive ratio:

$$p(t) = \begin{cases} h^{1-\frac{t}{t_0}} & t \in [0, t_0) \\ 1 & t \in [t_0, T] \end{cases}$$

This concludes the proof. □

Corollary 5.1. *The above posted-price mechanism achieves a constant competitive ratio of* $\frac{1-e^{-\lambda(T-t_0)}}{1-e^{-\lambda T}}$.

Proof (Corollary 5.1). The proof is trivial.

$$\frac{\mathbb{E}[R(\tilde{v})]}{\text{bench}(\tilde{v})} = \frac{k \cdot \tilde{v}}{k_{bench} \cdot \tilde{v}} = \frac{k}{k_{bench}} = \frac{\mathbb{E}[R(1)]}{\text{bench}(1)} = \frac{\int_0^{T-t_0} \lambda e^{-\lambda t}dt}{\int_0^{T} \lambda e^{-\lambda t}dt} = \frac{1 - e^{-\lambda(T-t_0)}}{1 - e^{-\lambda T}}$$

This concludes the proof. □

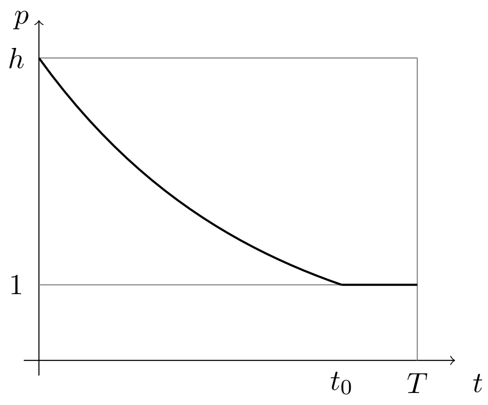

*Figure 3.4: The Upper Bound (*I*)*

The proof of the Theorem 2 has the same structure as the previous one. Here we have a linear *discounting rate* $\eta(t) = 1 - \frac{t}{T}$.

Proof (Theorem 2). Our aim is to define a pricing strategy $p(t) : [0,T] \to \mathbb{R}^+$ *s.t.* the Expected Revenue can be written as $\mathbb{E}[R(\tilde{v})] = k \cdot \tilde{v}$ for some k. We already know that we should offer the minimum price for $[t_0, T]$. Hence, we can restrict the problem to the pricing strategy $p(t) : [0, t_0] \to \mathbb{R}^+$. The Expected Revenue is $\mathbb{E}[R(\tilde{v})] = e^{\lambda t^*} \int_{t^*}^{t_0} p(t) \cdot \lambda e^{-\lambda t} dt + e^{\lambda t^*} \int_{t_0}^{T} \left(1 - \frac{t}{T}\right) \lambda e^{-\lambda t} dt$, while $\tilde{v}\left(1 - \frac{t^*}{T}\right) = p\left(t^*\right)$. We have the following equation to solve, where the function $p(t)$ is unknown.

$$e^{\lambda t^*} \int_{t^*}^{t_0} p(t) \cdot \lambda e^{-\lambda t} dt + e^{\lambda t^*} \int_{t_0}^{T} \left(1 - \frac{t}{T}\right) \lambda e^{-\lambda t} dt = k \frac{T}{T - t^*} p\left(t^*\right) \qquad (3.8)$$

Let us derive with respect to t^*. The left side:

$$\frac{\partial \mathbb{E}[R(\tilde{v})]}{\partial t^*} = e^{\lambda t^*} \frac{\partial G\left(t^*\right)}{\partial t^*} + \lambda \left[e^{\lambda t^*} \int_{t^*}^{t_0} p(t) \lambda e^{-\lambda t} dt + \lambda e^{\lambda t^*} \int_{t_0}^{T} \left(1 - \frac{t}{T}\right) \lambda e^{-\lambda t} dt \right]$$

Where $G\left(t^*\right) = \int_{t^*}^{t_0} p(t) \lambda e^{-\lambda t} dt = -\int_{t_0}^{t^*} p(t) \lambda e^{-\lambda t} dt = \int_{t_0}^{t^*} g(t) dt$. We can apply the fundamental theorem of calculus. Moreover the square brackets term is exactly the expected revenue. The derivative of the right side of the 3.8 is:

$$\frac{\partial}{\partial t^*} \left(\frac{kT}{T - t^*} p\left(t^*\right) \right) = \frac{kT}{T - t^*} p'\left(t^*\right) + \frac{kT}{\left(T - t^*\right)^2} p\left(t^*\right)$$

We can write the derivative *with respect to* t^* of (3.8) getting the following differential equation:

$$p'\left(t^*\right) = \left[\lambda - \frac{\lambda \left(T - t^*\right)}{kT} - \frac{1}{T - t^*} \right] p\left(t^*\right) \qquad (3.9)$$

The solution of (3.9) is:

$$p(t) = A \cdot e^{\int \left[\lambda - \frac{\lambda(T-t)}{kT} - \frac{1}{T-t}\right] dt} = A \cdot e^{\lambda\left(1-\frac{1}{k}\right)t + \frac{\lambda}{2kT}t^2 + \ln(T-t)}$$

Adding the boundary conditions *s.t.* $p(0) = h$ and $p(t_0) = 1 - \frac{t_0}{T}$, we can derive $A = \frac{h}{T}$ and $k = \lambda t_0 \frac{2T-t_0}{2T(\lambda t_0 + \ln(h))}$.
Finally we have to define t_0. It must be the time instant *s.t.* $\mathbb{E}[R(1)] = k$. Hence:

$$\int_0^{T-t_0} \left(1 - \frac{t}{T}\right) \lambda e^{-\lambda t} dt = 1 - \frac{1}{\lambda T}\left(1 + \lambda t_0 - e^{-\lambda(T-t_0)}\right) = k$$

We can finally write the deterministic posted-price mechanism that achieves the best competitive ratio:

$$p(t) = \begin{cases} h \cdot \left(1 - \frac{t}{T}\right) \cdot e^{\lambda\left(1-\frac{1}{k}\right)t + \frac{\lambda}{2kT}t^2} & t \in [0, t_0) \\ 1 - \frac{t}{T} & t \in [t_0, T] \end{cases}$$

This concludes the proof. □

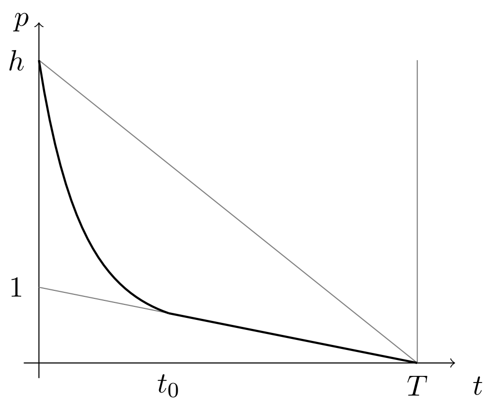

*Figure 3.5: The Upper Bound (*II*): mechanism $\mathcal{M}_1$*

Corollary 5.2. *$\mathcal{M}_1$ achieves a constant competitive ratio of* $\frac{1-\frac{1}{\lambda T}\left(1+\lambda t_0 - e^{-\lambda(T-t_0)}\right)}{1-\frac{1}{\lambda T}\left(1-e^{-\lambda T}\right)}$.

Proof (Corollary 5.2). The proof is trivial.

$$\frac{\mathbb{E}[R(\tilde{v})]}{\text{bench}(\tilde{v})} = \frac{k \cdot \tilde{v}}{k_{bench} \cdot \tilde{v}} = \frac{k}{k_{bench}} = \frac{\mathbb{E}[R(1)]}{\text{bench}(1)} = \frac{\int_0^{T-t_0} \left(1 - \frac{t}{T}\right) \lambda e^{-\lambda t} dt}{\int_0^{T} \left(1 - \frac{t}{T}\right) \lambda e^{-\lambda t} dt}$$

$$= \frac{1 - \frac{1}{\lambda T}\left(1 + \lambda t_0 - e^{-\lambda(T-t_0)}\right)}{1 - \frac{1}{\lambda T}\left(1 - e^{-\lambda T}\right)}$$

This concludes the proof. □

Finally we give the proof of the Theorem 3. This is a generalization of the previous two. Hence it has the same structure.

Proof (Theorem 3). Our aim is to define a pricing strategy $p(t) : [0,T] \to \mathbb{R}^+$ *s.t.* the expected revenue can be written as $\mathbb{E}[R(\tilde{v})] = k \cdot \tilde{v}$ for some k. We already know that we should offer the minimum price for $[t_0, T]$. Hence, we can restrict the problem to the pricing strategy $p(t) : [0, t_0] \to \mathbb{R}^+$. The Expected Revenue is $\mathbb{E}[R(\tilde{v})] = e^{\lambda t^*} \int_{t^*}^{t_0} p(t) \cdot \lambda e^{-\lambda t} dt + e^{\lambda t^*} \int_{t_0}^{T} \eta(t) \lambda e^{-\lambda t} dt$. while $\tilde{v}\eta(t^*) = p(t^*)$. Hence $\tilde{v} = \frac{1}{\eta(t^*)} p(t^*)$. We call $\zeta(t) = \frac{1}{\eta(t)}$. We have the following equation to solve, where the function $p(t)$ is unknown.

$$e^{\lambda t^*} \int_{t^*}^{t_0} p(t) \cdot \lambda e^{-\lambda t} dt + e^{\lambda t^*} \int_{t_0}^{T} \eta(t) \lambda e^{-\lambda t} dt = k\zeta(t^*) p(t^*) \tag{3.10}$$

We derive the equation 3.10 and apply the Fundamental theorem of calculus, getting the following differential equation:

$$-\lambda p(t^*) + \lambda k \zeta(t^*) p(t^*) = k\zeta(t^*)(t^*) p'(t^*) + k\zeta'(t^*) p(t^*)$$

$$\Rightarrow p'(t^*) = \left[\lambda - \frac{\lambda}{k\zeta'(t^*)} - \frac{\zeta'(t^*)}{\zeta(t^*)}\right] p(t^*) \tag{3.11}$$

The solution of (3.11) is:

$$p(t) = A \cdot e^{\int \left[\lambda - \frac{\lambda}{k\zeta'(t)} - \frac{\zeta'(t)}{\zeta(t)}\right] dt}$$

Adding the boundary conditions *s.t.* $p(0) = h$ and $p(t_0) = \eta(t_0)$, we can derive A and k. We also define t_0 *s.t.* $\mathbb{E}[R(1)] = k$.

We can finally write the deterministic posted-price mechanism that achieves the best competitive ratio:

$$p(t) = \begin{cases} A \cdot e^{\int \left[\lambda - \frac{\lambda}{k\zeta'(t)} - \frac{\zeta'(t)}{\zeta(t)}\right] dt} & t \in [0, t_0) \\ \eta(t) & t \in [t_0, T] \end{cases}$$

This concludes the proof. □

Corollary 5.3. *The above posted-price mechanism achieves a constant competitive ratio of* $\frac{\int_0^{T-t_0} \eta(t) \lambda e^{-\lambda t} dt}{\int_0^T \eta(t) \lambda e^{-\lambda t} dt}$.

Proof (Corollary 5.3). The proof is trivial.

$$\frac{\mathbb{E}[R(\tilde{v})]}{\text{bench}(\tilde{v})} = \frac{k \cdot \tilde{v}}{k_{bench} \cdot \tilde{v}} = \frac{k}{k_{bench}} = \frac{\mathbb{E}[R(1)]}{\text{bench}(1)} = \frac{\int_0^{T-t_0} \eta(t)\lambda e^{-\lambda t} dt}{\int_0^{T} \eta(t)\lambda e^{-\lambda t} dt}$$

This concludes the proof. □

3.4 Pricing and Convolution

The aim of this section is to show a different approach to the *Identical Valuation* problem. We show a possible interpretation of the Expected Revenue of a posted-price mechanism. We discuss the result in the following theorem. The proof is very simple and immediate.

Theorem 6. *In Identical Valuation scenario, consider a posted-price mechanism* $p(t) : [0, D] \to \mathbb{R}^+$ s.t. $\forall\ \tilde{v} \in [1, h]$*, there exists an unique intersection point* t^* *between* $v(t)$ *and* $p(t)$*. For any probability density function over the arrivals of the agents* $g(t)$ *and for any decreasing valuation rate* $\eta(t)$*;*
The expected revenue of the mechanism is the convolution of the functions $P(t)$ *and* $G(t)$ *where* $P(t) = p(-t) \cdot [H(D-t) - H(-t)]$, $G(t) = g(t) \cdot [H(t) - H(t-D)]$ *and* $H(t)$ *is the Heaviside step function.*

$$\mathbb{E}[R(\tilde{t})] = (P * G)(\tilde{t}) = \int_{-\infty}^{\infty} P(\tilde{t} - \tau) \cdot G(\tau) d\tau \tag{3.12}$$

where $\tilde{t} = -t^*$.

Proof (Theorem 6). The convolution equation is simply derived by the definition of the expected revenue. We call $\tilde{p}(t) = p(-t)$

$$\mathbb{E}[R(\tilde{t})] = \int_0^{D-t^*} p\left(t + t^*\right) g(t) dt =$$

$$= \int_0^{D-t^*} \tilde{p}\left(-t - t^*\right) g(t) dt = \int_0^{D+\tilde{t}} \tilde{p}(-t + \tilde{t}) g(t) dt =$$

$$= \int_{-\infty}^{\infty} \tilde{p}(-t + \tilde{t})[H(D-t) - H(-t)] \cdot g(t)[H(t) - H(t-D)] dt = (P * G)(\tilde{t})$$

This concludes the proof. □

This theorem suggests an interpretation of the expected revenue of a pricing strategy. The expected revenue is the weighted average of the function $P(t)$ - the reverse pricing curve - at the time $\tilde{t}$ where the weighting is given by $G(-t)$ - the reverse probability density function - shifted by amount $\tilde{t}$.

3.5 Variation of the Model: Abrupt Changes

In this section we show a possible variation of the model. We take into account the effects of the seasonability on the market. We consider these effects through abrupt changes in the valuations of the buyers, *i.e.* in the *discounting rate*.
Let us consider the following modified model. There are N τ-lengthed phases in which the valuations of the agents are continuous over time. From a phase to the next one, an abrupt change occurs. In each phase i there is a maximum valuation v_i^{max} and a minimum valuation v_i^{min}, where $v_i^{min} \geq v_{i+1}^{min}$. Without loss of generality, we assume v_i^{max} be $h \cdot v_i^{min}$, $\forall i$. Indeed, h can be the maximum ratio between valuations over the phases. We define the *undiscounted valuation* $\tilde{v}$ as the ratio between the valuation of an agent in a certain phase normalized with respect to the minumum valuation of that phase. In this new model, the *Identical Valuation* scenario suggests that each agent has a private undiscounted valuation $\tilde{v}$.
In each phase i, the valuations have time-dependency $\tilde{v}v_i^{min}\eta_i(t)$, where $\eta_i(t)$ is the continuous *discounting rate* of the i^{th} phase. In principle these functions should change from a phase to another one. Initially we consider the undiscounted case, in which $\eta_i(t) = 1$, for $i = 1, \dots, N$. Secondly, we will consider the linear and general discounting rate case.
Notice that the only assumption made on v_i^{min} is that they are non-increasing in i. This means that the magnitudes of these parameters could be generic. We consider two possible situations. In the first one, we assume $v_i^{min} > hv_{i+1}^{min}$, hence we call it *non-overlapping phases* case. In the second situation, we do not make this assumption.

Briefly, we discuss the role of the *benchmark*. Also in this scenario, the clairvoyant mechanism knows the undiscounted valuation $\tilde{v}$ of the agents. This means that in each phase i, she can offer the price $\tilde{v}v_i^{min}\eta_i(t)$, collecting:

$$bench = \tilde{v} \cdot \left[v_1^{\min} \int_0^{\tau} \lambda e^{-\lambda t} dt + \cdots + v_N^{\min} \int_{(N-1)\tau}^{N\tau} \lambda e^{-\lambda t} dt \right]$$

Where *bench* is the expected revenue of the benchmark mechanism. Note that the square brackets term is constant with respect to $\tilde{v}$. Hence, also in this case the expected revenue of the benchmark has a linear dependency with respect to the undiscounted valuations.

$$bench = k_{bench}\tilde{v} \tag{3.13}$$

The *Backward Induction* approach here seems to be the best approach to solve this problem. Unfortunately, it turns out to be not the case. Indeed, it is not possible to consider the phases independently from each other. The performance dramatically changes whether considering only one or several phases.

3.5.1 Non-Overlapping Phases

We study the case in which $v_i^{min} > hv_{i+1}^{min}$. This means that the agents' valuations are well distinct from different phases, Figure (3.6). Given a certain phase i, it seems reasonable to offer the minimum price v_i^{min}. Indeed, we will never have the possibility of posting such a price in future. We define the pricing strategy of the i^{th} phase to be: $p_i(t) = \tilde{v}\tilde{p}_i(t)$. We can think the single phase as a special instance of the model defined in section 3.1. For these reason, a first step is to apply the optimal mechanism (3.1) in the single phase, we call this mechanism Δ_1. One can object that a better choice should be that one of posting the pricing strategy of (3.1) in the overall time horizon. We define such a mechanism Δ_N.

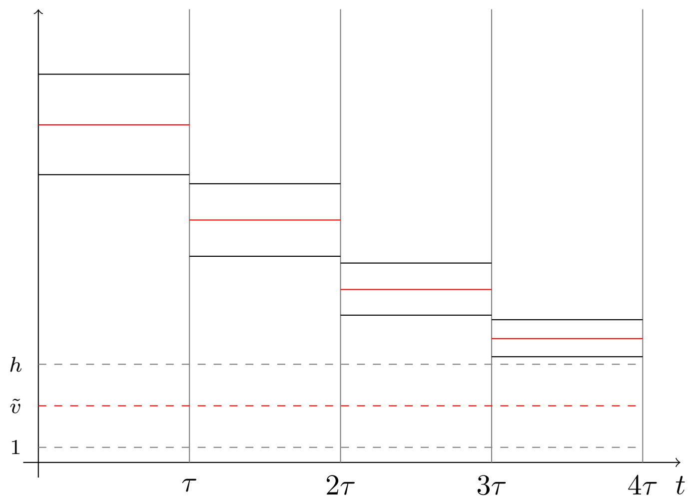

Figure 3.6: Abrupt Changes: Non-Overlapping Phases

For simplicity we compare these two choices in the case $N = 2$. Then the results will be generalized.

We define mechanism Δ_1 as follows:

$$\Delta_1 : \tilde{p}_i(t) = \begin{cases} h^{1^{\frac{t-(i-1)\tau}{t_1}}} & t \in [(i-1)\tau, (i-1)\tau + t_1) \\ 1 & t \in [(i-1)\tau + t_1, i\tau) \end{cases} \qquad i = 1, 2 \tag{3.14}$$

Where t_1 is such that $[t_1, \tau]$ is the time interval in which the bottom price is posted, as prescribed by mechanism (3.1).

Mechanism Δ_2 is defined as follows:

$$\Delta_2 : p(t) = \begin{cases} v_1^{\min} \cdot h^{1\frac{t}{t_2}} & t \in [0, \tau) \\ v_2^{\min} \cdot h^{1-\frac{t}{t_2}} & t \in [\tau, t_2) \\ v_2^{\min} & t \in [t_2, 2\tau] \end{cases} \tag{3.15}$$

Where t_2 is such that $[t_2, 2\tau]$ is the time interval in which the bottom price is posted, as prescribed by mechanism (3.1). See Figure (3.7).

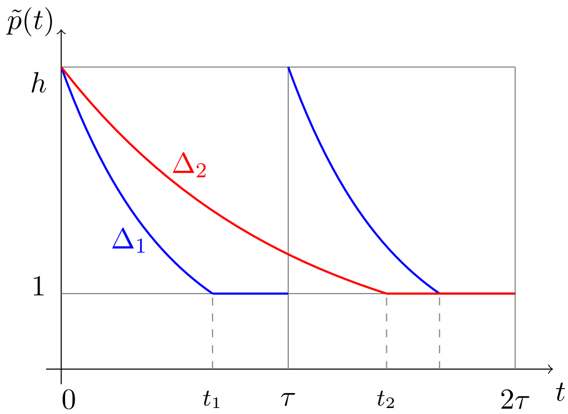

Figure 3.7: Mechanisms Δ_1 and Δ_2

We make a comparison of these two mechanisms. We denote the expected revenue of Δ_1 and Δ_2 as $\mathbb{E}[\mathcal{R}_{\Delta_1}(\tilde{v})]$ and $\mathbb{E}[\mathcal{R}_{\Delta_2}(\tilde{v})]$, respectively, when agents' undiscounted valuation is $\tilde{v}$; we give the following propositions:

Proposition 1. $\mathbb{E}[\mathcal{R}_{\Delta_1}(1)] \geq \mathbb{E}[\mathcal{R}_{\Delta_2}(1)]$.

Proposition 2. $\mathbb{E}[\mathcal{R}_{\Delta_2}(h)] \geq \mathbb{E}[\mathcal{R}_{\Delta_1}(h)]$.

We cannot say that one of these two mechanisms has a greater competitve ratio than the other one. However, because of these propositions, we are sure that no one of these two mechanisms always dominates the other one.

Proof (Proposition 1). When $\tilde{v} = 1$, the two mechanisms can sell the item only in the time interval in which they set the bottom price. Δ_1 offers the bottom price for an amount of time of $2(\tau - t_1)$, Δ_2 for $(2\tau - t_2)$. We claim $2(\tau - t_1) \geq (2\tau - t_2)$. Note that t_1 and t_2 are defined as follows:

$$\frac{\lambda}{\ln h} t_1 = e^{\lambda(\tau - t_1)} - 1 \tag{3.16}$$

$$\frac{\lambda}{\ln h} t_2 = e^{\lambda(2\tau - t_2)} - 1 \tag{3.17}$$

Suppose by contradiction that $(2\tau - t_2) \geq 2(\tau - t_1)$. It follows that $e^{(2\tau - t_2)} \geq e^{2(\tau - t_1)}$ and consequently

$$\frac{\lambda}{\ln h} t_2 \geq e^{2(\tau - t_1)} - 1 \tag{3.18}$$

From 3.16, it follows:

$$e^{\lambda(\tau - t_1)} = \frac{\lambda}{\ln h} t_1 + 1 \Rightarrow e^{2\lambda(\tau - t_1)} = \left(\frac{\lambda}{\ln h} t_1\right)^2 + 2\frac{\lambda}{\ln h} t_1 + 1$$

Hence, from 3.18:

$$\frac{\lambda}{\ln h} t_2 \geq \left(\frac{\lambda}{\ln h} t_1\right)^2 + 2\frac{\lambda}{\ln h} t_1 \Rightarrow t_2 \geq \frac{\lambda}{\ln h} t_1^2 + 2t_1$$

$$\Rightarrow 2\tau - \frac{\lambda}{\ln h} t_1^2 - 2t_1 \geq 2\tau - t_2 \geq 2(\tau - t_1)$$

Where the last inequality follows from the hypothesis, hence:

$$-\frac{\lambda}{\ln h} t_1^2 \geq 0 \tag{3.19}$$

And the contradiction is reached, being $\lambda > 0$ the parameter of the Poisson process. Hence, the claim is true. Note now that $\mathbb{E}[\mathcal{R}_{\Delta_2}(1)] = v_2^{min} \int_0^{2\tau - t_2} \lambda e^{-\lambda t} dt$ and $\mathbb{E}[\mathcal{R}_{\Delta_1}(1)] = v_1^{min} \int_0^{\tau - t_1} \lambda e^{-\lambda t} dt + v_2^{min} \int_{\tau - t_1}^{2(\tau - t_1)} \lambda e^{-\lambda t} dt$. We can graphically represent these two values:

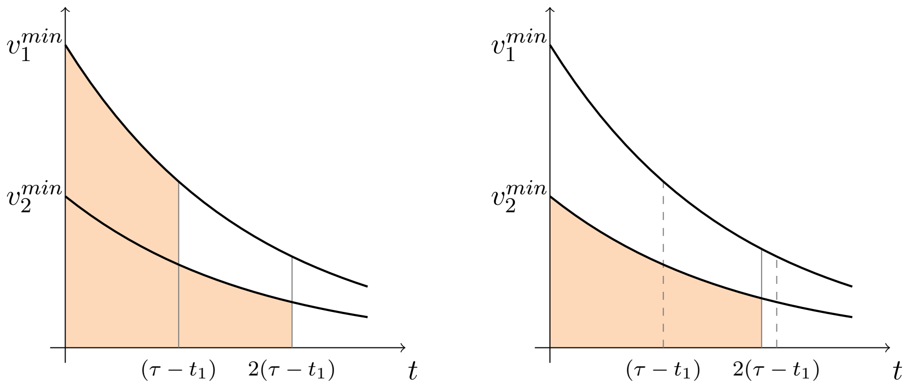

(a) Mechanism Δ_1 *(b) Mechanism Δ_2*

Figure 3.8: Expected revenue in $\tilde{v} = 1$

Since the claim is true, the integral time of the first one is always smaller than the second one. It is now evident that $\mathbb{E}[\mathcal{R}_{\Delta_1}(1)] \geq \mathbb{E}[\mathcal{R}_{\Delta_2}(1)]$. This concludes the proof. □

Before proposing the proof of Proposition (2), we report here a theorem by (Steffensen, 1925).

Theorem 7. *(Steffensen, 1925) Let g_1 and g_2 be functions defined on $[a, b]$ such that*

$$\int_a^x g_1(t)dt \geq \int_a^x g_2(t)dt \quad \text{for all} \quad x \in [a, b] \tag{3.20}$$

and

$$\int_a^b g_1(t)dt = \int_a^b g_2(t)dt \tag{3.21}$$

Let f be a decreasing function on $[a, b]$, then

$$\int_a^b f(x)g_1(x)dx \geq \int_a^b f(x)g_2(x)dx. \tag{3.22}$$

.

Proof (Proposition 2). Suppose that $v_2^{min} = v_1^{min}$. Note that, in the second phase, Δ_1 is always above Δ_2. Hence, this assumption is in favour of Δ_1. From the proof of Proposition (1), it follows that $2(\tau - t_1) > 2\tau - t_2 \Rightarrow t_2 > 2t_1$. Consider a new mechanism $\Delta_1^{'}$ such that the corresponding $t_2 = 2t_1^{'}$. It is clear that $\mathbb{E}[\mathcal{R}_{\Delta_1^{'}}(h)] \geq \mathbb{E}[\mathcal{R}_{\Delta_1}(h)]$. Let us consider another mechanism $\Delta_1^{''}$, which is equal to $\Delta_1^{'}$, but it switches the time intervals $[t_1, \tau)$ and $[\tau, \tau + t_1)$. From the definition of the expected revenue of a pricing strategy when $\tilde{v} = h$ it is evident that $\mathbb{E}[\mathcal{R}_{\Delta_1^{''}}(h)] \geq \mathbb{E}[\mathcal{R}_{\Delta_1^{'}}(h)]$. Consequently

$$\mathbb{E}[\mathcal{R}_{\Delta_1^{''}}(h)] \geq \mathbb{E}[\mathcal{R}_{\Delta_1}(h)] \tag{3.23}$$

We have to prove now that in $\tilde{v} = h$ the expected performance of Δ_2 exceeds the expected performance of $\Delta_1^{''}$.
Note that $\tilde{p}_{\Delta_2}(t)$ and $\tilde{p}_{\Delta_1^{''}}(t)$ satisfy Inequality (3.20) and Equality (3.21). Where $a = 0$ and $b = 2\tau$. Being $\lambda e^{-\lambda t}$ the decreasing function f of the Theorem 7, it follows that:

$$\int_0^{2\tau} \lambda e^{-\lambda t} \tilde{p}_{\Delta_2}(t)dt \geq \int_a^b \lambda e^{-\lambda t} \tilde{p}_{\Delta_1^{''}}(t)dt \tag{3.24}$$

Where the left side of Inequality (3.24) is the expected revenue of the mechanism Δ_2 in $\tilde{v} = h$, the right side that one of $\Delta_1^{''}$ in $\tilde{v} = h$. From Inequality (3.23), the claim of the theorem is proved. □

Mechanism Δ_1 is the mechanism (3.1) applied to time time horizon of one phase, Δ_2 is the same mechanism applied to time horizon of two phases. We can now extend this reasoning, and consider a generic mechanism Δ_i. The previous propositions can be generalized:

Proposition 3. $\mathbb{E}[\mathcal{R}_{\Delta_i}(1)] \geq \mathbb{E}[\mathcal{R}_{\Delta_{2i}}(1)]$.

Proposition 4. $\mathbb{E}[\mathcal{R}_{\Delta_{2i}}(h)] \geq \mathbb{E}[\mathcal{R}_{\Delta_i}(h)]$.

The proof of Proposition (3) easily follows the proof of Proposition (1).

Sketch of Proof (Proposition 4). We apply exactly the same reasonings of the proof of Proposition (2). Suppose all the minimum valuations v_j^{min} to be equal to $v_{\frac{i}{2}+1}^{min}$. Suppose $t_{2i} = 2t_i^{'}$ and a new mechanism $\Delta_i^{''}$ built switching the pricing strategies of the time intervals $[t_i, \frac{i}{2}\tau)$ and $[\frac{i}{2}\tau, \frac{i}{2}\tau + t_1)$. Then, the steps are the same of the proof of Proposition (2). $\square$

None of these mechanisms that dominates the other ones. We now propose a new mechanism. The basic idea is to divide the undiscounted valuations $\tilde{v} \in [1, h]$ into N sections. Each phase handles just one section. The first phase manages the highest valuations' section, the second one manages the second highest valuations' section and so on. We call this new mechanism Δ and it can handle also the overlapping phases.

3.5.2 Overlapping Phases

In this section the only assumption is that the minimum valuations v_i^{min} are non-increasing in i. The first result is a generalization of Lemma (4).

Lemma 8. *In the Identical Valuation Setting, a decreasing pricing strategy always dominates an increasing one.*

Proof (Lemma 8). We have already proved the claim of this theorem for the prices posted inside a certain phase, Lemma 4. We have to consider now the prices from a phase to another one.
Consider three undiscounted valuations $\tilde{v}_1$, $\tilde{v}_2$ and $\tilde{v}_3$ such that $\tilde{v}_1 < \tilde{v}_2 < \tilde{v}_3$ and two consecutive phase. We make a comparison between two mechanisms evaluating them in the infinitesimal time instant $[i\tau - dt, i\tau]$ just before the end of the i^{th} phase and the infinitesimal time instant $(i\tau, i\tau + dt]$ just after the beginning of the i^{th} phase. The first mechanism posts two constant prices p_i and p_{i+1} consequently; while the second mechanism posts p_{i+1} in $[i\tau - dt, i\tau]$ and then p_i in $(i\tau, i\tau + dt]$. Suppose that $p_i < p_{i+1}$ such that in

both the phases p_{i+1} lies below the valuation $\tilde{v}_3$ and above the valuation $\tilde{v}_2$, and p_i lies below the valuation $\tilde{v}_2$ and above the valuation $\tilde{v}_1$, as the figure 3.9 shows.

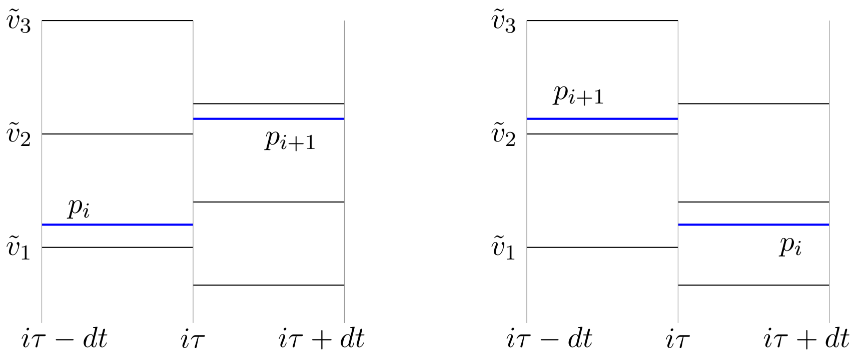

(a) The first mechanism *(b) The second mechanism*

Figure 3.9: Non-increasing mechanism: Phases

We now show that the expected revenue of the second mechanism is always greater than the one of the first mechanism in the considered undiscounted valuations. In the continuous case the external valuations converge to the central one, and the duration τ tends to zero.

$$\begin{aligned}
\mathbb{E}\left[R_1\left(\tilde{v}_1\right)\right] - \mathbb{E}\left[R_2\left(\tilde{v}_1\right)\right] &= \left[p_i\left(1 - e^{\lambda\tau}\right) + p_{i+1}\left(1 - e^{\lambda\tau}\right)e^{\lambda\tau}\right] + \\
&\quad - \left[p_{i+1}\left(1 - e^{\lambda\tau}\right) + p_i\left(1 - e^{\lambda\tau}\right)e^{\lambda\tau}\right] < 0 \\
\mathbb{E}\left[R_1\left(\tilde{v}_2\right)\right] - \mathbb{E}\left[R_2\left(\tilde{v}_2\right)\right] &= p_{i+1}\left(1 - e^{\lambda\tau}\right) - p_{i+1}\left(1 - e^{\lambda\tau}\right) = 0 \\
\mathbb{E}\left[R_1\left(\tilde{v}_3\right)\right] - \mathbb{E}\left[R_2\left(\tilde{v}_3\right)\right] &= 0
\end{aligned}$$

Note that, as in Lemma 4, the undiscounted valuations can be partitioned in three sets. The corresponding expected revenue always is one of the tree computations above. Hence, the expected revenue of the second mechanism is always greater than the one of the first mechanism. □

We now present mechanism Δ in the case of $\eta_i(t) = 1$, $\forall i$.

Mechanism Δ : given N phases, Δ divides the phases in three sets. The first set $\mathcal{S}_1$ is composed by the first $m-1$ phases, the second set $\mathcal{S}_2$ is composed by only the m^{th} phase, the third set $\mathcal{S}_3$ by the last $N-m$ phases. The pricing strategy adopted in a certain phase depends on the set at which that phase belongs. Consider that the actual price posted by the mechanism is $p_i(t) = v_i^{min}\tilde{p}_i(t)$, where each price $\tilde{p}_i(t)$ is defined with respect a relative time metric, $\tilde{p}_i : [0, \tau] \to [1, h]$. This means that $p_i(\bar{t})$ is the price posted at time $(i-1)\tau + \bar{t}$.

We defined Δ as follows:

$$\Delta : \quad \tilde{p}_i(t) = \begin{cases} A_i e^{\lambda\left(1-\frac{v_i^{min}}{k}\right)t} & \text{if } i \in \mathcal{S}_1 \\ A_m e^{\lambda\left(1-\frac{v_m^{\min}}{k}\right)t} \mathbb{I}_{[0,t_0)}\{t\} + \mathbb{I}_{[t_0,\tau)}\{t\} & \text{if } i \in \mathcal{S}_2 \\ 1 & \text{if } i \in \mathcal{S}_3 \end{cases} \tag{3.25}$$

Where the m parameters A_i are specified such that $\tilde{p}_i(i\tau) = \tilde{p}_{i+1}(i\tau)$. These constraints are fundamental. They do not imply a continuous pricing strategy, but they impose that all the undiscounted valuations $\tilde{v}$ will be covered. For the same reason, the parameter k is specified by imposing $\tilde{p}_1(0) = h$. Finally parameter t_0 is defined as

$$k = e^{\lambda((m-1)\tau+t_0)} \left[v_m^{\min} \int_{t_0}^{m\tau} \lambda e^{\lambda t} dt + \sum_{j=m+1}^{N} v_j^{\min} \int_{j\tau}^{(j+1)\tau} \lambda e^{\lambda t} dt \right] \tag{3.26}$$

The basic idea is that dividing the valuations in phases, each of the phases handles a portion of valuations independently from the others. In this way, we can adapt definition 7.

Definition 8. *Given a pricing strategy $\tilde{p}(t)$ and a valuation $\tilde{v}$ we call t^* the time instant s.t. $t^* \in [0, m\tau + t_0]$ and $\tilde{p}(t^*) = \tilde{v}$.*

Theorem 9. *In Identical Valuation setting, Δ reaches a constant competitive ratio of $\frac{k}{k_{bench}}$.*

Proof (Theorem 9). Let us distinguish the three sets of phases:

- $\mathcal{S}_3$: the set of the last phases in which the seller posts the bottom prices. The number of phases belonging to this set is not yet defined. It depend on t_0.
- $\mathcal{S}_2$: we follow the proof of Theorem (1). The expected revenue of the mechanism depends on $\tilde{v}$ and, consequently, on t^*. Remember that each phase handles his own portion of valuations, hence his own portion of t^*. The contribution of the m^{th} phase in the total expected revenue is the left side of the following equality. We impose such a contribution to be equal to $k\tilde{p}_m(t^*)$.

$$e^{\lambda t^*} v_m^{\min} \int_{t^*}^{t_0} \tilde{p}_m(t) \cdot \lambda e^{-\lambda t} dt + e^{\lambda(t^*-t_0)} k \tag{3.27}$$

Deriving with respect to t^* and using the fundamental theorem of calculus, we can provide the following differential equation:

$$\tilde{p}'_m(t^*) = \lambda\left(1 - \frac{v_m^{min}}{k}\right)p(t^*) \tag{3.28}$$

Whose solution is:

$$\tilde{p}_m(t) = A_m \cdot e^{\lambda\left(1 - \frac{v_m^{min}}{k}\right)t} \tag{3.29}$$

We add the boundary condition *s.t.* $\tilde{p}_m(t_0) = 1$, defining A_m.

- $\mathcal{S}_1$: the contribution of the $(m-1)^{th}$ phase to the expected revenue of the mechanism can be expressed as follows:

$$e^{\lambda t^*} v_{m-1}^{\min} \int_{t^*}^{(m-1)\tau} \tilde{p}_m(t) \cdot \lambda e^{-\lambda t} dt + e^{\lambda t^*}\left[v_m^{\min} \int_{(m-1)\tau}^{t_0} \tilde{p}_m(t) \cdot \lambda e^{-\lambda t} dt + e^{-\lambda t_0} k\right] \tag{3.30}$$

  Note the similarity between Equation (3.27) and (3.30). Only the square brackets term increases of a constant. Applying the same procedure, we obtain:

$$\tilde{p}_{m-1}(t) = A_{m-1} \cdot e^{\lambda\left(1 - \frac{v_{m-1}^{min}}{k}\right)t} \tag{3.31}$$

  We add the boundary condition *s. t.* $\tilde{p}_{m-1}((m-1)\tau) = \tilde{p}_m((m-1)\tau)$, defining A_{m-1}. The same procedure is iterate up to the first phase where the boundary conditions are $\tilde{p}_1(\tau) = \tilde{p}_2(\tau)$ and $\tilde{p}_1(0) = h$, defining A_1 and k.

Hence, mechanism Δ is built such that it has a linear expected revenue with respect the undiscounted valuation $\tilde{v}$, where the linear coefficient is k. Also the benchmark has the same linear behaviour, (3.13). Hence, Δ reaches a constant competitive ratio of $\frac{k}{k_{bench}}$. This concludes the proof. □

A final comment: mechanism Δ can be implemented also in the case of *linear* and *general* discounting rate, following the same steps of Theorem (9) and applying the computation of Theorem (2) and (3), respectively.

Chapter 4

Identical Valuation: Optimization

This chapter is dedicated to a mathematical programming approach to the problem of *Identical Valuation*. The main result in this context is the *Maximum Violation Algorithm*, that will be generalized to the *Random Valuation* scenario in the next chapter. This is the first reason why we study the *Identical Valuation* setting also from an optimization point of view. The second is to meet a common and realistic need: usually sellers can't modify the posted-price in every time instant. Hence a continuous time pricing function cannot be used in practice.
We study the case with linear *discounting rate*, *i.e.* $\eta(t) = 1 - \frac{t}{T}$. As already discussed, this case can represent a very common realistic economic situation. Moreover, this scenario is easily extendible to the degenerate *discounting rate* $\eta(t) = 1$ case. Indeed, the latter is a good approximation of the first one when we have a very high T and an arbitrary D. Hence, all the results of this chapter must be considered valid also for the degenerate *discounting rate* case.
This is a non-linear optimization problem.

4.1 Piecewise Linear Optimization

We consider a finite number of initial private valuation $\tilde{v} \in [1, h]$. The *undiscounted valuation vector* $\mathbf{v}= [\tilde{v}_0, \ldots, \tilde{v}_{n-1}]$ is composed of n discrete undiscounted valuation *s.t.* $\tilde{v}_0 = 1$ and $\tilde{v}_{n-1} = h$. We want to find a pricing curve $p(t)$, *s.t.* $p(0) = h$ and $p(t_0) = 1 - \frac{t_0}{T}$, that maximizes the competitive ratio. For any valuation $\tilde{v}_i \cdot (1 - \frac{t}{T})$ - a straight line, we assume only one intersection point with the pricing curve $p(t)$. We call t_i the time instant

when the pricing curve intersects the valuation $\tilde{v}_i$. We collect all these time instances in the vector $\mathbf{t}= [t_0, \ldots, t_{n-1}]$. Note that $t_{n-1} = 0$
In this first section we assume a piecewise linear curve $p(t)$, composed by n segments - the last one is $1 - \frac{t}{T}$ for $t \in [t_0, T]$. The points of discontinuity occur at the time instants t_i. We define each segment as follows:

$$r_i(t) = A_i \cdot t + B_i \quad t_i \leq t < t_{i-1} \tag{4.1}$$

Each segment have to respect two constraints: the starting and ending points of every segment has to lie on two consecutive valuations. Hence, for $i = 1, \cdots, n-1$:

$$r_i\left(t_i\right) = A_i \cdot t_i + B_i = \tilde{v}_i \left(1 - \frac{t_i}{T}\right)$$

$$r_i\left(t_{i-1}\right) = A_i \cdot t_{i-1} + B_i = \tilde{v}_i \left(1 - \frac{t_{i-1}}{T}\right)$$

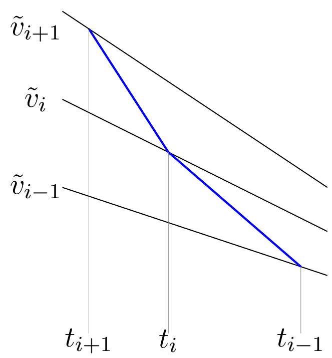

Figure 4.1: Piecewise Linear Pricing Curve

We are looking for the posted-price mechanism that maximizes the competitive ratio. This means that is a *maxmin* problem:

$$\max_{p} \left\{ \min_{\tilde{v}_i \in \mathbf{v}} c\left(\tilde{v}_i\right) \right\} \tag{4.2}$$

Where $c(\tilde{v}_i) = \frac{\mathbb{E}[R(\tilde{v}_i)]}{\text{opt}(\tilde{v}_i)}$. Our idea is to solve (4.2) by observing that $\mathbb{E}\left[R\left(\tilde{v}_i\right)\right] = \mathbb{E}\left[R\left(\tilde{v}_i, t_i\right)\right]$. We can write (4.2) as:

$$\max_{\mathbf{t}} \left\{ \min_{\tilde{v}_i \in \mathbf{v}} c\left(\tilde{v}_i\right) \right\} \tag{4.3}$$

The expected revenue of the piecewise linear mechanism is the followig:

$$
\begin{aligned}
\mathbb{E}\left[R\left(\tilde{v}_i, t_i\right)\right] = e^{\lambda t_i} \quad \cdot & \left\{\sum_{j=0}^{i-1}\left[A_{i-j}\left(e^{-\lambda t_{i-j}}\left(t_{i-j}+\frac{1}{\lambda}\right)-e^{-\lambda t_{i-j-1}}\left(t_{i-j-1}+\right.\right.\right.\right. \\
& \left.\left.\left.+\frac{1}{\lambda}\right)\right)+B_{i-j}\left(e^{-\lambda t_{i-j}}-e^{-\lambda t_{i-j-1}}\right)\right]+ \\
& \left.+e^{-\lambda D}\left(\frac{D}{T}-1+\frac{1}{\lambda T}\right)+e^{-\lambda t_0}\left(1-\frac{t_0}{T}-\frac{1}{\lambda T}\right)\right\}
\end{aligned}
$$

The mathematical programming problem is the following:

$$
\begin{array}{lcl}
 & \mathbf{max} \quad \alpha & \\
\textbf{subject to:} & c(\tilde{v}_i, t_i) \geq \alpha & \forall i = 0, n-1 \\
 & A_i \cdot t_i + B_i - \tilde{v}_i\left(1 - \frac{t_i}{T}\right) = 0 & \forall i = 1, n-1 \\
 & A_i \cdot t_{i-1} + B_i - \tilde{v}_i\left(1 - \frac{t_{i-1}}{T}\right) = 0 & \forall i = 1, n-1 \\
 & t_i - t_{i-1} \leq 0 & \forall i = 1, n-1 \\
 & t_0 \leq D, \quad t_{n-1} = 0 &
\end{array}
$$

Where the variables are: α, A_i, B_i and t_i. There are $2n$ inequality constraints and $2(n-1)$ equality constraints.
We can rewrite the problem reducing the number of variables. We derive the expressions of A_i and B_i from the second and third constraints, getting:

$$
A_i = \frac{\tilde{v}_i - \tilde{v}_{i-1} - \frac{1}{T}\left(\tilde{v}_i t_i - \tilde{v}_{i-1} t_{i-1}\right)}{t_i - t_{i-1}}
$$

$$
B_i = \tilde{v}_i\left(1 - \frac{t_{i-1}}{T}\right) - \frac{\tilde{v}_i - \tilde{v}_{i-1} - \frac{1}{T}\left(\tilde{v}_i t_i - \tilde{v}_{i-1} t_{i-1}\right)}{t_i - t_{i-1}} t_{i-1}
$$

These values are then insert in the expression of the expected revenue. The variables of the new problem are only α and t_i, the number of constraints is equal to $2n$. The final *Piecewise Linear Optimization* problem is the following:

$$
\begin{array}{lcl}
 & \mathbf{max} \quad \alpha & \\
\textbf{subject to:} & c(\tilde{v}_i, t_i) \geq \alpha & \forall i = 0, n-1 \\
 & t_i - t_{i-1} \leq 0 & \forall i = 1, n-1 \\
 & t_0 \leq D, \quad t_{n-1} = 0 &
\end{array}
$$

4.2 Maximum Violation Algorithm (I)

The Piecewise Linear Optimization solves $\max_p \{\min_{\tilde{v}_i \in \mathbf{v}} c(\tilde{v}_i)\}$, that is a *maxmin* problem,. Unfortunately, the search space $\mathbf{v}$ is a finite set of n elements. We can't consider an infinite set of undiscounted valuations, because for each element of the set there is an inequality constraint. The idea of our algorithm is to perform an *intelligent discretization* of the undiscounted valuations, *i.e.* starting from a (low) finite number of undiscounted valuations, define the piecewise linear pricing curve that maximizes the competitive ratio in the valuations considered, then search - in a continuous space - for the worst performance valuation (the maximum violation), and run a piecewise linear optimization considered the finite set of undiscounted valuations of the previous step and the worst performance valuation. We itearte this procedure, until a certain tolerances is reached. We call this algorithm the *Maximum Violation Algorithm (I)*, or *MVA(I)*

Algorithm 1 `Maximum Violation Algorithm(I)`

1: $\mathbf{v} \leftarrow [1, h]$;
2: $(\alpha, \mathbf{t}) \leftarrow$ *PiecewiseLinearOptimization*$(\boldsymbol{v}, \lambda, T, D)$;
3: $\bar{v} \leftarrow \min_v c(v, \mathbf{t})$;
4: $\beta \leftarrow c(\bar{v}, \mathbf{t})$;
5: **if** $\alpha - \beta \geq \varepsilon$ **then**
6: $\quad \mathbf{v} \leftarrow \mathbf{v} \cup \{\bar{v}\}$;
7: $\quad$ **go to** 2;
8: **end if**
9: **return** $\mathbf{t}, \beta$

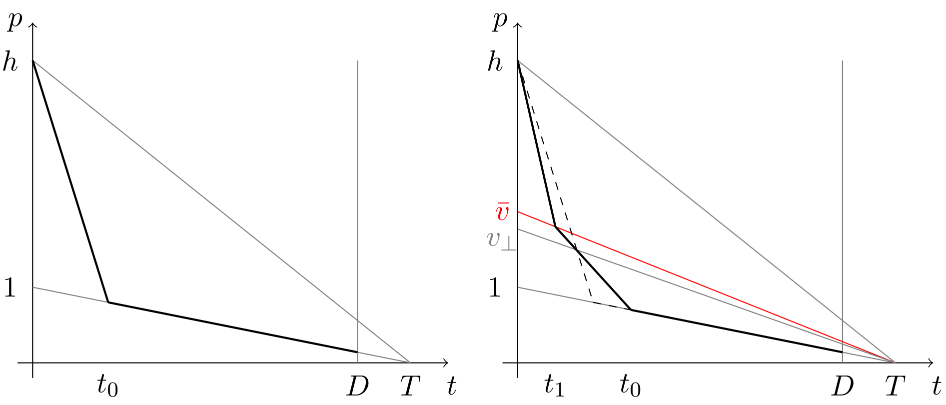

(a) Step 1 *(b) Step 2*

Figure 4.2: Maximum Violation Algorithm (I): first two steps

4.2.1 Analysis of the Algorithm

We study the behaviour of the *MVA* from a step to the next one. We take into account the first two steps, shown in Figure 4.2. First, notice that at *i-th* step $\alpha^{(i)}$ represents an upper bound of the competitive ratio, while $\beta^{(i)}$ a lower bound. As i grows, also the number of constraints grows, hence the upper bound $\alpha^{(i)}$ decreases. In the following analysis we want to show that the lower bound $\beta^{(i)}$ increases while i grows.
Take the first two steps. Since $\alpha^{(1)} \geq \alpha^{(2)}$, $t_0^{(2)} \geq t_0^{(1)}$. t_0 is proportional to the expected revenue when all the customers have the bottom valuation. The higher is t_0, the smaller is the upper bound. On the other hand, in the undiscounted valuation $\bar{v}$, we reach a higher Expected Revenue *wrt* the previous step. Indeed, fixing $t_0^{(2)}$, we move $r_1(t_1)$ towards zero, which is the best choice for $\bar{v}$ with a fixed t_0. It is clear that the optimization performs a trade-off, we gain something in the undiscounted valuation $\bar{v}$ - the worst one in the previous step - loosing in the undiscounted valuation 1 - the best one in the previous step.
What happen in between? We call $\tilde{v}_\perp$ the undiscounted valuation in which the two pricing strategies cross. From $\bar{v}$ to $\tilde{v}_\perp$ the performance of $p^{(2)}$ is better. From $\tilde{v}_\perp$ to 1, $p^{(2)}$ gains more and more *wrt* $p^{(1)}$. The undiscounted valuation 1 represents the point in which there is the maximum distance between $\mathbb{E}[R^{(1)}]$ and $\mathbb{E}[R^{(2)}]$. Hence $\forall v$ *s.t.* $1 < v \leq \bar{v}$:

$$\mathbb{E}[R^{(1)}(1)] - \mathbb{E}[R^{(2)}(1)] \geq \mathbb{E}[R^{(1)}(\tilde{v})] - \mathbb{E}[R^{(2)}(\tilde{v})]$$

$$\frac{\mathbb{E}[R^{(1)}(1)] - \mathbb{E}[R^{(2)}(1)]}{k_{opt}} \geq \frac{\mathbb{E}[R^{(1)}(\tilde{v})] - \mathbb{E}[R^{(2)}(\tilde{v})]}{\tilde{v} \cdot k_{opt}}$$

$$\frac{\mathbb{E}[R^{(1)}(1)]}{k_{opt}} - \frac{\mathbb{E}[R^{(2)}(1)]}{k_{opt}} \geq \frac{\mathbb{E}[R^{(1)}(\tilde{v})]}{\tilde{v} \cdot k_{opt}} - \frac{\mathbb{E}[R^{(2)}(\tilde{v})]}{\tilde{v} \cdot k_{opt}}$$

But notice that $\frac{\mathbb{E}[R^{(2)}(1)]}{k_{opt}}$ is greater or equal than the upper bound $\alpha^{(2)}$.

We consider now the range $[\bar{v}, h]$. The optimization performs a trade-off between the h - the best one in the previous step - and $\bar{v}$ - the worst one in the previous step. This trade-off is performed through the variable t_1. Moving t_1 towards zero, as already discussed, the performance in $\bar{v}$ increases, while in h decreases. Fixing the other variables, the comparison between $\mathbb{E}[R^{(1)}(v)]$ and $\mathbb{E}[R^{(2)}(v)]$ in the range $[\bar{v}, h]$ is the same comparison discussed before in the range $[1, \tilde{v}_\perp]$, but with reversed roles. This means that $\bar{v}$ is the undiscounted valuation in which the difference $\mathbb{E}[R^{(2)}] - \mathbb{E}[R^{(1)}]$ is maximum. This also means that h is the undiscounted valuation in which the difference $\mathbb{E}[R^{(1)}]-\mathbb{E}[R^{(2)}]$ is maximum. But, again, this happens exactly in the upper bound.

4.3 Comparison and Experimental Results

The *Maximum Violation Algorithm (I)* gives a guarantee on the result. The outputs of the algorithm are a posted-price mechanism and its competitive ratio. The value of the competitive ratio is not an approximation. This is the great advantage *wrt* the *Piecewise Linear Optimization*, whose outputs are a pricing strategy and its optimistic estimate of the competitive ratio. We cannot give any guarantee, except a measure of the robustness of the solution, given by the number of undiscounted valuations considered (hence the number of constraints). Clearly the guarantee insured by the *MVA* has a cost in terms of time.
In this section we show some experimental results, using the solver *BARON*. We consider three different settings:

- $h = 1.5$, $\lambda = 3$, $T = D = 5$;
- $h = 2$, $\lambda = 3$, $T = 7$, $D = 5$;
- $h = 3$, $\lambda = 5$, $T = 7$, $D = 3$;

Notice that, even if the three settings are quite different, they are in some sense comparable. Indeed, the expected number of customers λD is constant.

	$\varepsilon = 0.05$	$\varepsilon = 0.01$	$\varepsilon = 0.005$	$\varepsilon = 0.005$
MVA	$\alpha = 0.847944$ $\beta = 0.827606$ time $= 8.58531$ $n = 3$	$\alpha = 0.841996$ $\beta = 0.835173$ time $= 32.569$ $n = 5$	$\alpha = 0.841218$ $\beta = 0.836614$ time $= 73.1012$ $n = 6$	$\alpha = 0.839921$ $\beta = 0.83893$ time $= 316.198$ $n = 13$
PL Opt	$\alpha = 0.848847$ time $= 11.6688$ $n = 3$	$\alpha = 0.841574$ time $= 20.0735$ $n = 6$	$\alpha = 0.840668$ time $= 20.0797$ $n = 8$	$\alpha = 0.839807$ time $= 20.1091$ $n = 18$

Table 4.1: Scenario I*: $h = 1.5$, $\lambda = 3$, $T = D = 5$*

It is clear from the data the main difference between these two approaches. The *MVA* pays the guarantee with a bigger computational time, while the computational time of the *PL* optimization never exceeds 22 seconds. However, notice that we have no way to be sure what is the actual tolerance achieved by the optimization. Indeed, in the tables above there is

	$\varepsilon = 0.05$	$\varepsilon = 0.01$	$\varepsilon = 0.005$	$\varepsilon = 0.005$
MVA	$\alpha = 0.826299$ $\beta = 0.788559$ time $= 4.20144$ $n = 3$	$\alpha = 0.812035$ $\beta = 0.80251$ time $= 59.8762$ $n = 6$	$\alpha = 0.809641$ $\beta = 0.805173$ time $= 123.578$ $n = 9$	$\alpha = 0.808444$ $\beta = 0.807451$ time $= 442.192$ $n = 18$
PL Opt	$\alpha = 0.820124$ time $= 15.7917$ $n = 4$	$\alpha = 0.810176$ time $= 20.0801$ $n = 10$	$\alpha = 0.808755$ time $= 20.3472$ $n = 18$	$\alpha = 0.808242$ time $= 20.2961$ $n = 39$

Table 4.2: Scenario II*: $h = 2$, $\lambda = 3$, $T = 7$, $D = 5$;*

	$\varepsilon = 0.05$	$\varepsilon = 0.01$	$\varepsilon = 0.005$	$\varepsilon = 0.005$
MVA	$\alpha = 0.811141$ $\beta = 0.788172$ time $= 4.5067$ $n = 5$	$\alpha = 0.803279$ $\beta = 0.79628$ time $= 111.92$ $n = 9$	$\alpha = 0.801826$ $\beta = 0.796886$ time $= 199.466$ $n = 13$	$\alpha = 0.800341$ $\beta = 0.799405$ time $= 574.941$ $n = 32$
PL Opt	$\alpha = 0.816666$ time $= 16.9787$ $n = 5$	$\alpha = 0.801973$ time $= 20.1041$ $n = 19$	$\alpha = 0.800799$ time $= 20.2005$ $n = 33$	$\alpha = 0.800203$ time $= 21.4115$ $n = 93$

Table 4.3: Scenario III*: $h = 3$, $\lambda = 5$, $T = 7$, $D = 3$;*

a hidden computational time used to check the tolerances, and such a computational time is always greater than the one of the algorithm. However, data also show that if we consider a very high number n of undiscounted valuation, the tolerances achieved are satisfactory. For example notice table 4.3, *PL* optimization achieves a $\varepsilon = 0.005$ tolerance with $n = 93$ discrete undiscounted valuations. The time required for this goal is almost 21.5 seconds.
Notice also the behaviour of the upper and lower bounds α and β. As previously discussed, as the algorithm steps increase, α always decreases and β always increases.

Finally we want to graphically show the outputs of the two approaches. In the figure 4.3 the two posted-price mechanisms are compared with the upper bound (3.2) $\mathcal{M}_1$, *i.e.* the best possible mechanism.

We use the first one of the three settings described above, with a small difference. To make more visible the plot, the value of the maximum undiscounted valuation is $h = 4$.
It is clear that both of the approaches achieve a very good approximation of the best mechanism possible, even with a very small number of constraints ($n = 4$). In particular *MVA* gets a result very close to that one of $\mathcal{M}_1$. It

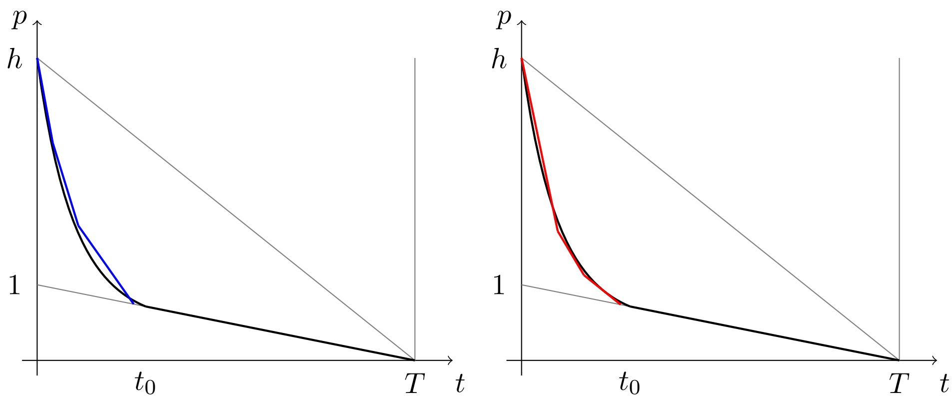

(a) PL Optimization: $n = 4$

(b) MVA: $n = 4$

Figure 4.3: The Upper Bound and The optimization

is also interesting to look at the competitive ratios. The best possible value is $c = 0.694$, Corollary 5.2. *PL* optimization reaches an optimistic value of 0.726. While the *MVA* outputs an upper bound $\alpha = 0.718$ and a lower bound $\beta = 0.647$. In this example it is evident that even if the *PL* optimization shows the higher competitive ratio, this is actually only an optimistic estimate. On the other hand, notice the quite good estimate given by the *MVA*. Estimate that should become much better in the next step $n = 5$, reaching a tight estimate of $\alpha = 0.707$ and $\beta = 0.686$.

4.4 Possible Extensions

This is a brief section that tries to answer the question *Why a piecewise linear pricing strategy?*. The possible answers are two. First, because it is simple and efficient. A piecewise linear pricing curve is an extremely manageable tool. The key point is that it is completely controllable through the intersection time points t_i. Second, because it can be modified and spread. Basically, one can follow two possible directions. The first one is adding variance, *i.e.* consider a bigger number of variables - with a constant number of discrete undiscounted valuations - in order to get a smoother pricing curve and have more degrees of freedom to modify it. The second one is trying to maintain the complexity of the problem, hence the number of variables and constraints, but changing the curve. We briefly show two non-linear mathematical programming examples, one for each direction.

Example 1 (SPLINE Optimization)**.** *This is an example of the first direction. We simply re-define the optimization problem adding a variable for* *each piece of curve. The segment of equation (4.1) in this case becomes:* $s_i(t) = A_i \cdot t^2 + B_i \cdot t + C_i \quad t_i \leq t < t_{i-1}$ *and the problem:*

$$\begin{aligned} \textbf{\textit{max}} \quad & \alpha \\ \textbf{\textit{subject to:}} \quad & c(\tilde{v}_i, t_i, B_i) \geq \alpha \quad \forall i = 0, n-1 \\ & t_i - t_{i-1} \leq 0 \qquad \forall i = 1, n-1 \\ & t_0 \leq D \end{aligned}$$

Where $c(\tilde{v}_i, t_i, B_i) = \frac{\mathbb{E}[R(\tilde{v}_i, t_i, B_i)]}{\mathrm{opt}(\tilde{v}_i)}$. *The Expected Revenue* $\mathbb{E}\left[R\left(\tilde{v}_i, t_i, B_i\right)\right]$ *is:*

$$e^{\lambda t_i} \left\{ \sum_{j=0}^{i-1} \left[\left(-A_{i-j} \frac{\lambda t_{i-j-1}^2 + 2\lambda t_{i-j-1} + 2}{\lambda^2} - B_{i-j} \frac{\lambda t_{i-j-1} + 1}{\lambda} - C_{i-j} \right) e^{-\lambda t_{i-j-1}} + \right.\right.$$
$$\left. + \left(A_{i-j} \frac{\lambda t_{i-j-1}^2 + 2\lambda t_{i-j-1} + 2}{\lambda^2} + B_{i-j} \frac{\lambda t_{i-j-1} + 1}{\lambda} + C_{i-j} \right) e^{-\lambda t_{i-j}} \right]$$
$$\left. + e^{-\lambda D} \left(\frac{D}{T} - 1 + \frac{1}{\lambda T} \right) + e^{-\lambda t_0} \left(1 - \frac{t_0}{T} - \frac{1}{\lambda T} \right) \right\}$$

While A_i *and* C_i *are easily derivable by the equations* $s_i(t_i) = \tilde{v}_i \left(1 - \frac{t_i}{T}\right)$ *and* $s_i(t_{i-1}) = \tilde{v}_i \left(1 - \frac{t_{i-1}}{T}\right)$. *Finally notice that also the* MVA *can be adapted to this framework.*

Example 2 (Polynomial Optimization)**.** *This is an example of the second direction. We define a new kind of pricing curve, depending on a set of variables that we have to optimize. We show here a simple polynomial pricing curve of order* m*:*

$$p(t) = \sum_{j=0}^{m} a_j t^j \quad t \in [0, t_0]$$

Where for $i = 0, \cdots, m$ a_j *are the variables to be optimized. We fix two of them, constraining the curve to assume values* h *in* $t = 0$ *and* $1 - \frac{t_0}{T}$ *and* $t = t_0$.

In the PL *optimization the variables defined by the mathematical programming are the time instants* t_i *and the discrete undiscounted valuations are the given parameters of the problem - as in Example 1. Here we switch the two roles. We give a vector of fixed time instants and the optimization procedure defines the price for each time instant, hence the corresponding undiscounted valuation.*

We can write the expected revenue of the mechanism as follows:

$$\mathbb{E}\left[R\left(t_i,a_j\right)\right]=e^{\lambda t_i}\int_{t_i}^{t_0}p\left(t+t_i\right)\lambda e^{-\lambda t}dt+e^{\lambda t_i}\int_{t_0}^{D}\left(1-\frac{t}{T}\right)\lambda e^{-\lambda t}dt=$$

$$-\sum_{j=1}^{m}a_j t_i^j\left(1-e^{-\lambda(t_0-t_i)}\right)-\sum_{j=0}^{m}a_j\frac{\sum_{k=0}^{j}\frac{j!}{k!}\left(\lambda t_0\right)^k}{\lambda^j}e^{-\lambda(t_0-t_i)}+$$

$$+\sum_{j=0}^{m}a_j\frac{\sum_{k=0}^{j}\frac{j!}{k!}\left(\lambda t_i\right)^k}{\lambda^j}+e^{\lambda t_i}\left[e^{-\lambda D}\left(\frac{D}{T}-1+\frac{1}{\lambda T}\right)+e^{-\lambda t_0}\left(1-\frac{t_0}{T}-\frac{1}{\lambda T}\right)\right]$$

Expressing the competitive ratio as $c(t_i,a_j)=\frac{\mathbb{E}[R(t_i,a_j)]}{\text{opt}(p(t_i))}$*, we can define the Polynomial optimization problem as follows:*

$$\begin{aligned} &\textbf{\textit{max}} \quad \alpha \\ \textbf{\textit{subject to:}} \quad & c(t_i,a_j)\geq\alpha \quad \forall i=0,n-1 \\ & p(0)=h \\ & p(t_0)=1-\tfrac{t_0}{T} \end{aligned}$$

4.5 Constrained Optimization

Very often, in a realistic economic scenario, a seller cannot modify the posted-price at every time instant. She must set a constant price at least for a given time interval δ. Consequently, she cannot use a continuous time pricing curve. Hence, a new model is needed. We study now a mathematical programming problem, similar to the *PL* optimization.

We have, again, a vector of undiscounted valuation $\mathbf{v}=[\tilde{v}_0,\ldots,\tilde{v}_{n-1}]$ composed of n discrete values *s.t.* $\tilde{v}_0=1$ and $\tilde{v}_{n-1}=h$. The goal is to define a constant price for each time slot of duration δ. After enumerating the $\frac{T}{\delta}$ slots, the price for each slot is p_i, $i=1,\cdots,\frac{T}{\delta}$.

We define two kinds of binary variables $x_{i,k}$ for each pair of price p_i and undiscounted valuation $\tilde{v}_k$ and y_i for each price p_i. $x_{i,k}$ is equal to 1 if the probability point mass distribution F is on $\tilde{v}_k$ and the item can be sold at the price p_i. Formally:

$$x_{i,k}=\begin{cases}1 & \text{if } \tilde{v}_k\left(1-\frac{i\delta}{T}\right)\geq p_i\\ 0 & \text{otherwise}\end{cases}$$

This formulation is traduced in the following constraints:

$$x_{i,k}\left(\tilde{v}_k\frac{T-i\delta}{T}-p_i\right)+(1-x_{i,k})\left(p_i-\tilde{v}_k\frac{T-i\delta}{T}\right)\geq 0 \quad \forall i,k$$

The binary variable y_i is 1 if the price the $i-th$ slot is before the time instant t_0. Where also in this case, $[t_0, T]$ is the time interval at which the minimum price is posted. Notice that t_0 is not a given parameter, but a variable to be optimized. Hence the y_i is defined as follows:

$$y_i = \begin{cases} 1 & \text{if } p_i \geq 1 - \frac{(i-1)\delta}{T} \\ 0 & \text{otherwise} \end{cases}$$

Traduced in the constraints:

$$y_i \left(p_i - 1 + \frac{(i-1)\delta}{T} \right) + (1 - y_i) \left(1 - p_i - \frac{(i-1)\delta}{T} \right) \geq 0 \quad \forall i$$

If the meaning of the first kind of binary variable $x_{i,k}$ is evident, the reason behind the binary variable y_i is not yet clear. Let us restrict our research to a special class of posted-price mechanisms. For each valuation $\tilde{v}(1 - \frac{t}{T})$ we consider only that mechanisms where $\exists$ a time slot $\bar{i}$ *s.t.* $p_i \geq \tilde{v}(1 - \frac{t}{T})$, $\forall i < \bar{i}$ and $p_i \leq \tilde{v}(1 - \frac{t}{T})$, $\forall i > \bar{i}$. This assumption is the same made in the chapter 3. Let us formulate the corresponding constraints. Given a price p_i, consider the corresponding undiscounted valuation $\tilde{v}$ *s.t.* $\tilde{v}(1 - \frac{(i-1)\delta}{T}) = p_i$, we say that the next price p_{i+1} has to be:

$$p_{i+1} \leq \tilde{v} \left(1 - \frac{(i+1)\delta}{T} \right) = p_i \frac{T - (i+1)\delta}{T - (i-1)\delta} \quad \forall i$$

This constraints hold in the range $[0, t_0]$. While from t_0 to T, the mechanism posts the minimum price, hence $p_i = 1 - \frac{i\delta}{T}$. And that is the meaning of y_i.

Once defined this class of mechanism, it is evident that if a certain undiscounted valuation $\tilde{v}_k$ is above a slot price p_i, then $\tilde{v}_k$ is above any other following slot prices p_j, $\forall j > i$. The same for y_i.

Hence, the *Constrained Optimization* problem is:

$$\begin{array}{lcr} & \textbf{max} \quad \alpha & \\ \textbf{s. t.:} & c(\tilde{v}_k) \geq \alpha & \forall i \\ & x_{i,k} \left(\tilde{v}_k \frac{T-i\delta}{T} - p_i \right) + (1 - x_{i,k}) \left(p_i - \tilde{v}_k \frac{T-i\delta}{T} \right) \geq 0 & \forall i, k \\ & y_i \left(p_i - 1 + \frac{(i-1)\delta}{T} \right) + (1 - y_i) \left(1 - p_i - \frac{(i-1)\delta}{T} \right) \geq 0 & \forall i \\ & y_i \left(p_{i+1} - p_i \frac{T-(i+1)\delta}{T-(i-1)\delta} \right) + (1 - y_i) \left(p_{i+1} - 1 + \frac{(i+1)\delta}{T} \right) \leq 0 & \forall i \\ & x_{i,k} \leq x_{i+1,k} & \forall i, k \\ & y_i \geq y_{i+1} & \forall i \end{array}$$

Where the competitive ratio is $c(\tilde{v}_k) = \frac{\mathbb{E}[R(\tilde{v}_k)]}{\mathrm{opt}(\tilde{v}_k)}$ and the Expected Revenue:

$$\mathbb{E}\left[R\left(v_k\right)\right] = \left(1 - e^{-\lambda\delta}\right) \sum_{i=1}^{T/\delta} x_{i,k} p_i e^{-\lambda\delta\Sigma_{j=1}^{i-1} x_{j,k}}$$

Example 3 (Constrained Optimization: a Room to Rent)**.** *We refer to the motivating example 1.1.1. Suppose to be at the end of the year. A seller wants to rent a room within the end of the next year. She has some knowledge of the market, for example she knows the range in which the agents' valuations lie. hence, she can evaluate the ratio between the maximum and minimum valuation. There exists a buyer who is willing to rent the same room at most two times the minimum buyers' valuation and there is no other agents with an higher valuations, hence* $h = 2$*. From the past years she also knows the expected number of agents coming in the market. Unfortunately, she can modify the posted price only twice a week. She has almost* 50 *weeks to rent the room in a year; hence, she has available* 100 *time slots. Let assume that in expectation* 15 *customers will be interested in renting the room over the year. Hence* $\lambda T = 15$*. Conventionally, suppose* $\lambda = 3$ *and* $T = 5$*.*
Solving the optimization problem (4.4) with BARON, we're able to reach a competitve ratio of 0.698*. Note that the optimal mechanism* $\mathcal{M}_1$ *in this scenario is able to reach* 0.784 *as constant competitive ratio.*

Chapter 5

Random Valuation

In the *Random Valuation* scenario, the undiscounted valuations $\tilde{v}$ are drawn from an unknown cumulative distribution function F. In this chapter, we assume F belonging to the family $\mathcal{F}$ of *MHR - Monotone Hazard Rate* - distributions. The hazard rate q of a CDF $F(x)$ is $q(x) = \frac{f(x)}{1-F(x)}$, defined for $F(x) \leq 1$. All the probability distributions, whose hazard rate is monotone non-decreasing in x, belong to the MHR family $\mathcal{F}$. The MHR assumption is common in economics and auction theory and this family includes a wide range of distributions.
In actuarial science the hazard rate is also called "force of mortality", in engineering science "failure rate", in economics its reciprocal is known as "Mills' ratio". The concept of virtual valuation suggests a very interesting interpretation of the hazard rate. We can think the undiscounted valuation of an agent $\tilde{v}_i$ as the maximum revenue obtainable from agent i. The reciprocal of the hazard rate can be interpretable as the inevitable revenue loss caused by not knowing $\tilde{v}_i$ in advance, the "information rent". The virtual valuation is the difference between $\tilde{v}_i$ and the Mills' ratio. This is the fundamental trade-off that a seller who doesn't know the buyer's willingness to pay must make. See (Roughgarden, 2013) for example.
In this chapter we consider the *linear discounting rate* case.

5.1 The Model

The *Random Valuation* scenario inherits the *Identical Valuation* model (3.1). We define random variable $\left(1 - \frac{t_i}{T}\right)$ as the time discount associated to agent i. It takes value from the discounting rate function. Where $t_i \sim \Gamma(i, \lambda)$ is gamma-distributed and represents the arriving time of the i^{th} agent in the Poisson process. A good question is whether the two random variables

$\left(1 - \frac{t_i}{T}\right)$ and $\tilde{v}_i$ are correlated. Do the arriving time of an agent influence the intrinsic willingness to pay of that agent? The arriving time of an agent is roughly the time at which the agent starts to need the item. Hence, there is no correlation. One can say that an agent with a higher undiscounted valuation is incentive to participate to the trade earlier with respect to an agent with a smaller valuation. But the no-correlation assumption makes the model an under estimation of the latter. Indeed, if high valuations show early, the time discounts for that valuations are close to one. Consequently there should be higher willingness to pay in the market and a seller should gain more. Hence, we assume the random variable $\tilde{v}_i$ and $\left(1 - \frac{t_i}{T}\right)$ to be independent.

We define random variable $X_{\lambda T}$ as the maximum undiscounted valuation of agents arriving in a T-lengthed time interval.

$$X_{\lambda T} = \max_{i=1:N(T)} \tilde{v}_i \tag{5.1}$$

Where random variable $N(T)$ represents the total number of agents arriving in $[0, T]$ according to the Poisson process. Hence, $X_{\lambda T}$ is the first order statistic of $N(T)$ samples from F. We call its CDF $F_{X_{\lambda T}}$.

We define $Y_{\lambda T}$ as the random variable of the maximum valuation of agents arriving in a T-lengthed interval. We call its CDF $F_{Y_{\lambda T}}$.

$$Y_{\lambda T} = \max_{i=1:N(T)} v_i = \max_{i=1:N(T)} \tilde{v}_i u_i \tag{5.2}$$

5.2 The Benchmark

This section is dedicated to the *onmiscient seller* that we use as the benchmark in our competitive analysis. The *benchmark* is the *clairvoyant offline* mechanism who knows all the input data before time 0, hence the realizations of the agents' valuations. This mechanism sells the item to the highest valuations agent, gaining exactly her willingness-to-pay. We call this mechanism *Y benchmark*. Thus, her expected revenue - let us call it *Ybench* - is equal to the expected value of the maximum valuation arriving until the deadline.

$$Ybench = \mathbb{E}\left[Y_{\lambda T}\right] \tag{5.3}$$

In this study, it is useful to refer to another ideal benchmark whose expected revenue is even higher than the previous one, let us call it *X benchmark*. Such a mechanism knows the valuations of the agents before they arrive in the market, but she is also able to sell the item in an offline manner, *i.e.* at time 0. This means that the discount associated is always 1. Thus, the expected revenue of X benchmark is equal to the expected value of the maximum undiscounted valuation arriving until the deadline.

$$Xbench = \mathbb{E}\left[X_{\lambda T}\right] \tag{5.4}$$

This section is dedicated to the computation of the cumulative density function $F_{Y_{\lambda T}}$ and $F_{Y_{\lambda T}}$, exposing some of their properties. First of all, we present a common result in a Poisson process, see for example (Ross et al., 1996).

Lemma 10. *(Poisson process) Given that $N(T) = n$, the n arrival times $t_1, \dots, t_n$ have the same distribution as the order statistics corresponding to n independent random variables uniformly distributed on the interval $(0, T)$.*

Observation 3. *Intuitively, we usually say that under the condition that n events have occurred in $(0, T)$, the times $t_1, \dots, t_n$ at which events occur, considered as unordered random variables, are distributed independently and uniformly in the interval $(0, T)$.*

We call $s_1, \dots, s_n$ the unordered arrival times of the agents conditioning to the event that n agents arrive in $(0, T)$. Their order statistics are $t_1, \dots, t_n$. We can conclude that s_i is uniformly distributed in $(0, T)$. Hence, we define u_i as the unordered time discount associated to a general agent i. Let us refer to the CDF and PDF of u_i as G_u and g_u, respectively. They have a very simple shape:

$$\begin{aligned} G_u(x) &= \mathbb{P}\left(u_i \leq x\right) = \mathbb{P}\left(1 - \frac{s_i}{T} \leq x\right) = 1 - \mathbb{P}\left(s_i \leq T(1-x)\right) = \\ &= 1 - \begin{cases} 1 & \text{if } x < 0 \\ \frac{T(1-x)}{T} & \text{if } x \in [0,1] \\ 0 & \text{if } x > 0 \end{cases} = \begin{cases} 0 & \text{if } x < 0 \\ x & \text{if } x \in [0,1] \\ 1 & \text{if } x > 0 \end{cases} \end{aligned}$$

$$g_u(x) = \mathbb{I}_{[0,1]}\{x\}$$

Where $\mathbb{I}_A$ is the indicator function defined on the set A.
Looking at G_u could be surprising notice that there is no dependency on T. Only the shape of the *discounting rate* affects the G_u. This seems curious. The probability distribution over the time discounts loses the information about how long is the time horizon. This brings to the following question. *Does not the time horizon affect the maximum value of the discounted valuation?*. We will answer this question later.

Hence, the random variable representing the private valuation v_i of agent i is expressible as the undiscounted valuation $\tilde{v}_i$ discounted by u_i, *i.e.* $v_i =$

$\tilde{v}_i u_i$. We stress the fact that agent i is not the i^{th} agent arrived. We are interested in computing the conditional cumulative distribution function $F_{V|N}(x,k) = \mathbb{P}\left(\tilde{v}_i\left(1-\frac{t_i}{T}\right) \le x | N = K\right)$, $\forall\ k \ge i$.

$$\begin{aligned}
F_{V|N}(x,k) &= \mathbb{P}\left(\tilde{v}\left(1-\frac{t_i}{T}\right) \le x | N = k\right) = \mathbb{P}(\tilde{v}u \le x) = \mathbb{P}\left(\tilde{v} \le \frac{x}{u}\right) = \\
&= \mathbb{I}_{[0,1)}\{x\} \iint_{\mathcal{D}'} f(\tilde{v})g(u)d\tilde{v}du + \mathbb{I}_{[1,h]}\{x\} \iint_{\mathcal{D}''} f(\tilde{v})g(u)d\tilde{v}du = \\
&= \mathbb{I}_{[0,1)}\{x\} \int_1^h f(\tilde{v}) \int_0^{x/\tilde{v}} g(u)dud\tilde{v} + \\
&\quad +\mathbb{I}_{[1,h]}\{x\} \left(\int_0^1 g(u)du \int_1^x f(\tilde{v})d\tilde{v} + \int_x^h f(\tilde{v}) \int_0^{x/\tilde{v}} g(u)dud\tilde{v}\right) \\
&= \mathbb{I}_{[0,1)}\{x\} \cdot x \int_1^h \frac{1}{\tilde{v}} f(\tilde{v})d\tilde{v} + \mathbb{I}_{[1,h]}\{x\} \left(F(x) + x\int_x^h \frac{1}{\tilde{v}} f(\tilde{v})d\tilde{v}\right)
\end{aligned}$$

The domains $\mathcal{D}'$ and $\mathcal{D}''$ are defined as:

$$\mathcal{D}' = \left\{0 \le u \le \frac{x}{\tilde{v}}, 1 \le \tilde{v} \le h\right\}$$

$$\mathcal{D}'' = \mathcal{D}_1'' \cup \mathcal{D}_2'' = \{0 \le u \le 1, 1 \le \tilde{v} \le x\} \cup \left\{0 \le u \le \frac{x}{\tilde{v}}, x < \tilde{v} \le h\right\}$$

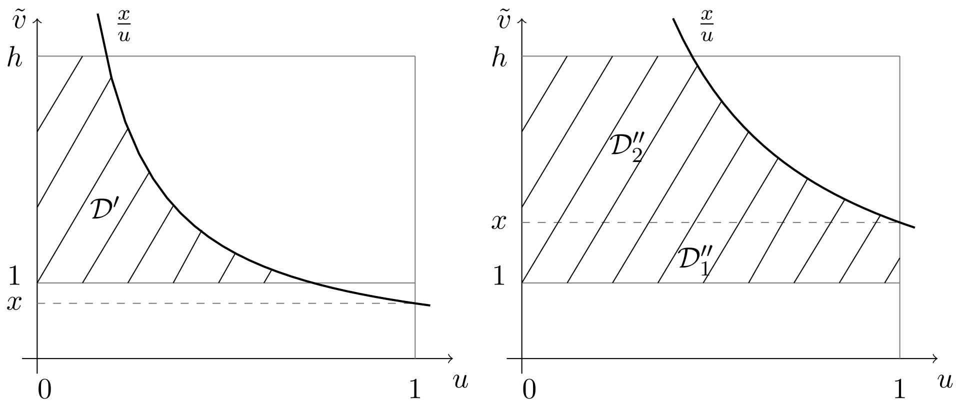

(a) $\mathcal{D}'$ for $x \in [0,1)$

(b) $\mathcal{D}''$ for $x \in [1,h]$

Figure 5.1: Domains of $F_{V|N}$

Hence, the conditional CDF $F_{V|N}(x,k)$ and PDF $f_{V|N}(x,k)$ are:

$$F_{V|N}(x,k) = \begin{cases} x \cdot \int_1^h \frac{1}{\tilde{v}} f(\tilde{v})d\tilde{v} & \text{if } x \in [0,1) \\ F(x) + x\int_x^h \frac{1}{\tilde{v}} f(\tilde{v})d\tilde{v} & \text{if } x \in [1,h] \end{cases} \tag{5.5}$$

$$f_{V|N}(x,k) = \begin{cases} \int_1^h \frac{1}{\tilde{v}} f(\tilde{v}) d\tilde{v} & \text{if } x \in [0,1) \\ \int_x^h \frac{1}{\tilde{v}} f(\tilde{v}) d\tilde{v} & \text{if } x \in [1,h] \end{cases} \tag{5.6}$$

Notice that for $x \in [0,1)$ the conditional distribution function has a typical uniform distribution behaviour. Indeed, $\int_1^h \frac{1}{\tilde{v}} f(\tilde{v}) d\tilde{v}$ is a constant term. It is worth mentioning that $F_{V|N}(1^-,k) = F_{V|N}(1^+,k)$. Notice also that for $x \in [1,h]$, $F_{V|N}(x,k) \geq F(x)$. This is correct, the discounted valuation of an agent is for sure smaller than the undiscounted one.
We want to underline the fact that these formulas are constant $\forall i$ and $\forall k \geq i$.
Let us discuss now a useful property.

Property 1. *The function $g(x_1,\ldots,x_n)$ is said to be symmetric if $g(x_1,\ldots,x_n) = g(x_{i_1},\ldots,x_{i_n})$ for any permutation of $(i_1,\ldots,i_n)$.*

Lemma 11. *The function $\mathbb{P}\left(\max_{i\in\{1,\ldots,j\}}\left\{\tilde{v}_i\big(1-\frac{t_i}{T}\big)\right\} \leq x\right)$ is symmetric, for any $j \geq 1$.*

Proof (Lemma 11). For $j = 1$ the proof is obvious.
Fo $j > 1$, notice that the random variables $\tilde{v}_i$ are *i.i.d.*. Hence, the maximum does not change with respect to the order of the random variables $\tilde{v}_i\big(1-\frac{t_i}{T}\big)$. Indeed:

$$\max_{i\in\{1,\ldots,j\}}\left\{\tilde{v}_i\big(1-\frac{t_i}{T}\big)\right\} = \max_{i\in\{i_1,\ldots,i_n\}}\left\{\tilde{v}_i\big(1-\frac{t_i}{T}\big)\right\}$$

for any permutation of $(i_1,\ldots,i_n)$. The thesis of the theorem directly follows. $\square$

Finally, we can compute the CDF of $Y_{\lambda T}$.

Theorem 12. *$Y_{\lambda T}$ has CDF:*

$$F_{Y_{\lambda T}}(x) = \sum_{j=0}^{\infty} \frac{(\lambda T)^j e^{-\lambda T}}{j!} F^j_{V|N}(x,j) \tag{5.7}$$

Proof (Lemma 12). For lemma 10 and lemma 11, it is possible to write the following, for a generic j:

$$\mathbb{P}\left(\max_{i\in\{1,\ldots,j\}}\left\{\tilde{v}_i\big(1-\frac{t_i}{T}\big)\right\} \leq x|\mathbb{P}(N(T)=j)\right) = \mathbb{P}\left(\max_{i\in\{1,\ldots,j\}}\left\{\tilde{v}_i\big(1-\frac{s_i}{T}\big)\right\} \leq x\right)$$

$$= \mathbb{P}\left(\max_{i\in\{1,\ldots,j\}}\{\tilde{v}_i u_i\} \leq x\right) = \mathbb{P}\left(\bigcap_{i=1}^{j}(v_i \leq x)\right) = \prod_{i=1}^{j}\mathbb{P}\left(v_i \leq x\right) = F^j_{V|N}(x,j)$$

Finally, being $Y_{\lambda T}$ the maximum valuations of agents arriving in a T-lengthed interval based on a Poisson process, its CDF is:

$$\begin{aligned} F_{Y_{\lambda T}}(x) &= \sum_{j=0}^{\infty} \mathbb{P}(N(T)=j)\mathbb{P}\left(\max_{i\in\{1,\dots,j\}}\left\{\tilde{v}_i\big(1-\frac{t_i}{T}\big)\right\} \leq x|\mathbb{P}(N(T)=j)\right) = \\ &= \sum_{j=0}^{\infty} \frac{(\lambda T)^j e^{-\lambda T}}{j!} F^j_{V|N}(x,j) \end{aligned}$$

This concludes the proof. □

We can now answer the open question: *Does not the time horizon affect the maximum value of the discounted valuation?*. Clearly it does! The magnitude of T affects neither G_u nor $F_{V|N}$. Nevertheless, it affects the weights of the summation. Given a certain number of agents, the maximum valuation of that agents has the same distribution independently by T. But this probability distribution is only a weighted term of the overall distribution. Such a weight depends on T.
It is quite interesting make a final remark.

Observation 4. *If the expected number of agents λT remains constant, the cumulative probability distribution over the maximum valuation $F_{Y_{\lambda T}}$ is constant, no matter the magnitude of T.*

Let us compare the (5.7) with the CDF of $X_{\lambda T}$:

$$F_{X_{\lambda T}}(x) = \sum_{j=0}^{\infty} \frac{(\lambda T)^j e^{-\lambda T}}{j!} F^j(x) \tag{5.8}$$

We derive a results about $F_{X_{\lambda T}}$. The random variable of the maximum in a τ-lengthed interval over the undiscounted valuations preserves the *Monotone Hazard Rate* property.

Lemma 13. *$F_{X_{\lambda\tau}}$ has non-decreasing monotone hazard rate.*

Proof (Lemma 13). The cumulative density function of $X_{\lambda\tau}$ is:

$$F_{X_{\lambda\tau}}(x) = \sum_{i=0}^{\infty} \frac{(\lambda\tau)^i e^{-\lambda\tau}}{i!}(F(x))^i = e^{-\lambda\tau}e^{\lambda\tau F(x)} = e^{-\lambda\tau(1-F(x))}$$

We compute the hazard rate of $F_{X_{\lambda\tau}}(x)$ and show it is non-decreasing.

$$h_{X_{\lambda\tau}}(x) = \frac{f_{X_{\lambda\tau}}(x)}{1 - F_{X_{\lambda\tau}}(x)} = \frac{\frac{d}{dx}F_{X_{\lambda\tau}}(x)}{1 - F_{X_{\lambda\tau}}(x)} = \frac{\lambda\tau f_{X_{\lambda\tau}}(x)e^{-\lambda\tau(1-F(x))}}{1 - e^{-\lambda\tau(1-F(x))}} \tag{5.9}$$

$$= \frac{\lambda\tau f(x)}{e^{\lambda\tau(1-F(x))} - 1} = \lambda\tau \frac{f(x)}{1 - F(x)} \frac{1 - F(x)}{e^{\lambda\tau(1-F(x))} - 1} \tag{5.10}$$

$$= \lambda\tau h(x) \frac{1 - F(x)}{e^{\lambda\tau(1-F(x))} - 1} \tag{5.11}$$

F is MHR, thus its hazard rate $h(x)$ is non-decreasing. Note that $F(x)$ is non-decreasing, so $y = 1 - F(x)$ in non-increasing. Proving that $\frac{1-F(x)}{e^{\lambda\tau(1-F(x))}-1}$ is non-decreasing in x is equivalent to show that $g(y) = \frac{y}{e^{\lambda\tau y}-1}$ is non-increasing in y. We study the first derivative of $g(y)$:

$$\frac{d}{dy}g(y) = \frac{e^{\lambda\tau y}(1 - \lambda\tau y) - 1}{(e^{\lambda\tau y} - 1)^2} < 0 \quad \forall y$$

This implies that $g(y)$ is non-increasing in y, hence $\frac{1-F(x)}{e^{\lambda\tau(1-F(x))}-1}$ is non-decreasing in x.
We conclude that $h_{X_{\lambda\tau}}(x)$ is monotone non-decreasing. $\square$

5.3 A Lower Bound for $\mathcal{M}_1$

$\mathcal{M}_1$ is the deterministic posted-price mechanism with the highest competitive ratio in the identical valuation setting:

$$p(t) = \begin{cases} h \cdot \left(1 - \frac{t}{T}\right) \cdot e^{\lambda\left(1-\frac{1}{k}\right)t + \frac{\lambda}{2kT}t^2} & t \in [0, t_0) \\ 1 - \frac{t}{T} & t \in [t_0, T] \end{cases}$$

where

$$k = 1 - \frac{1}{\lambda T}\left(1 + \lambda t_0 - e^{-\lambda(T-t_0)}\right)$$

and t_0 is such that

$$\lambda t_0 \frac{2T - t_0}{2T\left(\lambda t_0 + \ln(h)\right)} = 1 - \frac{1}{\lambda T}\left(1 + \lambda t_0 - e^{-\lambda(T-t_0)}\right)$$

In this section we study the performances of mechanism $\mathcal{M}_1$ in the random valuation setting. Although $\mathcal{M}_1$ was developed to be the optimal mechanism in the identical valuation setting, we show that it can achieve good performances also in this other scenario. In particular, Theorem 16 provides a lower bound on the competitive ratio of $\mathcal{M}_1$ that is computed with respect to a *clairvoyant* mechanism which knows the valuations of the agents in advance. We show that our bound is valid for every non-decreasing monotone hazard rate distribution F.

In order to derive a constant lower bound for the mechanism $\mathcal{M}_1$, we follow the basic idea suggested by (Babaioff et al., 2011) in their *Equal Sample for Every Scale* mechanism. In their setting, the number of agents that sooner or later arrive in the market is known. Moreover there is no deadline and no discounting rate of the valuations. This simply means there is no interest in *when* the item is sold, but only that it will be sold in future. Our challenge is much more complex for at leat three reasons. 1) we don't know how many agents will arrive in the market, we have only an estimate of the expected value of this number. 2) there is a deadline, we have to sell the item within a given time horizon, otherwise we gain 0 revenue. 3) the valuations are *time-variant*, hence time plays a very important role in our setting.

Note that the pricing strategy of $\mathcal{M}_1$ is monotone non-increasing. Image a τ-lengthed sliding window moves on the time horizon $[0, T]$, with $\tau \leq T$. Instead of comparing the performance of our mechanism in the time horizon $[0, T]$ with respect to the benchmark, we consider the performance of $\mathcal{M}_1$ only in an optimizing positioned sliding window and we compare it to the overall offline benchmark. It is clear that the expected performance of the mechanism in the overall time horizon is at least its expected performance only in the sliding window.
It will be useful computing the expected revenue of the benchmarks in the sliding window. Note that the expected revenue of Y *benchmark* depends both on the duration τ and on the initial time instant of the sliding window t: $\mathbb{E}\left[Y_{\lambda\tau,t}\right]$. The expected revenue of X *benchmark* depends only on τ. We use now a result by (Zheng et al., 2016).

Lemma 14. *(Zheng et al., 2016)*

$$\frac{\mathbb{E}[X_{\lambda\tau}]}{\mathbb{E}[X_{\lambda\tau'}]} \geq \frac{\ln \lambda\tau}{\ln \lambda\tau'}, \quad \forall \tau \leq \tau'$$

The proof in the appendix.

Let us call $\frac{\ln \lambda\tau}{\ln \lambda T} = (1 - \varepsilon)$. This implies the following expressions:

$$\mathbb{E}[X_{\lambda\tau}] \geq (1 - \varepsilon)\, \mathbb{E}[X_{\lambda T}] \tag{5.12}$$

$$\mathbb{E}[Y_{\lambda\tau,t]}] \geq \left(1 - \frac{t+\tau}{T}\right) \mathbb{E}[X_{\lambda\tau}] \geq \left(1 - \frac{t+\tau}{T}\right)(1 - \varepsilon)\, \mathbb{E}[X_{\lambda T}] \tag{5.13}$$

We give two definitions:

Definition 9. *Consider mechanism $\mathcal{M}_1$ and a time interval $I_{s,\tau} = [s, \tau + s] \subseteq [0, T]$ of length τ, where $s \in [0, T - \tau]$. The ratio between the price at the starting point and the price at the ending point is:*

$$\kappa(s) = \frac{p_s}{p_{\tau+s}} = \frac{T - s}{(T - \tau - s)e^{\lambda\left(1-\frac{1}{k}\right)\tau + \frac{\lambda}{2kT}(\tau^2+2s\tau)}}$$

Note that $\kappa(s)$ depends on the position of the interval on $[0, T]$.

Definition 10. *We denote the maximum ratio between the prices posted by $\mathcal{M}_1$ in a τ-lengthed interval by:*

$$\kappa = \max_{s\in[0,T-\tau]} \kappa(s)$$

Lemma 15. *Consider pricing mechanism $\mathcal{M}_1$. There exists at least a time interval $I_{s,\tau} = [s, \tau + s]$ such that each price p_t with $t \in I_{s,\tau}$ is in the range $\left[\frac{\mathbb{E}[X_{\lambda T}]\left(1-\frac{\tau+s}{T}\right)(1-\epsilon)}{\kappa}, \mathbb{E}[X_{\lambda T}]\left(1 - \frac{\tau+s}{T}\right)(1 - \epsilon)\right]$*

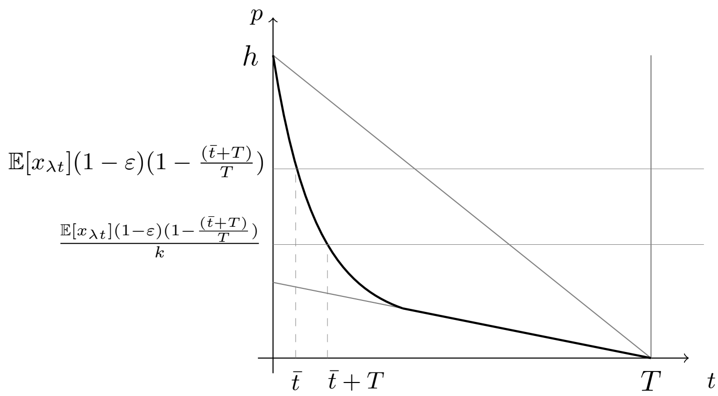

Figure 5.2: $\mathcal{M}_1$: Position of the sliding window

Proof (Lemma 15). We define the interval $I_{\bar{s},\tau}$ choosing $\bar{s}$ such that $p_{\bar{s}} = \mathbb{E}[X_{\lambda T}]\left(1 - \frac{\tau+\bar{s}}{T}\right)(1 - \epsilon)$. From Definition 10 we know that $\kappa(\bar{s}) \le \kappa$. Hence,

$$p_{\tau+\bar{s}} = \frac{p_{\bar{s}}}{\kappa(\bar{s})} = \frac{\mathbb{E}[X_{\lambda T}]\left(1 - \frac{\tau+\bar{s}}{T}\right)(1 - \epsilon)}{\kappa(\bar{s})} \ge \frac{\mathbb{E}[X_{\lambda T}]\left(1 - \frac{\tau+\bar{s}}{T}\right)(1 - \epsilon)}{\kappa}$$

The pricing strategy of mechanism $\mathcal{M}_1$ is monotone non-increasing, thus, $\forall t \in [\bar{s}, \tau + \bar{s}]$ we have:

$$p_t \in [p_{\bar{s}}, p_{\tau+\bar{s}}] \subseteq \left[\frac{\mathbb{E}[X_{\lambda T}]\left(1 - \frac{\tau+\bar{s}}{T}\right)(1 - \epsilon)}{\kappa}, \mathbb{E}[X_{\lambda T}]\left(1 - \frac{\tau + \bar{s}}{T}\right)(1 - \epsilon)\right]$$

The equalities are satisfied if $\bar{s} = \arg\max_{s\in[0,T-\tau]} \kappa(s)$. In this case there exists a unique interval satisfying Lemma 15. $\square$

Finally, we can derive the constant lower bound for the mechanism $\mathcal{M}_1$.

Theorem 16. *Let $(\lambda T)^\epsilon \geq \log_\kappa h$ for $\epsilon \in (0,1)$ and consider agent undiscounted valuations drawn i.i.d. from a distribution F with non-decreasing monotone hazard rate and linearly discounted by $\left(1-\frac{t}{T}\right)$. Mechanism $\mathcal{M}_1$ has a competitive ratio of $\frac{\left(1-(\lambda T)^{-\epsilon}-\frac{t_0}{T}\right)(1-\epsilon)}{\kappa e}$ that is constant with respect to the choice of F. $\bar{s}$ and κ are parameters defined by the problem instance, hence, depending on T, λ and h.*

Note that the monotone-hazard rate assumption is only required for F, the distribution of valuations. No such requirement is needed for the distribution of the random variable Z of the discounted valuations.

Proof (Theorem 16). Let $(\lambda T)^\epsilon \geq \log_\kappa h$. Since $1-\epsilon = \frac{\ln(\lambda\tau)}{\ln(\lambda T)}$, it follows $\lambda\tau = (\lambda T)^{1-\epsilon}$, for $\epsilon \in (0,1)$. From Lemma 15, mechanism $\mathcal{M}_1$ guarantees the existence of an interval $I_{s,\tau} \subseteq [0,T]$ such that $p_t \in \tilde{I} = \left[\frac{\mathbb{E}\left[X_{\lambda T}\right]\left(1-\frac{\tau+s}{T}\right)(1-\epsilon)}{\kappa}, \mathbb{E}[X_{\lambda T}]\left(1-\frac{\tau+s}{T}\right)(1-\epsilon)\right]$, $\forall t \in I_{s,\tau}$.

We have,

$$
\begin{aligned}
\mathbb{E}[\mathcal{R}(\mathcal{M}_1)] &= \mathbb{E}[\mathcal{R}(\mathcal{M}_1)]_{[0,T]} && (5.14)\\
&\geq \mathbb{E}[\mathcal{R}(\mathcal{M}_1)]_{[s,\tau+s]} && (5.15)\\
&\geq p_{\tau+s}\mathbb{P}(Y_{\lambda\tau,s} \geq p_\tau) && (5.16)\\
&\geq p_{\tau+s}\mathbb{P}\Big(X_{\lambda\tau}\Big(1-\frac{\tau+s}{T}\Big) \geq \mathbb{E}\left[X_{\lambda T}\right]\Big(1-\frac{\tau+s}{T}\Big)(1-\epsilon)\Big) && (5.17)\\
&= p_{\tau+s}\mathbb{P}(X_{\lambda\tau} \geq \mathbb{E}\left[X_{\lambda T}\right](1-\epsilon)) && (5.18)\\
&= p_{\tau+s}\mathbb{P}\left(X_{\lambda\tau} \geq \mathbb{E}\left[X_{\lambda T}\right]\frac{\ln(\lambda\tau)}{\ln(\lambda T)}\right) && (5.19)\\
&\geq p_{\tau+s}\mathbb{P}(X_{\lambda\tau} \geq \mathbb{E}\left[X_{\lambda\tau}\right]) && (5.20)\\
&\geq \frac{p_{\tau+s}}{e} \geq \frac{\mathbb{E}\left[X_{\lambda T}\right]\left(1-\frac{\tau+s}{T}\right)(1-\epsilon)}{\kappa e} && (5.21)\\
&\geq \frac{\mathbb{E}\left[X_{\lambda T}\right]\left(1-\frac{\tau+t_0}{T}\right)(1-\epsilon)}{\kappa e} && (5.22)
\end{aligned}
$$

where $\mathbb{E}[\mathcal{R}(\mathcal{M}_1)]_{[a,b]}$ is the expected revenue of the mechanism in an interval $[a,b]$. Inequality (5.17) holds because of Inequality 5.13. Inequality (5.20) follows from Lemma 14. A result by Barlow and Marshall (1964) implies that the probability of exceeding its expectation is at least $\frac{1}{e}$ for a non-decreasing monotone hazard rate distribution. We use this result in Inequality (5.21). Finally, note that the sliding window $I_{s,\tau}$ in the worst case has its starting point at $s = t_0$. Indeed, from t_0 on, the mechanism posts the minimum price. Hence, we have Inequality (5.22).

Now we compute a lower bound for the competitive ratio of the mechanism $\mathcal{M}_1$:

$$c_{\mathcal{M}_1} = \frac{\mathbb{E}[\mathcal{R}(\mathcal{M}_1)]}{\mathbb{E}[\mathcal{R}(\mathcal{M}^*)]} \geq \frac{\mathbb{E}[\mathcal{R}(\mathcal{M}_1)]}{\mathbb{E}[Y_{\lambda T}]} \tag{5.23}$$

$$= \frac{\mathbb{E}[X_{\lambda T}]}{\mathbb{E}[Y_{\lambda T}]} \frac{\left(1 - \frac{\tau + s}{T}\right)(1-\epsilon)}{\kappa e} \tag{5.24}$$

$$\geq \frac{\left(1 - \frac{\tau + t_0}{T}\right)(1-\epsilon)}{\kappa e} \tag{5.25}$$

It is easy to see that $\frac{\mathbb{E}[X_{\lambda T}]}{\mathbb{E}[Y_{\lambda T}]} \geq 1$ and $\lim_{\lambda T \to \infty} \frac{\mathbb{E}[X_{\lambda T}]}{\mathbb{E}[Y_{\lambda T}]} = 1$, hence we can write Inequality 5.25.

Recalling the condition $\lambda\tau = (\lambda T)^{1-\epsilon}$,

$$\tau = T^{1-\epsilon}\lambda^{-\epsilon}$$

We can write the bound as

$$c_{\mathcal{M}_1} = \frac{\left(1 - (\lambda T)^{-\epsilon} - \frac{t_0}{T}\right)(1-\epsilon)}{\kappa e}$$

where t_0, ε and κ are constants depending on the parameters of the problem, which are T, λ, h. □

5.4 A Lower Bound for $\mathcal{M}_2$

In this section we define a mechanism $\mathcal{M}_2$ where the offered price $p(t)$ is a *staircase function*. We evenly partition $[0, t_0]$ in $\lceil \log_\delta h \rceil$ τ-lengthed subintervals $I_i = [(i-1)\tau, i\tau]$, for $i \in \{1, \ldots, \lceil \log_\delta h \rceil\}$. Note that $\tau = \frac{t_0}{\lceil \log_\delta h \rceil}$. The constant price posted in interval I_i is denoted by $p(I_i)$ and δ is a parameter in the range $(1, h]$. Both t_0 and δ are parameters that can be optimized.
Mechanism $\mathcal{M}_2$ offers price $\frac{h}{\delta^i}(1 - \frac{i\tau}{T})$ to agents arriving in the time interval I_i, for $i \in \{1, \ldots, \lfloor \log_\delta h \rfloor\}$, $(1 - \frac{\lceil \log_\delta h \rceil \tau}{T})$ to agents arriving in the time

interval $I_{\lceil \log_\delta h \rceil}$ and $(1 - \frac{i\tau}{T})$ for agents arriving in the time interval I_i for $i = \lfloor \log_\delta h \rfloor, \dots, \rfloor \frac{T}{\tau} \lfloor$. Note that if h and δ are *s.t.* $\lceil \log_\delta h \rceil = \lfloor \log_\delta h \rfloor$, then $\frac{h}{\delta^{\lceil \log_\delta h \rceil}} = 1$ and the two intervals are actually the same.

Note that $\mathcal{M}_2$ for $[\lfloor \frac{T}{\tau} \rfloor, T]$ is not defined. This is not an oversight. Actually, the lower bound built for this mechanism does not consider the interval $[t_0, T]$. The reason is quite simple, a seller does not want to sell the item too late, otherwise the associated discount is too high. It is quite evident that, in a real scenario, the pricing strategy is defined also in such an interval. This clearly increases the revenue of the mechanism, because it can be possible that no agents arrive within t_0. Nevertheless, for the purposes of the computation of this bound it is not meaningful.

There are at least two reasons for posting a staircase function. *First*, as already mentioned in the section (4.5), a constant price in a sub-interval meets a specific need of some sellers which cannot change the price continuously in time. *Second*, this is an extension of the *Equal Sample for Every Scale* by (Babaioff et al., 2011), in which constant prices are offered to a sub-set of agents.

We now prove a similar result to that one of 15 for mechanism $\mathcal{M}_1$.

Lemma 17. *Consider pricing mechanism $\mathcal{M}_2$. There exists a price p_i^* such that:*

$$p_i^* \in \Big[\frac{\mathbb{E}\left[X_{\lambda T}\right]\big(1 - \frac{i\tau}{T}\big)(1-\epsilon)}{\delta}, \mathbb{E}\left[X_{\lambda T}\right]\big(1 - \frac{i\tau}{T}\big)(1-\epsilon)\Big]$$

for some $i \in \{1, \dots, \lceil \log_\delta h \rceil\}$.

Proof (Lemma 17). We prove by contradiction the existence of such a price p_i^*. We start by supposing that $\nexists\, i$ such that $p_i^* \in \tilde{I}_i$.
Note that ν is the lower bound for the maximum undiscounted valuation in a τ-lengthed interval. Hence, $\nu \in [1, h]$. Considering Lemma 14 this lower bound is expressible with respect to the maximum undiscounted valuation in a T- lengthed interval as $\mathbb{E}[X_{\lambda T}](1-\epsilon)$. This means that either $\nu \in [\frac{h}{\delta}, h]$ or $\nu \in [1, \frac{h}{\delta})$. If $\nu \in [\frac{h}{\delta}, h]$, then $\exists p_i^* = p_1$ such that $p_1 \in [\frac{\nu}{\delta}\big(1 - \frac{\tau}{T}\big), \nu\big(1 - \frac{\tau}{T}\big)]$. Hence, it must hold $\nu \in [1, \frac{h}{\delta})$. Then, either $\nu \in [\frac{h}{\delta^2}, \frac{h}{\delta}]$ or $\nu \in [1, \frac{h}{\delta^2})$. But if $\nu \in [\frac{h}{\delta^2}, \frac{h}{\delta}]$, then $\exists p_i^* = p_2$ such that $p_2 \in [\frac{\nu}{\delta}\big(1 - \frac{2\tau}{T}\big), \nu\big(1 - \frac{2\tau}{T}\big)]$.
We iterate this procedure until the $\lfloor \log_\delta h \rfloor$-th step.
Note that $p_{\lfloor \log_\delta h \rfloor}$ is for sure smaller than $\delta\big(1 - \frac{\lfloor \log_\delta h \rfloor \tau}{T}\big)$. Note also that from the step before, it must hold that $\nu \in [1, \frac{h}{\delta^{\lfloor \log_\delta h \rfloor - 1}})$. Hence, either $\nu \in (\frac{h}{\delta^{\lfloor \log_\delta h \rfloor - 1}}, \frac{h}{\delta^{\lfloor \log_\delta h \rfloor}}]$ or $\nu \in [1, \frac{h}{\delta^{\lfloor \log_\delta h \rfloor}})$. If the former holds, then $p_i^* =$

$p_{\lfloor \log_\delta h \rfloor}$. If the latter holds, then $(1 - \frac{\lceil \log_\delta h \rceil \tau}{T}) \in [\frac{\nu}{\delta}(1 - \frac{\tau}{T}), \nu(1 - \frac{\tau}{T})]$, but this is the final offered price within t_0. The contradiction is reached. $\square$

We can finally report the main theorem of this section, in which we show that mechanism $\mathcal{M}_2$ has a competitive ratio that is lower bounded by a constant, independently from the *unknown* probability distribution from which the valuations are sampled.

Theorem 18. *Let $(\lambda T)^\epsilon \geq \log_\delta h$ for $\epsilon \in (0, 1)$ and consider agent undiscounted valuations drawn i.i.d. from a distribution F with non-decreasing monotone hazard rate and linearly discounted by $\left(1 - \frac{t}{T}\right)$. Mechanism $\mathcal{M}_2$ has a competitive ratio of $\frac{\left(1-\frac{t_0}{T}\right)(1-\epsilon)}{\delta e}$ that is constant with respect to the choice of F.*

Proof (Lemma 18). Let $\lambda\tau = (\lambda T)^\epsilon$. For Lemma 17, there exists a price $p_i^* = p(I_i) = \frac{h}{\delta^i}(1 - \frac{i\tau}{T})$ such that $p_i^* \in \tilde{I}_i = \Big[\frac{\mathbb{E}[X_{\lambda T}]\left(1-\frac{i\tau}{T}\right)(1-\epsilon)}{\delta}, \mathbb{E}[X_{\lambda T}]\big(1 - \frac{i\tau}{T}\big)(1-\epsilon)\Big]$, for some $i \in \{1, \ldots, \lfloor \log_\delta h \rfloor\}$. The expected revenue of the mechanism in $[0, T]$ is at least that attained in a τ-lengthed interval, with $\tau \leq T$. Being $Y_{\lambda\tau, i\tau}$ the random variable of the maximum discounted valuation of agents arriving in the τ-lengthed interval $I_i = [(i - 1)\tau, i\tau]$. Notice that $Y_{\lambda\tau, i\tau}$ and $Y_{\lambda\tau, i'\tau}$ are different random variables if $i \neq i'$ because the valuations discount changes over time.

$$
\begin{aligned}
\mathbb{E}[\mathcal{R}(\mathcal{M}_2)] &= \mathbb{E}[\mathcal{R}(\mathcal{M}_2)]_{[0,T]} && (5.26)\\
&\geq \mathbb{E}[\mathcal{R}(\mathcal{M}_2)]_{[(i-1)\tau, i\tau]} && (5.27)\\
&\geq p_i^* \mathbb{P}(Y_{\lambda\tau, i\tau} \geq p_i^*) && (5.28)\\
&\geq p_i^* \mathbb{P}\Big(X_{\lambda\tau}\Big(1 - \frac{i\tau}{T}\Big) \geq \mathbb{E}\left[X_{\lambda T}\right]\Big(1 - \frac{i\tau}{T}\Big)(1-\epsilon)\Big) && (5.29)\\
&= p_i^* \mathbb{P}(X_{\lambda\tau} \geq \mathbb{E}\left[X_{\lambda T}\right](1-\epsilon)) && (5.30)\\
&= p_i^* \mathbb{P}\left(X_{\lambda\tau} \geq \mathbb{E}\left[X_{\lambda T}\right] \frac{\ln(\lambda\tau)}{\ln(\lambda T)}\right) && (5.31)\\
&\geq p_i^* \mathbb{P}(X_{\lambda\tau} \geq \mathbb{E}\left[X_{\lambda\tau}\right]) && (5.32)\\
&\geq \frac{p_i^*}{e} \geq \frac{\mathbb{E}\left[X_{\lambda T}\right]\left(1 - \frac{i\tau}{T}\right)(1-\epsilon)}{\delta e} && (5.33)\\
&\geq \frac{\mathbb{E}\left[X_{\lambda T}\right]\left(1 - \frac{t_0}{T}\right)(1-\epsilon)}{\delta e} && (5.34)
\end{aligned}
$$

where $\mathbb{E}[\mathcal{R}(\mathcal{M}_1)]_{[a,b]}$ is the expected revenue of the mechanism on an interval $[a, b]$. Inequality (5.29) holds because $X_{\lambda\tau}\left(1 - \frac{i\tau}{T}\right)$ is the maximum valuation drawn from F in the τ-lengthed i nterval weighted by the maximum discount, that is the minimum weight in that interval. $Y_{\lambda\tau,i\tau}$ is the maximum discounted valuation in I_i which means that is some valuation wighted by a discount $\left(1-\frac{t}{T}\right) \geq \left(1-\frac{i\tau}{T}\right)$. Inequality (5.32) follows from Lemma 14. A result by Barlow and Marshall (1964) implies that the probability of exceeding its expectation is at least $\frac{1}{e}$ for a non-decreasing monotone hazard rate distribution. We use this result in Inequality (5.33).

Now we compute a lower bound for the competitive ratio of the mechanism $\mathcal{M}_2$:

$$c_{\mathcal{M}_2} = \frac{\mathbb{E}[\mathcal{R}(\mathcal{M}_2)]}{\mathbb{E}[\mathcal{R}(\mathcal{M}^*)]} \geq \frac{\mathbb{E}[\mathcal{R}(\mathcal{M}_2)]}{\mathbb{E}[Y_{\lambda T}]} \tag{5.35}$$

$$= \frac{\mathbb{E}[X_{\lambda T}]}{\mathbb{E}[Y_{\lambda T}]} \frac{\left(1-\frac{i\tau}{T}\right)(1-\epsilon)}{\delta e} \tag{5.36}$$

$$\geq \frac{\left(1-\frac{t_0}{T}\right)(1-\epsilon)}{\delta e} \tag{5.37}$$

This concludes the proof. □

5.5 Experimental Results

In this section we provide a *simulation* of a real scenario. We refer to the *Motivating Example* 1.1.1. In Milan, a seller aims to rent a single room but she does not know the private valuations of the agents. However she can estimate the range of the initial undiscounted valuations $[1, h]$. Considering data from the metropolitan area of Milan, the value chosen for h is 2.8. We simulate a real setting in which agents arrive according to a Poisson process of parameter $\lambda = 2$ in a period of time $T = 20$. Both the number of agents and the arrival times are random variables. Each agent draw a sample $\tilde{v}$ from a probability distribution F, then her private valuation is discounted with respect to her arrival time. We consider 10 different probability distributions F. We are interested in evaluating the performance of $\mathcal{M}_1$ and $\mathcal{M}_2$. The parameters (δ, t_0) of mechanism $\mathcal{M}_2$ are optimized with respect the theoretical bound on the competitive ratio in the considered setting. In Figure 5.3, it is shown the two mechanisms facing an instance of the problem. The simulation is performed in *MATLAB*.

In this simulation we have considered 250 instances of the problem, 25 for each distribution. For each instance the selling price is compared with

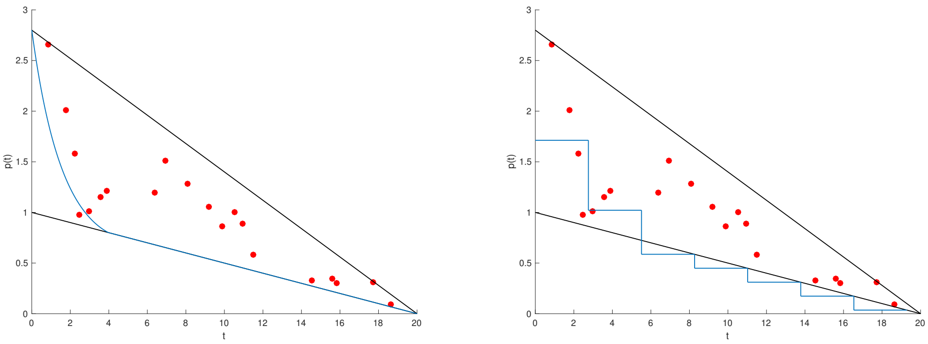

(a) Mechanism $\mathcal{M}_1$ *(b) Mechanism $\mathcal{M}_2$*

This is an instance of the problem. Each circle represents an agent, specifically her private valuation and arrival time. The blue plots are the pricing strategies of the two mechanisms. The black straight lines are the maximum and the minimum valuations.

Figure 5.3: Simulation: an Instance of the Problem

the *Y benchmark*, whose selling price is always the maximum valuations of the agents for that specific instance. The ratios between these two selling prices are collected in a dataset, shown in Figure 5.4.

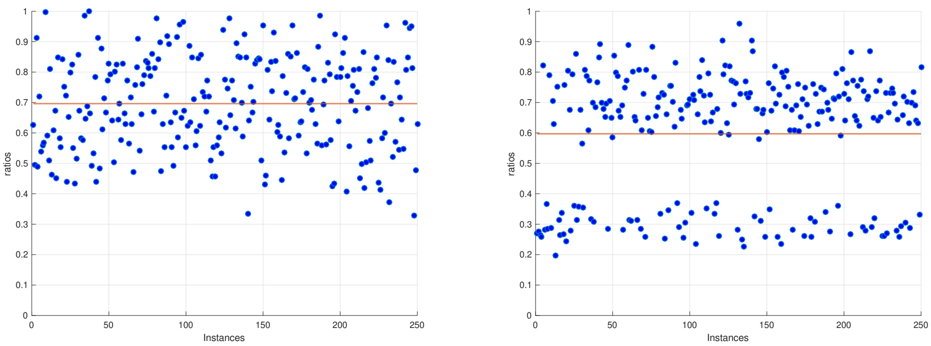

(a) Mechanism $\mathcal{M}_1$ *(b) Mechanism $\mathcal{M}_2$*

This is a dataset of 250 instances from the Simulation. Each circle represents the ratio between the selling price of the mechanism and the selling price of the *Y benchmark*. The red line is the mean computed over the 250 ratios.

Figure 5.4: Simulation: Comparison of the two Mechanisms

First of all, note that the expected performance of mechanism $\mathcal{M}_1$ exceeds the one of $\mathcal{M}_2$ of almost 0.1. Indeed, $\mathcal{M}_1$ reaches a mean competitive ratio of 0.6959, while $\mathcal{M}_2$ of 0.5997. Then, note that in Figure 5.4b there are two well-distinct regions. Indeed, $\mathcal{M}_2$ has only 7 possible selling prices and it is able to sell the room at least at its second highest price. Instead the selling prices of $\mathcal{M}_2$ are much more spread. Finally, for the most of the instances, note that $\mathcal{M}_1$ reaches at least a 0.5-*approximation* of the Y *benchmark*.

Chapter 6

Random Valuation: Optimization

This chapter is dedicated to a mathematical programming approach to the problem of *Random Valuation*. The analysis is a natural extension of the Chapter 4, where a first version of the *Maximum Violation Algorithm* was presented. Here we modify the algorithm in order to handle the RV setting. The algorithm will be quite different, but it keeps the same structure and the basic idea. We study the case with linear *discounting rate*, *i.e.* $\eta(t) = 1 - \frac{t}{T}$. As already said, this case is easily extendible to the degenerate *discounting rate* case $\eta(t) = 1$.
The mathematical model is almost the same of that presented in Chapter 5. But here we cannot use the continuous assumption on the initial private valuations of the agents. We consider a discrete set of undiscounted valuation $\tilde{v}_0, \ldots, \tilde{v}_{n-1}$. Therefore, also the unknown probability distribution F has to be considered a discrete probability distribution - whose PDF is denoted with $\rho_0, \ldots, \rho_{n-1}$, leading to a new formulation of the expected revenue of the benchmark. Moreover, there is a huge difference with respect to the Chapter 5: we faces the problem of the *unknown* distribution proposing a *distributionally robust* optimization. This approach leads to a new challenge not yet been addressed, the evaluation of a mechanism given a probability distribution F. This is not a trivial issue and we suggest two ways to deal with this challenge. Finally, we propose the *Maximum Violation Algorithm (II)*, or *MVA (II)*.

6.1 The Benchmark

The expected revenue of the benchmark mechanism is still the expected value of the maximum valuation of the agents inside the time horizon.

$$bench = \mathbb{E}[Y_{\lambda T}] = \mathbb{E}[\max_{i=1:N(T)} v_i] = \mathbb{E}[\max_{i=1:N(T)} \tilde{v}_i u_i] \tag{6.1}$$

Hence, let us compute $\mathbb{E}[Y_{\lambda T}]$ following the steps of (5.2). First, we derive the expression of $F_{V|N}$, the conditional CDF of v_i, then we compute $\mathbb{E}[Y_{\lambda T}]$.

$$\begin{aligned}
F_{V|N}(x,k) &= \mathbb{I}_{[0,1]}\{x\} \sum_{i=0}^{n-1} \rho_i \int_0^{x/\tilde{v}_i} g_u(u)du + \\
&\quad \mathbb{I}_{[1,h]}\{x\} \left(\sum_{i=0}^{\bar{i}} \rho_i \int_0^1 g_u(u)du + \sum_{i=\bar{i}+1}^{n-1} \rho_i \int_0^{x/\tilde{v}_i} g_u(u)du \right) = \\
&= \mathbb{I}_{[0,1]}\{x\} \cdot x \sum_{i=0}^{n-1} \frac{\rho_i}{v_i} + \mathbb{I}_{[1,h]}\{x\} \left(\sum_{i=0}^{\bar{i}} \rho_i + x \sum_{i=\bar{i}+1}^{n-1} \frac{\rho_i}{v_i} \right)
\end{aligned}$$

Where $\bar{i} = \left\lfloor (x-1)\frac{n-1}{h-1} \right\rfloor$. Observing that $F_{Y_{\lambda T}}$ is a mixture distribution, from Equation (5.7), where the mixture weights are $a_j = \frac{(\lambda T)^j e^{-\lambda T}}{j!}$, we compute the expected revenue of $Y_{\lambda T}$:

$$\begin{aligned}
\mathbb{E}[Y_{\lambda T}] &= \sum_{j=0}^{\infty} a_j \mathbb{E}\left[F^j_{V|N}(x,k)\right] = \sum_{j=0}^{\infty} a_j \int_0^h 1 - F^j_{V|N}(x,k)dx = \\
&= h - \sum_{j=0}^{\infty} a_j \left[\frac{1}{j+1} \left(\sum_{i=0}^{n-1} \frac{\rho_i}{v_i} \right)^j + \int_1^h \left(\sum_{i=0}^{\bar{i}} \rho_i + x \sum_{i=\bar{i}+1}^{n-1} \frac{\rho_i}{v_i} \right)^j dx \right]
\end{aligned}$$

To solve the integral, notice that $\forall x \in [\tilde{v}_k, \tilde{v}_{k+1}]$ the index $\bar{i}$ does not change. Hence, we can rewrite the term as follows:

$$\int_1^h \left(\sum_{i=0}^{\bar{i}} \rho_i + x \sum_{i=\bar{i}+1}^{n-1} \frac{\rho_i}{v_i} \right)^j dx = \sum_{k=0}^{n-2} \int_{v_k}^{v_{k+1}} \left(\sum_{i=0}^{k} \rho_i + x \sum_{i=k+1}^{n-1} \frac{\rho_i}{v_i} \right)^j dx$$

Finally, the expected revenue of the benchmark can be written as follows:

$$\mathbb{E}\left[Y_{\lambda T}\right] = h - \sum_{j=0}^{\infty} \frac{(\lambda T)^j e^{-\lambda T}}{j!} \left\{ \frac{1}{j+1} \left(\sum_{i=0}^{n-1} \frac{\rho_i}{v_i} \right)^j + \right.$$
$$\left. + \sum_{k=0}^{n-2} \left[\frac{\left(\sum_{i=0}^{k} \rho_i + v_{k+1} \sum_{i=k+1}^{n-1} \frac{\rho_i}{v_i}\right)^{j+1} - \left(\sum_{i=0}^{k} \rho_i + v_k \sum_{i=k+1}^{n-1} \frac{\rho_i}{v_i}\right)^{j+1}}{(j+1) \sum_{i=k+1}^{n-1} \frac{\rho_i}{v_i}} \right] \right\} \tag{6.2}$$

6.2 An Evaluation Procedure

We try now to answer a simple question: *Given a probability distribution F, how can we evaluate a mechanism?* We propose two possible answers. The first solution, discussed in this section, is to extend the pricing strategy $p(t) = p(t, \tilde{v})$ and computing the expected revenue of $p(t, \tilde{v})$ consequently; the second one, discussed in the next section, follows a new possible interpretation of the discrete problem.
Clearly, the mentioned extension does not imply that a seller should be able to post different prices for different undiscounted valuations. Indeed, it is just a useful notation *s.t.* $p(t) = p(t, \tilde{v})\ \forall \tilde{v}$. In a posted-price mechanism the seller faces the agents one-by-one. Let us build the expected revenue of a mechanism facing the agents one-by-one. First of all, consider the i^{th} agent with undiscounted valuation $\tilde{v}_i$ and arrival time $t_i \sim \Gamma(i, \lambda)$. Let us notice that:

$$\mathbb{P}\left(\tilde{v}_i \leq x\right) = \mathbb{P}\left(\tilde{v}_i \left(1 - \frac{t_i}{T}\right) \leq x \left(1 - \frac{t_i}{T}\right)\right) = \mathbb{P}\left(v_i \leq x \left(1 - \frac{t_i}{T}\right)\right)$$

Consequently: $\mathbb{P}\left(\tilde{v}_i \leq f(x)\right) = \mathbb{P}\left(v_i \leq f(x)\left(1 - \frac{t_i}{T}\right)\right)$.

If the item is still available, the necessary condition for an agent to buy the item is that her private valuation $v_i(t_i)$ is above the posted price at the time t_i. Let us rewrite this probability $\mathbb{P}\left(v_i \leq p(t_i)\right) = \mathbb{P}\left(\tilde{v}_i \leq \frac{p(t_i)}{1 - \frac{t_i}{T}}\right)$.
Hence, given a posted price mechanism $p(t)$, if the item is still available when the i^{th} agent enters in the market, the probability that the agent i buys the item is:

$$\mathbb{P}\left((\tilde{v}_i, t_i) \in \mathcal{D}\right) \tag{6.3}$$

Where $\mathcal{D} = \left\{1 \leq \tilde{v}_i \leq h, T\frac{p(t)}{T-t} \leq t_i \leq T\right\}$. Figure 6.1.

Consequently, the probability that the pair $(\tilde{v}_i, t_i)$ belongs to the complementary domain $\mathcal{D}^{\perp}$ is the probability that the agent does not buy the

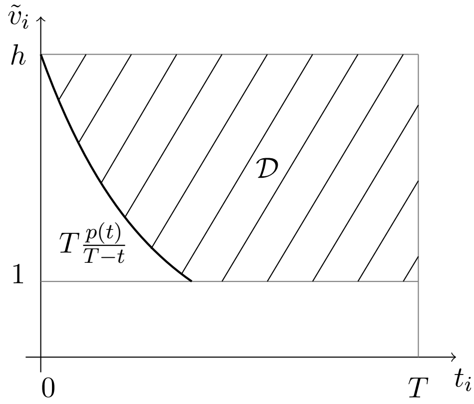

Figure 6.1: Domain $\mathcal{D}$

item (again, if the item is still available). Notice that the sum of the two probabilities is not 1, but it is equal to the probability that the i^{th} agent enters in the market within the deadline.
We can now derive the expression of the expected revenue of a certain mechanism $\mathbb{E}[\mathcal{R}]$. Let us expand the pricing strategy, imposing $p(t, \tilde{v}) = p(t) + 0 * \tilde{v}$. The pricing curve can be now evaluated in the domain $\mathcal{D}$, Figure 6.2. We define $\mathbb{E}[\mathcal{R}_i]$ as the expected revenue of the mechanism facing the i^{th} agent and assuming that the at time t_i the item is still available.

$$\mathbb{E}[\mathcal{R}_i] = \iint_{\mathcal{D}} p\left(\tilde{v}_i, t_i\right) f\left(\tilde{v}_i\right) f_{\Gamma,i}\left(t_i\right) d\tilde{v}_i dt_i \tag{6.4}$$

Where $f_{\Gamma,i}$ is the PDF of the random variable $t_i \sim \Gamma(i, \lambda)$.

The total expected revenue of the mechanism is therefore:

$$\mathbb{E}[\mathcal{R}] = \sum_{j=0}^{\infty} \mathbb{E}\left[\mathcal{R}_j\right] \cdot \mathbb{P}\left(\bigcap_{i=1}^{j} \left(\left(\tilde{v}_i, t_i\right) \in \mathcal{D}^{\perp}\right)\right) \tag{6.5}$$

6.3 A New Possible Interpretation: Markov Chains

In this section we introduce a new model. This model allows us to give a new interpretation of our setting, suggesting a very intuitive representation. Our posted-pricing problem can be seen as a big trade-off among posted prices and attracting customers. First of all, we analyse this model and we show that it is equivalent to the original one. Then, we use the proposed model to evaluate a certain mechanism. We will use this method to build the *Maximum Violation Algorithm (II)*.
Let us call $\mathcal{H}_1$ the model used so far: a Poisson process with parameter λ

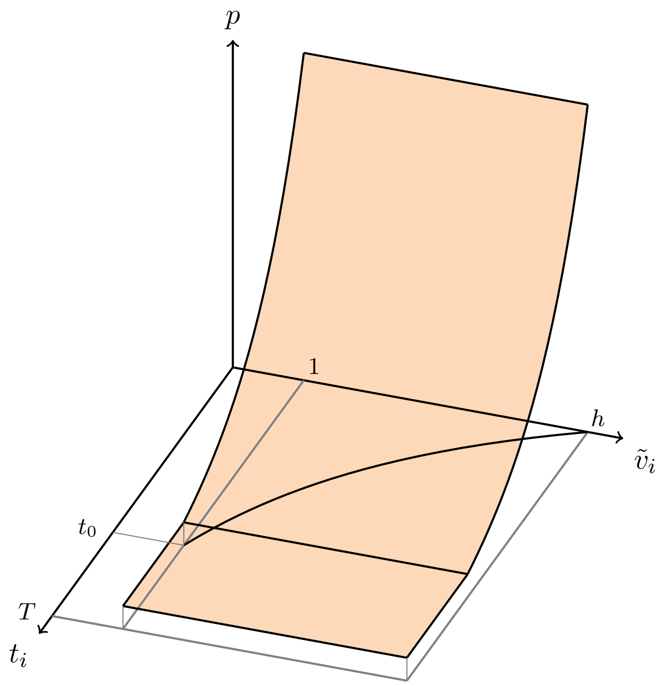

Figure 6.2: Evaluation of the pricing curve $p(t)$

manages the arrivals of the agents with private valuations $\tilde{v}_i \cdot \eta(t_i)$, where $\tilde{v}_i$ are *i.i.d.* from a discrete CDF F. At each undiscounted valuation $\tilde{v}_i \in [\tilde{v}_0, \ldots, \tilde{v}_{n-1}]$ is associated a probability $\rho_i \in [\rho_1, \ldots, \rho_{n-1}]$. We propose a new model, let us call it $\mathcal{H}_2$: consider n independent Poisson processes with rates $\lambda\rho_i$. Each of them manages the arrivals of agents with a specific undiscounted valuation.

Theorem 19. *Models $\mathcal{H}_1$ and $\mathcal{H}_2$ are equivalent.*

Proof (Theorem 19). Let randomly split the Poisson process into n subprocesses with probabilities $[\rho_1, \ldots, \rho_{n-1}]$. Each arrival is switched independently of each other arrival and independently of the arrival epochs. First consider a small increment $(t, t+\delta]$. The original process has an arrival in this incremental interval with probability $\lambda\delta$. Because of the independent increment property of the original Poisson process and the independence of the division of each arrival between the subprocesses, the new processes have the independent increment property, and from above have the stationary increment property. Thus each process is Poisson.
It remains to prove the independence of the processes.
Let $N_i(I_i)$ be the number of arrivals from subprocess i in the interval I_i.

Arrivals in non overlapping intervals are certainly independent. There may be independence only among $N_i(I)$ where I is an interval shared between several agents. But these represent the random split of the total number of arrivals $N(I)$ from the original process into n sets, with $N(I)$ distributed as a $Poisson(\lambda|I|)$. And the Poisson distribution is an infinitely-divisible distribution. Hence, $N_i(I)$ are independent. This concludes the proof. $\square$

How does a mechanism should handle the model $\mathcal{H}_2$? We try now to answer this question. Let us consider a mechanism whose pricing strategy $\bar{p}(t)$ is *s.t.* $\bar{p}(t) = (1 - \frac{t}{T})$ for $t \in [t_0, T]$ and satisfies the following: given $\bar{t}$ *s.t.* $\bar{p}(\bar{t}) = \tilde{v}_k(1 - \frac{\bar{t}}{T})$, then $\nexists\ t > \bar{t}$ *s.t* $\bar{p}(t) > \tilde{v}_k(1 - \frac{\bar{t}}{T})$, $\forall k$. Let us call t_k the first time instant *s.t.* $\bar{p}(t_k) = \tilde{v}_k(1 - \frac{t_k}{T})$. This means that t_k is the time instant at which the k^{th} subprocess is *activated*, *i.e.* from t_k on, if an agent arrives on subprocess k, she will get the item. This assumption ensures that if a subprocess is activated at a the time $\bar{t}$, it remains activated $\forall t \geq \bar{t}$. Let us call this assumption *activation assumption.* Model $\mathcal{H}_2$ allows us to talking about classes of agents: class a_k is composed of agents with undiscounted valuation $\tilde{v}_k$ and arrival rate $\lambda\rho_k$. We can shift the focus from the probability over the valuations, to the arrival frequency of classes of agents.

We are now interested in the behaviour of agents facing a mechanism $\bar{p}(t)$. To represent such a behaviour we use a *discrete-parameter Markov Chain.*

- *States*: there are $\frac{(n+1)(n+2)}{2}$ states specifying whether the item is *available* or *sold* and, if the latter holds, the class of agent that bought the item. We have two sets of states for each time instant t_i, for $i = 0, \ldots, n-1$ and two sets for T. A set represents a sold item, the other one the unsold item. In the latter there is always a unique state. In the former there are $n - 1 - i$ states, one for each class of agents that potentially bought the item. s_1 is the initial state. $s_{2,n-1}$ is the state for the item sold within the time t_{n-2} to an agent of class a_{n-1}; s_2^* for the item unsold within the time t_{n-2}. $s_{3,k}$ the state representing the item sold within the time t_{n-3} to the class of agent a_k; s_3^* the item unsold within the time t_{n-3}, and so on.

- *Transition Probabilities*: first notice that all the states corresponding to the sold item are terminal states. Hence the transition probabilities from those states are all equal to zero. The non-zero probabilities are related to the transitions from an unsold state to a sold one.

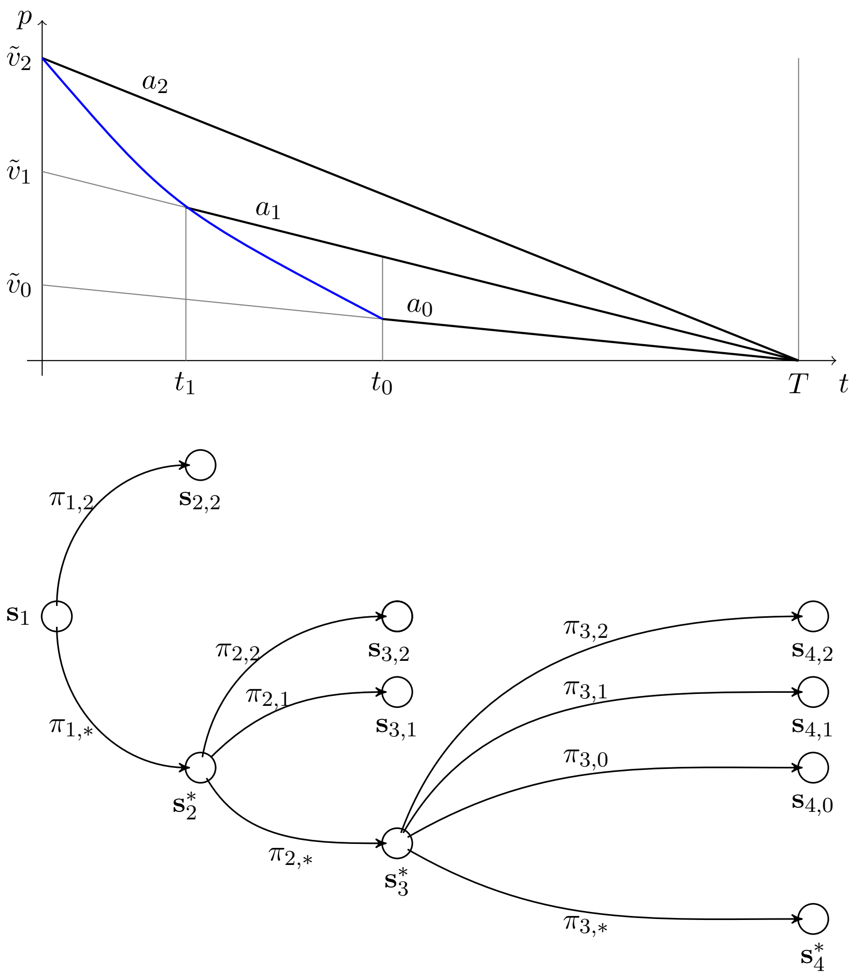

Figure 6.3: Markov Chain (1)

It is worth underlying the fact that this *Markov chain* describes the behaviour of the agents facing with a general mechanism $\bar{p}(t)$; the only assumption made on the pricing strategy is the activation assumption. As the example of Figure 6.3 suggests, the states lying on the same vertical line are parallel states. Each vertical line corresponds to a time instant t_i, for $i = 0, \ldots, n-1$ and T. Quite intuitively, at the time t_i, the Markov chain should stay in one of the parallel states s_{n-i}.
Let us compute the transition probabilities. We show the first two steps, then we generalize the computation. We refer to the case of Figure 6.3, where $n = 2$. The first step is straightforward. At time t_1 the probability that the item is sold corresponds to the arrival probability of an agent of class a_2 within t_1. Hence, $\pi_{1,2} = 1 - e^{-\lambda_2 t_1}$, while the $\pi_{1,2} = e^{-\lambda_2 t_1}$ is the

complementary probability, *i.e* the probability that zero agents of class a_2 arrive within t_1.
The second step. The probability to transition from s_2^* to $s_{3,2}$ is the probability to sell the item to an agent of class a_2 in $[t_1, t_0]$, knowing the item is still available in t_1. Hence, it is equal to the probability that an agent of class a_2 arrives before any agent of class a_1, in $[t_1, t_0]$, knowing that no agents of class a_2 arrive in $[0, t_2]$. The same reasoning can be done for the computation of $\pi_{2,1}$ and $\pi_{2,*}$.

Proposition 5. $\pi_{2,2} = \frac{\lambda_2}{\lambda_1+\lambda_2}\left(1 - e^{-(\lambda_1+\lambda_2)(t_0-t_1)}\right)$

Proof (Proposition 5). Let $\tilde{t}_2$ and $\tilde{t}_1$ be the *random variables* defining the arrival time instants of the first agent of classes a_2 and a_1, respectively. The two random variables are independent exponentially distributed. We first notice that:

$$\mathbb{P}\left(\tilde{t}_2 \le \tilde{t}_1 + t_1 | \tilde{t}_2 \ge t_1\right) = \frac{\mathbb{P}\left(t_1 \le \tilde{t}_2 \le \tilde{t}_1 + t_1\right)}{\mathbb{P}\left(\tilde{t}_2 \ge t_1\right)} =$$

$$= \frac{\mathbb{P}\left(\tilde{t}_2 \le \tilde{t}_1 + t_1\right) - \mathbb{P}\left(\tilde{t}_2 \le t_1\right)}{\mathbb{P}\left(\tilde{t}_2 \ge t_1\right)} = \frac{1 - e^{-\lambda\left(\tilde{t}_1 + t_1\right)} - 1 + e^{-\lambda t_1}}{e^{-\lambda t_1}} =$$

$$= 1 - e^{-\lambda \tilde{t}_1} = \mathbb{P}\left(\tilde{t}_2 \le \tilde{t}_1\right)$$

Then, we can derive $\pi_{2,2}$.

$$\begin{aligned}
\pi_{2,2} &= \mathbb{P}\left(\tilde{t}_2 \le \tilde{t}_1 \cap \tilde{t}_2 \le t_0 - t_1\right) = \int_0^{t_0-t_1} f_1\left(\tilde{t}_1\right) d\tilde{t}_1 \int_0^{t_1} f_2\left(\tilde{t}_2\right) d\tilde{t}_2 + \\
&\quad + \int_{t_0-t_1}^{\infty} f_1\left(\tilde{t}_1\right) d\tilde{t}_1 + \int_0^{t_0-t_1} f_2\left(\tilde{t}_2\right) d\tilde{t}_2 = \\
&= \int_0^{t_0-t_1} \lambda_1 e^{-\lambda_1 \tilde{t}_1} \left[-e^{-\lambda_2 \tilde{t}_2}\right]_0^{t_1} d\tilde{t}_1 + \int_{t_0-t_1}^{\infty} \lambda_1 e^{-\lambda_1 \tilde{t}_1} \left[-e^{-\lambda_2 \tilde{t}_2}\right]_0^{t_0-t_1} d\tilde{t}_1 = \\
&= \left(1 - \frac{\lambda_1}{\lambda_1+\lambda_2}\right)\left(1 - e^{-(\lambda_1+\lambda_2)(t_0-t_1)}\right) \\
&= \frac{\lambda_2}{\lambda_1+\lambda_2}\left(1 - e^{-(\lambda_1+\lambda_2)(t_0-t_1)}\right)
\end{aligned}$$

This concludes the proof. □

Hence:

$$\pi_{2,2} = \frac{\lambda_2}{\lambda_1+\lambda_2}\left(1 - e^{-(\lambda_1+\lambda_2)(t_0-t_1)}\right) \tag{6.6}$$

$$\pi_{2,1} = \frac{\lambda_1}{\lambda_1 + \lambda_2}\left(1 - e^{-(\lambda_1+\lambda_2)(t_0-t_1)}\right) \tag{6.7}$$

$$\pi_{2,*} = e^{-(\lambda_1+\lambda_2)(t_0-t_1)} \tag{6.8}$$

A general step. We simply adapt the same principles, considering that an agent from a certain class buys the item if she arrives before any agents from another class.

Proposition 6. $\pi_{k,i} = \frac{\lambda_i}{\sum_{j=n-1-k}^{n-1} \lambda_j}\left(1 - e^{-\sum_{j=n-1-k}^{n-1} \lambda_j (t_{n-k+1}-t_{n-k})}\right)$

Proof (Proposition 6). We simply generalise the Proposition 5. Let $\tilde{t}_j$ be the *random variable* defining the arrival time instant of the first agent of class a_j.

$$\pi_{k,i} = \mathbb{P}\left(\tilde{t}_i \leq \min\left\{\tilde{t}_{n-k}, \tilde{t}_{n-k+1}, \ldots, \tilde{t}_{i-1}, \tilde{t}_{i+1}, \ldots, \tilde{t}_{n-1}\right\} \cap \tilde{t}_i \leq t_{n-k+1}\right)$$

But notice that the random variables $\tilde{t}_j$ are *i.i.d.* with exponential distribution. Hence, we can write:

$$\mathbb{P}\left(\min\left\{\tilde{t}_{n-k}, \tilde{t}_{n-k+1}, \ldots, \tilde{t}_{i-1}, \tilde{t}_{i+1}, \ldots, \tilde{t}_{n-1}\right\} > t\right) = \mathbb{P}\left(\tilde{t}_{n-k} > t \cap \ldots \cap \tilde{t}_{n-1} > t\right)$$

$$\prod_{j=n-k}^{i-1} \mathbb{P}\left(\tilde{t}_j > t\right) \prod_{j=i+1}^{n-1} \mathbb{P}\left(\tilde{t}_j > t\right) = \prod_{j=n-k}^{i-1} e^{-\lambda_j t} \prod_{j=i+1}^{n-1} e^{-\lambda_j t}$$

$$= e^{-\left(\sum_{j=n-k}^{i-1} \lambda_j + \sum_{j=i+1}^{n-1} \lambda_j\right)t}$$

But this means that the random variable of the minimum $\tilde{t}_{min}$ has an exponential distribution with parameter $\lambda_{min} = \sum_{j=n-k}^{i-1} \lambda_j + \sum_{j=i+1}^{n-1} \lambda_j$. Consequently, $\pi_{k,i}$ can be written as follows:

$$\pi_{k,i} = \mathbb{P}\left(\tilde{t}_i \leq \tilde{t}_{min} \cap \tilde{t}_i \leq t_{n-k+1} - t_{n-k}\right)$$

That is exactly the case of Proposition 5, where $\tilde{t}_1$ is replaced by $\tilde{t}_{min}$ and $(t_0 - t_1)$ with $(t_{n-k+1} - t_{n-k})$. Hence:

$$\pi_{k,i} = \frac{\lambda_i}{\lambda_i + \lambda_{min}}\left(1 - e^{-(\lambda_i+\lambda_{min})(t_{n-k+1}-t_{n-k})}\right)$$

This concludes the proof. □

Finally, we can derive the transition probabilities for the generic step:

$$\pi_{k,i} = \frac{\lambda_i}{\sum_{j=n-1-k}^{n-1} \lambda_j}\left(1 - e^{-\sum_{j=n-1-k}^{n-1} \lambda_j (t_{n-k+1}-t_{n-k})}\right) \tag{6.9}$$

$$\pi_{k,*} = e^{-\sum_{j=n-1-k}^{n-1} \lambda_j (t_{n-k+1}-t_{n-k})} \tag{6.10}$$

A comment: in the discussed Markov chain there is neither price nor reward. But the chosen pricing strategy modifies the probabilities associated to the transitions. *How?* Imposing values to the activation time instants t_k. A seller cannot choose or even know the arrival rates of the different classes of agents, but she can decide how long she should wait for the highest valuations classes. It is more clear and practically intuitive the *big dilemma* of our problem, which can be seen as a great *trade-off* between posted prices and attracted agents classes. Take the first step, the seller offers a very high price, but she offers that price only to a small portion of agents, the higher valuation class. The trade-off is between the very high posted price and the discarded classes of agents, hence the cumulative probability of selling to that agents. At the second step, the price is slightly smaller, but this allows to attract one more agent class. Up to the final step, in which all the classes should buy the item, but with a very low price.

6.3.1 Another Evaluation Procedure

We use this new interpretation to evaluate a certain mechanism. First notice that the discussed Markov chain describes the agents behaviour facing a mechanism. But this behaviour is not completely observable by the seller. If the item is sold, the seller is not able to distinguish which class of agents bought the item; she can only know which classes do not buy it. Hence, we present now a simplified Markov chain, in which are presented only the observable states by the seller. Take the $\bar{k}^{th}$ step, for $i = n-k, n-k+1, \ldots, n-k$, the states $s_{\bar{k},i}$ represent all the possible attracted classes that should buy the item. But this is a hidden information. Hence, we simply unify states $s_{\bar{k},i}$ in a unique $s_{\bar{k}}$ state. We have:

- *States*: there are $2n+1$ states specifying whether the item is *available* or *sold*. We have two states for each time instant t_i, for $i = 0, \ldots, n-1$ and two states for T. A state represents a sold item, the other one the unsold item. s_1 is the initial state. s_2 is the state for the item sold within the time t_{n-2}; s_2^* for the item unsold within the time t_{n-2}, and so on.

- *Transition Probabilities*: first notice that all the states corresponding to the sold item are terminal states. Hence the transition probabilities from those states are all equal to zero. The non-zero probabilities are related to the transitions from an unsold state to a sold one. The probabilities $\pi_{i,*}$ are the same. We call π_i the transition probabilities from s_i^* to s_{i+1}: $\pi_i = \sum_{j=n-i}^{n-1} \pi_{i,j} = \left(1 - e^{-\sum_{j=n-1-k}^{n-1} \lambda_j (t_{n-k+1} - t_{n-k})}\right)$. π_i is equal to the probability that at least one agent from the activated classes arrives.

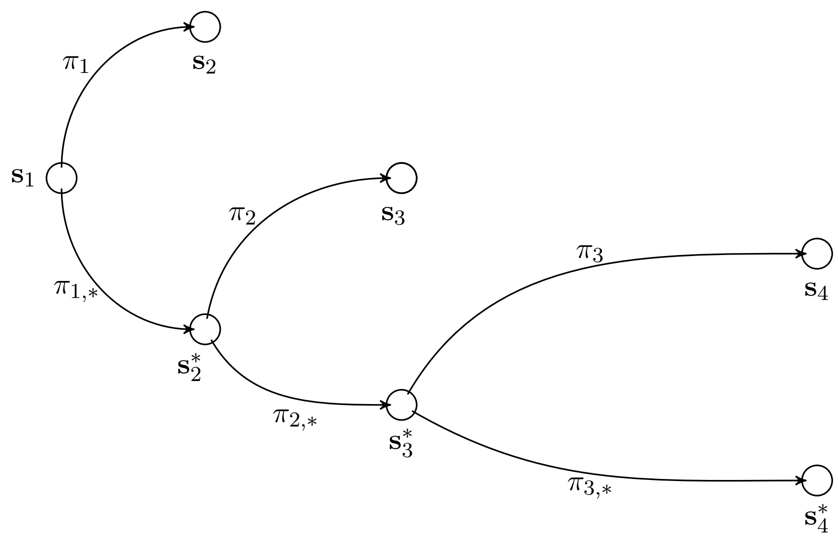

Figure 6.4: Markov Chain (2)

Given a discrete CDF F over the undiscounted valuations and a pricing strategy $p(t)$, which satisfies the activation assumption, we can now evaluate its expected revenue in the model $\mathcal{H}_2$, following the Markov chain (6.3.1). Since, for Theorem (19) models $\mathcal{H}_1$ and $\mathcal{H}_2$ are equivalent, the found expected revenue also holds for the original model $\mathcal{H}_1$.
First, consider a mechanism posting constant prices $\tilde{v}_i(1 - \frac{t_{i-1}}{T})$ for $t \in [t_i, t_{i-1}) = r_i$, satisfying the activation assumption. Its expected revenue is computed following the possible paths in the Markov chain:

$$\mathbb{E}\left[\mathcal{R}_{\text{const}}(F)\right] = \pi_1 r_{n-1} + \pi_{1,*}\left(\pi_2 r_{n-2} + \pi_{2,*}\left(\pi_3 r_{n-3} + \pi_{3,*}(\dots)\right)\right) =$$

$$= \pi_1 r_{n-1} + \pi_{1,*}\pi_2 r_{n-2} + \pi_{1,*}\pi_{2,*}\pi_3 r_{n-3} + \pi_{1,*}\pi_{2,*}\pi_{3,*}(\dots) + \dots$$

For a generic mechanism satisfying the activation assumption, its expected revenue is:

$$\mathbb{E}[\mathcal{R}(F)] = \sum_{i=1}^{n-1}\prod_{j=1}^{i-1}\pi_{j*}\int_0^{t_{n-i-1}-t_{n-i}} p\left(t + t_{n-i}\right)\sum_{j=n-i}^{n-1}\lambda_j e^{-\sum_{j=n-i}^{n-1}\lambda_j t}dt + \\ + \prod_{j=1}^{n-1}\pi_{j,*}\int_0^{T-t_0} p\left(t+t_0\right)\lambda e^{-\lambda t}dt \tag{6.11}$$

6.4 Maximum Violation Algorithm (II)

The previous sections of this chapter are tools used to develop the *Maximum Violation Algorithm* in the case of Random Valuation. The first version of this algorithm - 4.2 - basically consists in iterating two steps. In the first one an optimization is performed, in the second the algorithm searches for the maximum violation - which in that case corresponds to a valuation. In the present framework, the basic idea is the same. The algorithm is composed of two steps. Starting from a certain probabibility distribution $F^{(i)}$, in the first step the algorithm searches for a mechanism that optimize the ratio between the expected revenue of that mechanism and that of the benchmark, given $F^{(i)}$. We are able to evaluate a mechanism given a probability distribution; indeed, we can compute its expected revenue - section 6.2 or 6.3.1 - and the expected revenue of the benchmark - section 6.1. In the second step, the algorithm finds the *maximum violation* - that, in this case, is the probability distribution $F^{(i+1)}$ that minimize the competitive ratio of the mechanism just defined. This procedure is then iterated, performing a *distributionally robust* optimization. There are different ways to search for this probability distribution. Our choice is to use model $\mathcal{H}_2$, build a database of *MHR* discrete probability distributions and then make a choice inside this database.

Let us consider a family of mechanisms which offer a price $p(t) = \tilde{v}_k(1 - \frac{t}{T})$ for $t \in [t_k, t_{k-1})$, $k = 0, \dots, n-1$ and $p(t) = 1 - \frac{t}{T}$ for $t \in [t_0, T]$. Such a family is defined by the vector $\mathbf{t} = [t_0, \dots, t_{n-2}]$ (Figure 6.5), with $t_0 \geq t_1 \geq \dots t_{n-2}$. The *MVA (II)* is going to optimize such a vector.

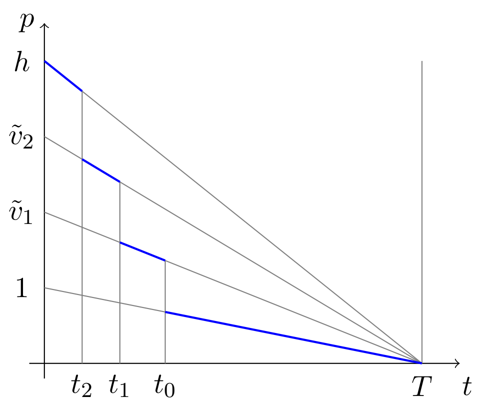

Figure 6.5: Family of Mechanisms $\boldsymbol{t}$

Note that this mechanisms' family satisfies the activation assumption.

Hence, given a mechanism $\mathbf{t}$ and a probability distribution F, we can compute the expected revenue of the mechanism using Equation (6.11). Let us call $\boldsymbol{\rho}$ the PDF of F. By imposing $\lambda_k = \lambda\rho_k$, we can compute:

$$\mathbb{E}[\mathcal{R}(\mathbf{t},\boldsymbol{\rho})] = \sum_{i=0}^{n-1} \tilde{v}_j \cdot \left(\prod_{j=i+1}^{n-1} e^{\lambda_j t_j}\right) \cdot \left[\frac{e^{-\sum_{k=i}^{n-1}\lambda_k t_{i-1}}\left(\sum_{k=i}^{n-1}\lambda_k\left(t_{i-1}-T\right)+1\right)}{T\sum_{k=i}^{n-1}\lambda_k} + \right.$$
$$\left. - \frac{e^{-\sum_{k=i}^{n-1}\lambda_k t_i}\left(\sum_{k=i}^{n-1}\lambda_k\left(t_i-T\right)+1\right)}{T\sum_{k=i}^{n-1}\lambda_k}\right] \tag{6.12}$$

Moreover, thanks to Equation 6.2 we can compute the expected revenue of the benchmark for the PDF $\boldsymbol{\rho}$. This allows us to define the following optimization problem. Given the parameters of the model (λ, T, h) and a finite set $\mathcal{P}$ of PDFs, we define the following *Random Valuation Optimization* problem:

$$\begin{aligned} & \textbf{max} \quad \alpha \\ \textbf{subject to:} \quad & \frac{\mathbb{E}[\mathcal{R}(\mathbf{t},\boldsymbol{\rho})]}{bench(\boldsymbol{\rho})} \geq \alpha && \forall \boldsymbol{\rho} \in \mathcal{P} \\ & t_i - t_{i-1} \leq 0 && \forall i = 1, n-1 \\ & t_0 \leq T \end{aligned} \tag{6.13}$$

Where the variables are α and $\mathbf{t}$.
Finally, we can present the *MVA(II)*. We denote by Π a arbitrary large set of the *MHR* probability distributions and with $\boldsymbol{\rho}^{(\mathbf{0})}$ a certain PDF. We define the following algorithm:

Algorithm 2 `Maximum Violation Algorithm(II)`

1: $\mathcal{P} \leftarrow \boldsymbol{\rho}^{(\mathbf{0})}$;
2: $(\alpha, \mathbf{t}) \leftarrow$ *RandomValuationOptimization*$(h, \lambda, T, \mathcal{P})$;
3: $\bar{\boldsymbol{\rho}} \leftarrow \min_{\boldsymbol{\rho}\in\Pi} c\left(\boldsymbol{\rho}, \mathbf{t}\right)$;
4: $\beta \leftarrow c(\bar{\boldsymbol{\rho}}, \mathbf{t})$;
5: **if** $\alpha - \beta \geq \varepsilon$ **then**
6: $\quad\mathcal{P} \leftarrow \mathcal{P} \cup \{\bar{\boldsymbol{\rho}}\}$;
7: $\quad$**go to** 2;
8: **end if**
9: **return** $\mathbf{t}, \beta$

Where $c(\bar{\boldsymbol{\rho}}, \mathbf{t}) = \frac{\mathbb{E}[\mathcal{R}(\mathbf{t},\boldsymbol{\rho})]}{bench(\boldsymbol{\rho})}$.

6.4.1 Experimental Results

We present here an application example of the *MVA(II)*. First, note the strongly non-linear dependency on the variables of the *Random Valuation Optimization* 6.13, specifically in the term of the expected revenue. Since this formulation is a non-linear program, also in this case *BARON* has to be employed as solver.
Let Π be composed by 2500 different discrete probability distributions over the $n = 10$ undiscounted valuations. As already discussed in 5.5, we consider $h = 2.8$, $\lambda = 2$ and $T = 20$.
The algorithm outputs the following results:

$$\mathbf{t} = [0.16, 0.27, 0.29, 0.373, 0.667, 0.813, 1.167541, 1.67541, 12.32563]$$

$$\beta = 0.489513$$

reaching a tolerance of $\varepsilon = 0.001$ in 12 iterations.

Chapter 7

Conclusions

In this thesis, for the first time in literature, we studied an *online*, *time-variant*, *model-free*, *revenue-maximization* pricing problem from an *economic mechanism design* point of view. We defined the optimal economic mechanism in different Identical Valuation scenarios, also proposing a method to handle the seasonability effects of the market. We introduced a mechanism following the *automated mechanism design* approach, which turned out to be very flexible and able to address different variations of the initial problem, for example the constrained model in which the seller can post only constant prices for a certain time interval. We then studied the Random Valuation scenario, proposing three different economic mechanisms. We proved that two of them have a constant lower bound of the competitive ratio. The third one is the output of a mathematical programming algorithm.

Our work handles a *single-item single-unit* pricing scenario. We believe that the methods studied in this thesis can be quite easily generalised. Let us refer to the *motivating example*, 1.1.1. Suppose a company aims to sell many units of rooms and apartments. None of these items are equal to the other ones. But they clearly share the same *features*, for example location, square footage, number of occupants and so on. Our economic mechanisms, combined together with some *Machine Learning* techniques of *feature selection*, can define the price of the single feature, the final prices of the rooms and apartments can be composed accordingly.

A final comment. In this thesis, we tried to solve a *pricing* problem very close to real economic scenarios. Reading this work, one might wonder: *Is that really a fair price?* But this is not really a recent question. Even in Ancient Greece, *Aristotle* asked himself exactly the same question. Let us close this work with a quotation from *John Kenneth Galbraith*, in his *Economics in Perspective: A Critical History* - *'Nothing has so engaged economic attention over the centuries as the need to persuade people that the price given by the market has a justification superior to all ethical concern'.* (Galbraith, 1987).

Bibliography

B. An, N. Gatti, and V. R. Lesser. Alternating-offers bargaining in one-to-many and many-to-many settings. *Annals of Mathematics in Artificial Intelligence*, 77(1-2):67–103, 2016.

L. M. Ausubel, P. Milgrom, et al. The lovely but lonely vickrey auction. *Combinatorial auctions*, 17:22–26, 2006.

M. Babaioff, L. Blumrosen, S. Dughmi, and Y. Singer. Posting prices with unknown distributions. In *In ICS*. Citeseer, 2011.

R. E. Barlow and A. W. Marshall. Bounds for distributions with monotone hazard rate, ii. *Annals of Mathematical Statistics*, 35(3):1258–1274, 09 1964. URL `https://doi.org/10.1214/aoms/1177703282`.

L. Blumrosen and T. Holenstein. Posted prices vs. negotiations: an asymptotic analysis. *EC*, 10:1386790–1386801, 2008.

N. Cesa-Bianchi, C. Gentile, and Y. Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2014.

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010.

V. Conitzer and T. Sandholm. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 132–141, 2004.

B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision support systems*, 43(1):192–198, 2007.

G. Farina and N. Gatti. Adopting the cascade model in ad auctions: Efficiency bounds and truthful algorithmic mechanisms. *Journal of Artificial Intelligence Research*, 59:265–310, 2017.

J. K. Galbraith. *Economics in perspective: A critical history*. Houghton Mifflin, 1987.

N. Gatti, A. Lazaric, M. Rocco, and F. Trovò. Truthful learning mechanisms for multi-slot sponsored search auctions with externalities. *Artificial Intelligence*, 227:93–139, 2015.

A. Gibbard. *Econometrica*. ISSN 00129682, 14680262. URL `http://www.jstor.org/stable/1914083`.

F. D. Giunta and N. Gatti. Bargaining over multiple issues in finite horizon alternating-offers protocol. *Annals Mathematics in Artificial Intelligence*, 47(3-4):251–271, 2006.

M. T. Hajiaghayi, R. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, pages 58–65, 2007.

D. He, W. Chen, L. Wang, and T.-Y. Liu. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.

R. Kleinberg and T. Leighton. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 594–605. IEEE, 2003.

R. Kleinberg, B. Waggoner, and E. G. Weyl. Descending price optimally coordinates search. *arXiv preprint arXiv:1603.07682*, 2016.

B. Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47, 2017.

A. M. Medina and M. Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 262–270, 2014.

D. S. Mitrinovic, J. Pecaric, and A. M. Fink. *Classical and new inequalities in analysis*, volume 61. Springer Science & Business Media, 2013.

M. Mohri and A. Munoz. Optimal regret minimization in posted-price auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, pages 1871–1879, 2014.

R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, USA, 2007. ISBN 0521872820.

J. Rong, T. Qin, and B. An. Dynamic pricing for reusable resources in competitive market with stochastic demand. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

S. M. Ross, J. J. Kelly, R. J. Sullivan, W. J. Perry, D. Mercer, R. M. Davis, T. D. Washburn, E. V. Sager, J. B. Boyce, and V. L. Bristow. *Stochastic processes*, volume 2. Wiley New York, 1996.

T. Roughgarden. Cs364a: Algorithmic game theory lecture 5: Revenue-maximizing auctions, 2013.

T. Sandholm and A. Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *International Workshop on Agent-Mediated Electronic Commerce*, pages 73–91. Springer, 2003.

M. A. Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *J. Econ. Theory*, pages 187–217, 1975.

Y. Shoham and K. Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.

J. Steffensen. On a generalization of certain inequalities by tchebychef and jensen. *Scandinavian Actuarial Journal*, 1925(3-4):137–147, 1925.

F. Trovo, S. Paladino, M. Restelli, and N. Gatti. Improving multi-armed bandit algorithms for pricing.

F. Trovo, S. Paladino, M. Restelli, and N. Gatti. Multi–armed bandit for pricing. In *Proceedings of the European Workshop on Reinforcement Learning (EWRL)*, 2015.

B. Zheng, L. Xiao, G. Yang, and T. Qin. Posting prices for a perishable item: (extended abstract). In *AAMAS*, 2016.

Appendix A

A.1 Proof of Lemma 14

We state the following variant of Chebyshev Inequality (Mitrinovic et al., 2013). It is useful for the proof of Lemma 14.

Lemma 20 (Mitrinovic et al. (2013))**.** *Suppose function $h(x)$ is positive and non-decreasing on $[a,b]$, function $g(x)$ is non-decreasing on $[a,b]$, and function $f(x)$ is continuous on $[a,b]$, then the following inequality holds,*

$$\frac{\int_a^b h(x)f(x)g(x)dx}{\int_a^b h(x)f(x)dx} \geq \frac{\int_a^b f(x)g(x)dx}{\int_a^b f(x)dx}$$

Lemma 20 is a variant of Chebyshev Inequality (Mitrinovic et al., 2013).

Proof (Lemma 14) Recall that $F_{X_{\lambda\tau}}(x) = e^{-\lambda\tau(1-F(x))}$ from proof of Lemma 13. We need to express the expected value of the maximum valuation of agents arriving in a τ-lengthed time interval, with $\tau \leq \tau' \leq T$, as follows:

$$\begin{aligned}
\mathbb{E}[X_{\lambda\tau}] &= \int_0^\infty x f_{X_{\lambda\tau}}(x)dx = \int_0^\infty 1 - F_{X_{\lambda\tau}}(x)dx = \int_0^\infty 1 - e^{-\lambda\tau(1-F(x))}dx \\
&= \int_0^\infty \frac{1-F(x)}{f(x)} \frac{1-e^{-\lambda\tau(1-F(x))}}{1-F(x)} dF(x) \\
&= \int_0^1 \frac{1}{H(F^{-1}(1-k))} \frac{1-e^{-\lambda\tau k}}{k} dk
\end{aligned}$$

Now we apply Lemma 20. F having non-decreasing monotone hazard rate implies that $h(k) = \frac{1}{H(F^{-1}(1-k))}$ is a non-decreasing function of k. Hence, $h(k)$ is non-decreasing and positive on $[0,1]$. $g(k) = \frac{1-e^{-\lambda\tau k}}{1-e^{-\lambda\tau' k}}$ is non-decreasing on $[0,1]$ and $f(k) = \frac{1-e^{-\lambda\tau' k}}{k}$ is continuous on $[0,1]$. We have,

$$\begin{aligned}
\frac{\mathbb{E}[X_{\lambda\tau}]}{\mathbb{E}[X_{\lambda\tau'}]} &= \frac{\int_0^1 \frac{1}{H(F^{-1}(1-k))} \frac{1-e^{-\lambda\tau k}}{k} dk}{\int_0^1 \frac{1}{H(F^{-1}(1-k))} \frac{1-e^{-\lambda\tau' k}}{k} dk} = \frac{\int_0^1 \frac{1}{H(F^{-1}(1-k))} \frac{1-e^{-\lambda\tau' k}}{k} \frac{1-e^{-\lambda\tau k}}{1-e^{-\lambda\tau' k}} dk}{\int_0^1 \frac{1}{H(F^{-1}(1-k))} \frac{1-e^{-\lambda\tau' k}}{k} dk} \\
&\geq \frac{\int_0^1 \frac{1-e^{-\lambda\tau k}}{k} dk}{\int_0^1 \frac{1-e^{-\lambda\tau' k}}{k} dk} = \frac{\int_0^{\lambda\tau} \frac{1-e^{-t}}{t} dt}{\int_0^{\lambda\tau'} \frac{1-e^{-t}}{t} dt} = \frac{Ein(\lambda\tau)}{Ein(\lambda\tau')} \\
&= \frac{\gamma - Ei(-\lambda\tau) + \ln(\lambda\tau)}{\gamma - Ei(-\lambda\tau') + \ln(\lambda\tau')} \geq \frac{\ln(\lambda\tau)}{\ln(\lambda\tau')}
\end{aligned}$$

where $Ein(x) = \int_0^x \frac{1-e^{-t}}{t} dt$ is the entire exponential integral function, $Ei(x) = \int_{-\infty}^x \frac{e^t}{t} dt$ is the exponential integral function and $\gamma \approx 0.577$ is the Euler's constant. $\square$