UCLouvain

École polytechnique de Louvain

Optimal Decision Making in Black-Box Online Auctions Markets

Author: **Jonathan Sarteel**
Supervisor: **Raphaël Jungers**
Readers: **Jean-Charles Delvenne, Pierre Devolder**
Academic year 2018–2019
Master [120] in Mathematical Engineering

Abstract

A commonly studied problem in the context of Algorithmic Game Theory is the problem of behavior optimization in infinitely repeated games. While the computation of equilibria in the case of non-repeated games is generally a well-defined problem, the computation of equilibria in infinitely repeated games can be quite difficult. Multiple methods have however been developed in order to tackle this problem, notably in the context of repeated online ad auctions, but most of them require strong assumptions about the behavior of the other players.

In this thesis, we propose a method for behavior optimization in repeated ad auctions basing ourselves on the assumption that the highest adverse bids for an ad slot are independent realizations of unknown random variables. This assumption is used to show that our problem of behavior optimization can be seen as a problem of censored regression of these unknown random variables. We use our estimation of the different distributions to formulate a probabilistic optimization problem. Finally, we validate the performance of our algorithm using simulated data and conclude this thesis by providing numerical results proving the convergence of our algorithm to the optimal bid.

Acknowledgments

I would like to take this opportunity to thank all the people who have contributed to this master thesis.
I would first like to thank my thesis supervisor Raphaël Jungers for his patience and his dedication. His constructive comments have helped me a lot and the research made during this master thesis could not have been performed without his continued support and assistance.
I would also like to thank Jean-Charles Delvenne and Pierre Devolder for agreeing to read my thesis despite their many other academic activities.
Finally, I would like to thank my friends and my family for their help and continued support during this final master year.

Contents

1 Introduction

1.1 Context

Since the beginning of the 2000's, online advertising has become an increasingly important part of the global economy. With $269 billion being spent yearly for online ads[1], this sector has become an important part of the advertising industry, making it unavoidable for the launch of any marketing campaign. This growing interest in the sector has led to the emergence of an important literature on the subject and to an increasingly precise model for online ad auctions, generally using models from Game Theory.

This thesis will focus on the problem of the computation of equilibria in infinitely repeated auctions. While the problem of computation of equilibria in the case of non-repeated games is generally a well-defined problem, the same problem in infinitely repeated games can be quite difficult to solve. Multiple methods have however been developed in order to tackle this problem, but they generally require strong assumptions about the behavior of the other players. We will here introduce a model based on the assumption that the highest adverse bids in repeated auctions are all realizations of an unknown random variable. This assumption is analogous to some classical models in finance, where, despite strategic behavior at the individual level, stochastic models of the stock prices have shown to be quite efficient for decision-making (see [1] by *Amin et al.*).

In this thesis, we will focus on two main contributions: the use of censored estimators for equilibrium computation and the formulation of our problem as a problem of probabilistic optimization in order to reach equilibrium behavior in repeated auctions. This is **not** a thesis about fundamentally developing new methods for censored estimation: the different methods we are going to use in this thesis are taken straight out of the literature on the subject, as we believe it is already sufficient for the application we have in mind. The goal of this thesis is mostly to

[1]https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/

show how these methods of censored estimators, that are, to our knowledge, mostly used in the context of life insurance models, can also be applied in other contexts and how they can help improve our knowledge about the convergence properties of repeated games from a Game Theoretical point of view.

While we are not the first to study the use of censored estimation for adverse bid estimation, we found that the two papers already written on the subject (see the works of *Amin et al.* in [1] and *Wu et al.* in [2]) could be improved: in fact, both papers focused on second-price auctions, even though they represent the easiest study case, and only the paper by *Amin et al.*[1] explain how the estimation of the clearing price can be used in bid optimization, even though this is a non trivial task. We found as well that the method developed by *Amin et al.* could be improved by viewing the problem as a *probabilistic optimization problem* (also called *chance-constrained optimization problem* in the literature) and not as a Markov Decision Process, since the latter requires a discretization that becomes more and more computationally expensive if we take part in a large number of auctions and does not allow for more complex constraints than a basic budget constraint. The objective of this thesis is thus to give a complete framework for the problem of probabilistic optimization in repeated auctions using censored estimation, to develop an algorithm using this framework and to prove the convergence of this algorithm. This is, to our knowledge, the first paper on the subject viewing the problem of bid optimization in repeated auctions as a probabilistic optimization problem and developing methods allowing to solve the problem efficiently.

1.2 Structure of the thesis and contributions

This thesis will be divided into multiple chapters.

In Chapter 2, we will introduce multiple concepts of Game Theory and then model online ad auctions as infinitely repeated Bayesian games. We will then explain in detail how these Game Theoretical models relate to real-life online ad auctions and how advertisers can use them in order to better understand the game they are playing in. This is mostly a state of the art of the mechanisms behind online ad auctions and of the equilibrium properties these auctions might have.

In the chapters following Chapter 2, we will present the new methods we have developed and our contributions to the literature on the subject. The specific contributions of this thesis to the literature on the subject are the following:

- **Censored estimation of the distributions of the highest adverse bids**. In Chapter 3, we will explain in detail our methodology for the estimation of the distributions of the highest adverse bids and how this methodology can be used for behavior optimization in repeated Bayesian games. In fact, it is possible as an advertiser to estimate the distribution of the highest adverse bids using only the results of the auctions he has taken part in. The methodology used in this chapter is to consider the results of repeated ad auctions as censored observation of random variables, which makes it then possible to estimate the distributions of these random variables using censored estimation methods. As we have said before, our contributions in this chapter are not the algorithms used (they are taken straight from the literature) but the adaptation of these algorithms to our problem and empirical results using simulations.

- **Probabilistic Optimization**. In Chapter 4, we model our problem of behavior optimization in repeated auctions as a probabilistic optimization problem and develop some algorithms that can be used in order to solve efficiently this optimization problem using the estimation of the highest adverse bids developed in Chapter 3. This model of our problem as a probabilistic optimization is part of our contribution to the literature: to our knowledge, we are the first to formulate the problem of bid optimization in repeated auctions as a probabilistic optimization problem.

2 | Game theoretical model

2.1 Basic concepts of Game Theory

Since this thesis is devoted to the use of Game Theory principles for the optimization of the behavior of an advertiser in online ad auctions, we are going to need to define some concepts of Game Theory first. In order to do so, we will use some reference works on the subject. For the definition of concepts purely related to Game Theory, we will use the reference work on the subject by *Myerson* in [3] and the course notes written by *Jungers et al.* in [4], based on Myerson's work. For the definition of the concepts more related to mechanism design and Algorithmic Game Theory in general, we will use the reference work on the subject written by *Nisan et al.* in [5]. These three works on the subject will be used during the course of this chapter in order to introduce all the necessary concepts.

We can start by explaining the concept of *auctions*. We use the following definition given in [4].

Definition 2.1.1: Auctions

An auction is a method of allocating resources among selfish agents by offering items for bid, taking bids, and then selling them to the highest bidder. There exist two participating sides in an auction: the side of the *seller*, which is the agent selling the resources and the side of the *buyers* or *bidders*, which are the agents submitting bids to the sellers in order to get one of multiple resources.

The intrinsic idea behind the concept of auctions is that every bidder in an auction has a *valuation* of the items that are sold. We use the definition given in [4].

Definition 2.1.2: Valuation

The valuation of an item for a bidder is the highest price a bidder is willing to purchase an item for.

In order to introduce some new properties, we will first need to introduce a rigorous model of repeated auctions. In order to do so, we introduce the concept of a game with strict incomplete information as defined in [5].

Definition 2.1.3: Game with strict incomplete information

A game with strict incomplete information is a game with n players with the following ingredients:

- Each player has a set of actions X_i
- Each player has a set of types T_i. A value $t_i \in T_i$ is the private information i has
- Each player has an utility function $u_i : T_i \times X_1 \times ... \times X_n \to \mathbb{R}$, where $u_i(t_i, x_1, ..., x_n)$ is the utility achieved by player i if his type is t_i and the profile of actions taken by all other players is $x_1, ..., x_n$
- Each player is a strategy, which is a function $s_i : T_i \to X_i$

Now that we have described the concept of games with strict incomplete information, we can introduce the concept of Bayesian games, which is closely related. The definition used originates from [5].

Definition 2.1.4: Bayesian games

Bayesian games are games with strict incomplete information where each player has beliefs with known probability distribution about the private information of each other player.

While this version of the definition is intuitive, it is a bit insufficient when it comes to the full modelling of a repeated Bayesian game. To be able to precisely model repeated Bayesian games, we are going to introduce a more formal definition based on the works of *Jungers et al.* in [4] and of *Forges et al.* in [6].

Definition 2.1.5: Repeated Bayesian game

A repeated Bayesian game is a structure of the form

$$\Gamma^r = (E, \Theta, (X_i, S_i, T_i, u_i, s_i, p_i)_{i \in E}, q),$$

where

- E is the set of *players*
- Θ is the set of *states of the world*
- X_i is the set of *actions* available to player i
- S_i is the set of *signals* that player i may receive
- T_i is the set of *types* of each player. A value $t_i \in T_i$ is the private information i has
- u_i is the *utility function* of player i. The function for each player is $u_i : T_i \times X_1 \times ... \times X_n \to \mathbb{R}$, where $u_i(t_i, x_1, ..., x_n)$ is the utility achieved by player i if his type is t_i and the profile of actions taken by all other players is $x_1, ..., x_n$
- $s_i : T_i \to X_i$ is the *strategy* of each player
- $p_i \in \Delta(\Theta|S_i)$ is the *probability distribution over the states of nature*. Each player has different views of the probability distribution over the states of the nature given the signal send to them but the exact state of nature is generally unknown to them
- $q : X \times \Theta \to S \times \Theta$ is the *transition function* between each distribution for each player

This is thus the general definition that we are going to use in order to model the auctions described in this thesis.

In the definition of repeated Bayesian games, we use the concept of the *strategy* of a player. We can further explain this concept using the definition given by *Polak* in [7].

Definition 2.1.6: Strategy of a player

A player's strategy is any of the options which he chooses in a setting where the outcome depends not only on their own actions but also on the actions of others.

The strategies of players are heavily dependent on the type of games used. Strategies of players can have a lot of different properties. One of the important properties of player's strategies is the concept of *weakly-dominant strategies*, defined in [3] in the following way.

Definition 2.1.7: Weakly-dominant strategy

A weakly-dominant strategy is a strategy that earns a player a payoff at least as high as any other strategy, regardless of the moves of the other players.

This leads to the definition of *Dominant Strategy Incentive Compatible* (or DSIC) auction mechanisms, defined in [5] in the following way.

Definition 2.1.8: Dominant Strategy Incentive Compatible mechanism (DSIC)

An auction is Dominant Strategy Incentive Compatible if truth-telling is a weakly-dominant strategy. In other words, this means that it is in your best interest to bid your true valuation in the auction.

The concept of DSIC mechanisms is interesting for the design of auctions: since we consider as the best outcome to an auction the outcome where the highest bidder is the one receiving the object, we might want the bidders to tell their true valuations in order to give the highest bidder the object at a "fair" price.

2.2 Types of online ad auctions

Now that we have talked in detail about the concept of auctions in the context of Game Theory, we can come back to our case of sponsored-search and online ad auctions. The process of selling ads on the internet can be viewed as an auction, where the items sold are ad slots on a website and where buying an ad slot gives you the permission of putting an ad in it. These ad auctions have a number of different agents participating in it. We are going to group these agents into three categories: the *publishers*, the *advertisers* and the *customers*.

The *publishers* are the sellers in our auction: they possess a number of ad slots on their website and are selling these ad slots to other companies who might be interested in them. An example of publishers in the context of online ad auctions are social networks and search engines, such as Google or Facebook.
The *advertisers* are the bidders in our auction: they have ads they want to be published on different websites and they are willing to pay money in order to have the possibility to do so.
Finally, the *customers* are the people who see the ads while browsing the internet and that sometimes click on it. While the customers do not have a direct role in ad auctions, they play an indirect role in them; in fact, they are the one giving

the value to an ad slot, by occasionally clicking on it and buying products on the advertised website.

Now that we have described the different agents taking part in the auctions, we can try to describe how these agents interact together. This section will thus be devoted to the study of the different types of auctions that can be used by the publishers to bill their customers. For each auction described, we will follow the same structure: we will first model the auction as a Bayesian repeated game using the definition used in the previous section, then describe the properties of the auction. We will focus on the properties that are relevant to our thesis: whether the auction is DSIC or not and the types of equilibrium that exist for the auction in the infinitely repeated case.

2.2.1 Single-item auctions

Single-item auctions are not the most used type of auctions for online ads (see following sections); they are however a good introduction to the subject as they are easier to understand than multiple-item auctions. In this section we will detail the two types of single-item auctions used in the context of online advertising: first- and second-price auctions.

Second-price auctions

We can intuitively describe the mechanism of the second-price auction using the work done in [4].

Definition 2.2.1: Second-price auctions

Let there be N bidders competing for a single item. Let us define the set of bidders as $E = \{1, ..., N\}$. and the set of bids received as $B = \{b_1, ..., b_N\}$ (supposing the bids and the bidders are ranked in decreasing order of bids). Then the highest bidder receives the item at price b_2.

As a repeated Bayesian game, the second-price auction can be modelled in the following way.

Definition 2.2.2: Second-price auctions as repeated Bayesian game

The second-price auction can be modelled as a repeated Bayesian game with a structure of the form:

$$\Gamma^r = (E, \Theta, (X_i, S_i, T_i, u_i, s_i, p_i)_{i \in E}, q),$$

where

- $E = \{1, ..., N\}$ is the total set of bidders to the auctions
- $\Theta = \{firstTurn, win(1), ..., win(N), stop\}$ represent all the states where the world might be. The state $firstTurn$ represents the state where the game has just begun and therefore has no previous winner. The states $win(i)$ represent the state where the previous auction was won by the player i
- $X_i = x_i \in (0, \infty)$ represent the set of moves that each player can play, which is the bid he makes for an item
- $S_i = \{firstTurn, win, lose, stop\}$ represent the signal send to each player after an auction, which is whether the game stops next turn or not and whether the player has won or lost last turn
- $T_i = t_i \in (0, \infty)$ represent the private information of player i which is the valuation he has for the item being sold
- u_i represent the utility function of each player and can be expressed by the following formula: $u_i(t_i, x_1, ..., x_N) = (t_i - \max_{j \neq i} x_j)\mathbb{1}\{x_i \geq \max_{j \neq i} x_j\}$, where $\mathbb{1}\{A\}$ represents the indicator function for the set A.
- s_i is the strategy of player i. Since we want to keep the model of the auction general, we are not going to make any particular assumption on this function
- p_i is the belief in the state of nature. If the state of nature is $firstTurn$ or $stop$ then each player knows the state of nature with certainty, in other words $p_i(firstTurn|s_i = firstTurn) = 1$ or $p_i(stop|s_i = stop) = 1$ for all players. If a player loses then he does not know exactly which player is the winner. We do know as well than $p_i(win(i)|s_i = win) = 1$
- q is the transition function. It can be expressed as:

$$q(x_1, ..., x_N, \theta) = \begin{cases} (firstTurn, ..., firstTurn) & \text{if first turn} \\ (stop, ..., stop, stop) & \text{if } \theta = stop \\ (lose, ..., win, ..., lose, win(i)) & \text{if } x_i = \max_{j \neq i} x_j \end{cases}$$

Once we have defined the auction formally as a Bayesian repeated game, we can try to state its properties. For that we are going to introduce some new concepts. The following definitions originates from [5].

Definition 2.2.3: Allocation rule

An allocation rule is the per auction dependent rule of the distribution of the items between the different bidders (for instance, the highest bidder win). A payment rule is a per auction dependent rule of the payments associated with the distribution of the items to the different bidders (for instance, the winner pays the price of the second highest bidder).

Definition 2.2.4: Implementable allocation rule

An allocation rule is implementable if there exists a payment rule such that the sealed-bid auction is DSIC.

Definition 2.2.5: Monotone allocation rule

An allocation rule is monotone if, for each bidder and each bid by the other bidders, the allocation is non decreasing in its bid. In other words, bidding higher in a monotone allocation rule can only lead you to win more items.

Now that we have defined some properties of allocation rules, we can use them to introduce Myerson's lemma, as it is defined in [3].

Lemma 2.2.1: Myerson's lemma

For a given environment:

1. An allocation rule is implementable if and only if it is monotone
2. If an allocation rule is monotone then there is a unique payment rule such that the sealed-bid mechanism is DSIC

Myerson's lemma has important corollaries for every type of auction. For that we are going to need the concept of quasilinear utility functions, as it is defined is [5].

Definition 2.2.6: Quasilinear utility function

A quasilinear utility function is a utility function such that, if the outcome is x and the payment is p_i, then the total utility of agent i is:

$$u_i(x) = v_i(x) - p_i$$

This definition leads us to conditions for the second-price auction to be DSIC, as proven in [3].

Corollary 2.2.1

Assuming quasilinear utility functions and private values for the agents, the second-price auction is DSIC.

This means that a dominant strategy in the case of repeated second-price auctions is simply for each player to bid its true valuation, regardless of the bids of the other players. This has consequences for the existence of equilibria in the case of infinitely repeated auctions. In order to describe the equilibrium properties of second-price auctions, we will introduce during this thesis multiple formal definition of equilibria in the case of infinitely repeated auctions. We will start by introducing the most well-known type of equilibrium in Game Theory, the *Nash equilibrium*, as defined in [3].

Definition 2.2.7: Nash equilibrium

Let Γ^r be a repeated Bayesian game with N players. A strategy profile $\sigma = (\sigma_1, ..., \sigma_N)$ is a Nash equilibrium of Γ^r iff

$$u_i(\sigma) \geq u_i(\sigma_{-i}, \tau_i) \quad \forall i \in N, \tau_i \in \Delta(s_i)$$

where (σ_{-i}, τ_i) denote the strategy profile in which the i^{th} component is τ_i and all the other components are as in σ and $\Delta(s_i)$ represents the set of all possible strategies for player i.
In other words, a strategy profile is a Nash equilibrium if and only if no player could increase his expected payoff by unilaterally deviating from the prediction of the randomized-strategy profile.

This definition can thus be used to prove the existence of an equilibrium in repeated second-price auctions. In fact, since it is a dominant strategy for each player to bid its true valuation, it is possible to see that for each player to bid its true valuation is a Nash equilibrium, since no player can improve his expected payoff by unilaterally deviating from this strategy. While this may seem like an interesting property, we will show here that this property is actually quite trivial, because of the consequences of a so-called *Folk theorem*. We will thus introduce this theorem. In order to do that, we are going to need some new concepts. We will use the definitions from [4].

Definition 2.2.8: Minimax value

The minimax value of a player is the highest value that the player can be sure to obtain without knowing the actions of the other players. In other words, if $\underline{v}_i$ is the minimax value of player $i \in E$, then

$$\underline{v}_i = \min_{x_{-i} \in X_{-i}} \max_{x_i \in X_i} u_i(x_{-i}, x_i)$$

where $X_{-i} = X_1 \times ... \times X_{i-1} \times X_{i+1} \times ... \times X_N$.

Definition 2.2.9: Enforceable payoff profile

A payoff profile $(u_1, ..., u_N)$ is enforceable if $\forall i \in E$, $u_i \geq \underline{v}_i$.

Definition 2.2.10: Feasible payoff profile

A payoff profile $(\bar{u}_1, ..., \bar{u}_N)$ is feasible is $\forall i \in E$, $\exists \alpha_x \in \mathbb{Q}^+$ such that $\sum_{x \in X} \alpha_x u_i(x) = \bar{u}_i$ and $\sum_{x \in X} \alpha_x = 1$.

These definitions can then be used to introduce a Folk Theorem for infinitely repeated games, as summarized in [4].

Theorem 2.2.1: Folk Theorem

In an infinitely repeated game, every payoff profile that is both enforceable and feasible is a Nash equilibrium.

This thus means that in an infinitely repeated game, the set of Nash equilibria is very large, possibly infinite! However, all Nash equilibria in an infinitely repeated game are not equivalent for all the players: a player might obtain a higher utility in some Nash equilibrium than in another. This observation confirms the need for methods for behavior optimization in repeated auctions, which is the main subject of this thesis.

Now that we have talked about the main properties of second-price auctions, we can start to talk about the second type of single-item auctions majoritarily used on the internet: first-price auctions.

First-price auctions

We can describe intuitively the mechanism behind first-price auctions using the definition given in [4].

Definition 2.2.11: First-price auctions

Let there be N bidders competing for a single item. Let us define the set of bidders as $E = \{1, ..., N\}$ and the set of bids as $B = \{b_1, ..., b_N\}$ (supposing the bids and the bidders are ranked in decreasing order of bids). Then the highest bidder receives the item at price b_1.

As a repeated Bayesian game, the first-price auction can be modelled in the following way.

Definition 2.2.12: First-price auctions as repeated Bayesian game

The first-price auction can be modelled as a repeated Bayesian game with a structure of the form:

$$\Gamma^r = (E, \Theta, (X_i, S_i, T_i, u_i, s_i, p_i)_{i \in E}, q),$$

where

- $E = \{1, ..., N\}$ is the total set of bidders to the auctions
- $\Theta = \{firstTurn, win(1), ..., win(N), stop\}$ represent all the states where the world might be. The state $firstTurn$ represents the state where the game has just begun and therefore has no previous winner. The states $win(i)$ represent the state where the previous auction was won by the player i
- $X_i = x_i \in (0, \infty)$ represent the set of moves that each player can play, which is the bid he makes for an item
- $S_i = \{firstTurn, win, lose, stop\}$ represent the signal send to each player after an auction, which is whether the game stops next turn or not and whether the player has won or lost last turn
- $T_i = t_i \in (0, \infty)$ represent the private information of player i which is the valuation he has for the item being sold
- u_i represent the utility function of each player and can be expressed by the following formula:
  $u_i(t_i, x_1, ..., x_N) = (t_i - x_i)\mathbb{1}\{x_i \geq \max_{j \neq i} x_j\}$
- s_i is the strategy of player i. Since we want to keep the model of the auction general, we are not going to make any particular assumption on this function
- p_i is the belief in the state of nature. If the state of nature is $firstTurn$ or $stop$ then each player knows the state of nature with certainty, in other

words $p_i(firstTurn|s_i = firstTurn) = 1$ or $p_i(stop|s_i = stop) = 1$ for all players. If a player loses then he does not know exactly which player is the winner. We do know as well than $p_i(win(i)|s_i = win) = 1$

- q is the transition function. It can be expressed as:

$$q(x_1, ..., x_N, \theta) = \begin{cases} (firstTurn, ..., firstTurn) & \text{if first turn} \\ (stop, ..., stop, stop) & \text{if } \theta = stop \\ (lose, ..., win, ..., lose, win(i)) & \text{if } x_i = \max_{j \neq i} x_j \end{cases}$$

Now that we have formally defined first-price auctions as a repeated Bayesian game, we can start by talking about its properties. We will start by talking about whether or not first-price auctions are a DSIC mechanism. This property can be studied by using Myerson's lemma (Lemma 2.2.1) to prove that under certain conditions, the first-price auction is not DSIC, as proved in [3].

Corollary 2.2.2

Assuming quasilinear utility functions and private values for the agents, the first-price auction is not DSIC.

It is thus not a dominant strategy in the first-price auction to bid its true valuation. Even worse, while the first-price auction is quite popular in a number of settings, the computation of equilibrium is extremely hard in the general case, and require strong assumptions about the valuations of the players. For instance, *Roughgarden* has shown in [8] that in certain cases we could observe a phenomenon of *bid inversion* with respect to the valuations, where the bidders with the lowest valuation wins the auction in the equilibrium case (if the valuations are drawn from different distributions for instance). In the same paper, it has also been shown that the repeated first-price auction is not guaranteed to converge to an equilibrium.

2.2.2 Multiple-item auctions

Now that we have detailed the main properties of single-item auctions, we can start talking about multiple-item auctions. As shown by a paper on the subject by *Grece* [9], multiple-item auctions are the most common type of auctions on the internet today. This dominance of multiple-item auctions make sense; in fact on a typical website, we are going to have multiple ad slots and the website would prefer to sell all its ads at the same time in order to reduce the latency.
The two types of auctions that will be studied in this section are the General Second-Price auction and the Vickrey-Clarke-Groves auction.

General Second-Price auction

Let us first describe intuitively the mechanism behind the General Second-Price (GSP) auction, as defined in [5].

Definition 2.2.13: General Second-Price (GSP) auctions

Let there be M items sold and N bidders. Let us define the set of items as $O = \{1, ..., M\}$, the set of bidders as $E = \{1, ..., N\}$. and the set of bids as $B = \{b_1, ..., b_N\}$ (supposing the bids and the bidders are ranked in decreasing order of bids). Then the bidder i (when $i \leq M$) receives the item i at the price b_{i+1} and the bidders j (when $j > M$) receives nothing.

As a repeated Bayesian game, the GSP auction can be modelled in the following way.

Definition 2.2.14: GSP auctions as repeated Bayesian game

The GSP auction can be modelled as a repeated Bayesian game with a structure of the form:

$$\Gamma^r = (E, \Theta, (X_i, S_i, T_i, u_i, s_i, p_i)_{i \in E}, q),$$

where

- $E = \{1, ..., N\}$ is the total set of bidders to the auctions
- $\Theta = \{firstTurn, win_k(i), stop\}$ represent all the states where the world might be. The state $firstTurn$ represents the state where the game has just begun and therefore has no previous winner. The states $win_k(i)$ represent all the states corresponding to the fact that the bidder i has won the item k. Each of these state correspond to a combination of m winners
- $X_i = x_i \in (0, \infty)$ represent the set of moves that each player can play, which is the bid he makes in the auction
- $S_i = \{firstTurn, win(1), ..., win(M), lose, stop\}$ represent the signal send to each player after an auction, which is whether the game stops next turn or not and whether the player has won the ad slot k last turn
- $T_i = \{t_i^{(1)}, ..., t_i^{(M)}\} \in (0, \infty)$ represent the private information of player i which is the valuation he has for the items being sold
- u_i represent the utility function of each player and can be expressed by the following formula:
  $u_i(t_i, x_1, ..., x_N) = \sum_{k=1}^{M} (t_i^{(k)} - \max_{j \in U_{k+1}} x_j) \mathbb{1}\{x_i \in U_k\}$
  U_k is defined recursively in the following way:

Let us define $B_1 = E$, $U_1 = \{x \in B_1 : x \geq b \quad \forall b \in B_1\}$. Then you have $B_{i+1} = B_i \setminus U_i$ and $U_{i+1} = \{x \in B_{i+1} : x \geq b \quad \forall b \in B_{i+1}\}$

- s_i is the strategy of player i. Since we want to keep the model of the auction general, we are not going to make any particular assumption on this function
- p_i is the belief in the state of nature. If the state of nature is $firstTurn$ or *stop* then each player knows the state of nature with certainty, in other words $p_i(firstTurn|s_i = firstTurn) = 1$ or $p_i(stop|s_i = stop) = 1$ for all players. If a player loses then he does not know exactly which other players are winners. We do know as well than $p_i(win_k(i)|s_i = win(k)) = 1$
- q is the transition function. It can be expressed as:

$$q(x_1, ..., x_N, \theta) = \begin{cases} (firstTurn, ..., firstTurn, firstTurn) & \text{if first turn} \\ (stop, ..., stop, stop) & \text{if } \theta = stop \\ (\sum_{k=1}^{M} \mathbb{1}\{x_1 \in U_k\} win_k(1) & \\ +\mathbb{1}\{x_1 \in X \setminus \cup_{k=1}^{M} U_k\} lose, ... & \\ , \sum_{k=1}^{M} \mathbb{1}\{x_N \in U_k\} win_k(N) & \\ +\mathbb{1}\{x_N \in X \setminus \cup_{k=1}^{M} U_k\} lose, win_k(i)) & \text{else} \end{cases}$$

An important property of GSP auctions is that they are not DSIC, as proved in [3].

Corollary 2.2.3

Assuming quasilinear utility functions and private values for the agents, the GSP auction is not DSIC.

A consequence of this property is that it is not a dominant strategy to bid your true valuation for an item. As a consequence, the paper [10] by *Edelman et al.* has shown that the use of that system has lead to the emergence of bidding robots developed by advertisers in order to pay less for impressions, something that is detrimental for the revenues of the publishers. In the same paper, however, it was shown that the repeated GSP auction does have some interesting properties. In order to introduce these properties, we first need to introduce some new concepts. We start by introducing the concept of best-response as defined in [5].

Definition 2.2.15: Best-response

Consider a player $i \in E$ and his strategy s_i. We say that the strategy s_i is a best-response if it maximizes the player utility on the set of possible strategies.

Using the concept of best-response, we can then define the concept of of an ex-post Nash equilibrium, as defined in [5].

Definition 2.2.16: Ex-post Nash equilibrium

A profile of strategies $s_1, ..., s_N$ is an ex-post Nash equilibrium in a game with incomplete information if for every $t_1, ..., t_N$ we have that actions the actions $s_1(t_1), ..., s_N(t_N)$ are in Nash equilibrium in the full information game defined by the t_i. This requires that $s_i(t_i)$ is a best-response to $s_i(t_{-i})$ for every possible value of t_i.

In other words, this concept of equilibrium consists to say that if at each turn in a repeated auction the strategy of all the players is a best-response against the strategy of all the other players, then we can say that the repeated auction has reached an ex-post Nash equilibrium. The concept of ex-post Nash equilibrium has been studied multiple time for online ad-auctions, both for GSP auctions and for VCG auctions. In [10] by *Edelman et al.*, the authors have also proved that the GSP auctions converged to an ex-post Nash equilibrium under certain conditions and that the expected revenue of the seller in this equilibrium was at least as high as in the dominant-strategy equilibrium of the Vickrey-Clarke-Groves auction. All these properties make the GSP auction very interesting to the seller, which explains its popularity in the world of online ad auctions.

Vickrey-Clarke-Groves auction

Now that we have explained the mechanism behind the GSP auction, we are going to introduce the Vickrey-Clarke-Groves (VCG) auction. We are first going to describe it intuitively then describe it more formally as an infinitely repeated Bayesian game, using the definition given in [5].

Definition 2.2.17: Vickrey-Clarke-Groves (VCG) auctions

Let there be M items sold and N bidders. Let us define the set of items as $O = \{1, ..., M\}$ and the set of bidders as $E = \{1, ..., N\}$. Let U_E^O be the utility the group the group of bidders E has for the items O. Then the items are allocated in order to maximize the total utility of the bidders and the price for each bidder i who wins the item k is equal to $U_{E\setminus\{i\}}^{O} - U_{E\setminus\{i\}}^{O\setminus\{k\}}$

Since the utility of an item is very dependent on the type of item being sold, we will restrict ourselves to a model of VCG auctions for online ad auctions. In order to do so, we will thus need to first find a way to formally define the valuation of a bidder for an ad slot k. To do that, we will first introduce the concept of *Click-Through Rate* of an ad slot as defined by *Zhu et al.* in [11].

Definition 2.2.18: Click-through rate (CTR)

The click-through-rate (CTR) of an ad is the number of clicks for an ad normalized by the number of times this ad was printed. In this thesis, we will suppose that the CTR is only dependent on the *position* of the ad: we will write the CTR of an ad at position k as β_k.

This concept will thus be useful in order to model the valuation of a bidder for an ad slot. In this thesis, we will suppose that an advertiser has an intrinsic valuation for a visit on his website, that we will write as v. The valuation for ad slot will thus be dependent both on the intrinsic valuation and on the CTR of the ad slot. In other word, the valuation of an advertiser for ad slot k will be equal to $\bar{v}_k = \beta_k v$, and that valuation represents the utility obtained by the advertiser in this auction. We will explain further in this chapter how these parameters can be estimated.

Now that we have described these concepts, we can model the VCG auction as a repeated Bayesian game.

Definition 2.2.19: VCG auctions as repeated Bayesian game

The VCG auction can be modelled as a repeated Bayesian game with a structure of the form:

$$\Gamma^r = (E, \Theta, (X_i, S_i, T_i, u_i, s_i, p_i)_{i\in E}, q),$$

where

- $E = \{1, ..., N\}$ is the total set of bidders to the auctions
- $\Theta = \{firstTurn, win_k(i), stop\}$ represent all the states where the world might be. The state $firstTurn$ represents the state where the game has just begun and therefore has no previous winner. The states $win_k(i)$ represent all the states corresponding to the fact that the bidder i has won the item k. Each of these state correspond to a combination of m winners
- $X_i = x_i \in (0, \infty)$ represent the set of moves that each player can play, which is the bid he makes in the auction
- $S_i = \{firstTurn, win(1), ..., win(M), lose, stop\}$ represent the signal send to each player after an auction, which is whether the game stops next turn or not and whether the player has won the ad slot k last turn
- $T_i = \{t_i^{(1)}, ..., t_i^{(M)}\} \in (0, \infty)$ represent the private information of player i which is the valuation he has for the items being sold
- u_i represent the utility function of each player and can be expressed by the following formula:

$u_i(t_i, x_1, ..., x_N) = \sum_{k=1}^{M}(t_i^{(k)} - \sum_{j=k+1}^{M+1} \max_{l \in U_{j+1}} x_l(\beta_{j-1} - \beta_j))\mathbb{1}\{x_i \in U_k\}$, where β_j is a known parameter for ad slot j (see Definition 2.2.18).
U_k is defined recursively in the following way:
Let us define $B_1 = E$, $U_1 = \{x \in B_1 : x \geq b \quad \forall b \in B_1\}$. Then you have $B_{i+1} = B_i \setminus U_i$ and $U_{i+1} = \{x \in B_{i+1} : x \geq b \quad \forall b \in B_{i+1}\}$

- s_i is the strategy of player i. Since we want to keep the model of the auction general, we are not going to make any particular assumption on this function
- p_i is the belief in the state of nature. If the state of nature is $firstTurn$ or $stop$ then each player knows the state of nature with certainty, in other words $p_i(firstTurn|s_i = firstTurn) = 1$ or $p_i(stop|s_i = stop) = 1$ for all players. If a player loses then he does not know exactly which other players are winners. We do know as well than $p_i(win_k(i)|s_i = win(k)) = 1$
- q is the transition function. It can be expressed as:

$$q(x_1, ..., x_N, \theta) = \begin{cases} (firstTurn, ..., firstTurn, firstTurn) & \text{if first turn} \\ (stop, ..., stop, stop) & \text{if } \theta = stop \\ (\sum_{k=1}^{M} \mathbb{1}\{x_1 \in U_k\} win_k(1) & \\ +\mathbb{1}\{x_1 \in X \setminus \cup_{k=1}^{M} U_k\} lose, ... & \\ , \sum_{k=1}^{M} \mathbb{1}\{x_N \in U_k\} win_k(N) & \\ +\mathbb{1}\{x_N \in X \setminus \cup_{k=1}^{M} U_k\} lose, win_k(i)) & \text{else} \end{cases}$$

Now that we have modelled repeated VCG auction formally, we can list its properties. Historically, the VCG auction was designed in order to create a DSIC mechanism for multiple item auctions: it has thus the property of being DSIC under certain conditions, as proved in [3].

Corollary 2.2.4

Assuming quasilinear utility functions and private values for the agents, the VCG auction is DSIC.

Now that we have shown that the VCG auctions is a DSIC mechanism, we can study the equilibrium properties associated with it. In [12], *Holzman et al.* have shown that under certain conditions, an infinitely repeated VCG auction converged to an ex-post Nash equilibrium. *Myerson* has also shown in [13] that under certain conditions, the VCG auction with a reserve price was the optimal auction in terms of revenue maximization for the seller.

However, if these properties are not satisfied, VCG auctions can have some undesirable properties. In [14], *Ausubel et al.* have proved that under certain conditions, the VCG auction brings no revenue to the seller, which is obviously a big problem since the seller is generally the one designing the auction. In the same paper, it is also shown that the revenues are in general not monotonic and that the bidders could improve their wins by siding together against the seller. All these undesirable properties are the main reasons why the VCG auction is still so underused today: in fact, as shown by *Grece* in [9], there does not seem to exist a major ad publisher using VCG auctions for ad allocations outside of Facebook.

2.3 Platforms for online ad auctions

Now that we have described in detail the different auctions used, we are going to show in detail how these auctions are used in real online ad auctions. As different publishers might use different types of auctions and different billing methods, we will do a short review in this chapter on the different methods used in function of the advertiser.

2.3.1 Biggest publishers and multiple-item auctions

As we have discussed before, there exists a large number of different types of auctions and billing methods that can be used in order to rank the different bids made by the advertisers. While the system used by the advertisers might seem to be quite difficult to understand, there exists a lot of literature talking about the different pricing schemes used by the publishers and the publishers themselves are in general quite open about the type of auctions they use. In this section we will thus explain the type of auctions used by the biggest publishers on the internet.

We can first start by describing the market of online ad auctions by comparing their share of the market. Sadly, there does not seem to exist exact data on the subject, as it can be quite difficult to estimate. There however exists estimations of the share of net digital ad revenue by company, notably one done by *Grece* in [9]. The data obtained is presented in Figure 2.1.

Table 3 Net Digital Ad revenues worldwide, by company 2014-2015, in USD billion

	2014	2015	2016	Growth 2015/2016	Share of worldwide net digital ad revenues 2016
Google	46,13	53,05	57,8	9%	30,9%
Facebook	11,49	17,08	22,37	31%	12,0%
Baidu	6,77	8,97	11,46	28%	6,1%
Alibaba	6,05	8,12	10,99	35%	5,9%
Tencent	1,24	2,48	4,48	81%	2,4%
Yahoo	3,46	3,28	2,83	-14%	1,5%
Microsoft	2,11	2,55	2,94	15%	1,6%
Twitter	1,26	1,99	2,61	31%	1,4%
IAC	1,21	1,21	1,24	2%	0,7%
Verizon (AOL and Millenial Media)	1,25	1,33	1,39	5%	0,7%
Amazon	0,88	0,94	0,99	5%	0,5%
Sohu	0,8	0,97	1,16	20%	0,6%
Pandora	0,73	0,95	1,22	28%	0,7%
LinkedIn	0,75	0,96	1,06	10%	0,6%
Youku Tudou	0,57	0,76	1,03	36%	0,6%
Sina	0,6	0,67	0,74	10%	0,4%
Yelp	0,35	0,48	0,58	21%	0,3%
Other	47,68	53,54	61,92	16%	33,1%
Total digital	133,33	159,33	186,81	17%	

Source: eMarketer based on company reports March 2016

Note : includes advertising that appears on desktop and laptop computers as well as mobile phones, tablets and other internet-connected devices, and includes various formats of advertising on those platforms; net ad revenues after company pays traffic acquisition costs (TAC) to partner sites

Figure 2.1: Net digital ad revenues worldwide by company[9]

While we can see that there exists a clear domination at the top by Facebook and Google, there exist quite a lot of other ad companies with significant market shares. We can obviously not detail the auction mechanism for each website individually, but we are going to explain and source the auction mechanism used by all the companies with a share bigger than 1%.

The mechanism that is the most used by the biggest online publishers seems to be the GSP auction. In fact, this type of auction is being used by Google (see [15] by *McLaughlin et al.* and [16] by *Lucier et al.*), by Baidu (see [17] by *Chen et al.*), by Alibaba (see [18] by *Bai et al.*), by Tencent (for Soso.com) (see [19] by *Wang et al.*), by Yahoo! (see [20] by *Cary et al.*), Microsoft (for bing.com) (see [16] by *Lucier et al.*) and by Linkedin (see [21] by *Agarwal et al.*). This thus makes it the most used type of auction in the market of online ad auctions by far.

The other main auction used by online publishers is the VCG auction. While we have shown before that the VCG auction has a lot of interesting properties, it seems to be used only by one major publisher: Facebook (and Instagram by extension) (see [22] by *Varian et al.* and [23] by *Bjorner et al.*). Since Facebook has shown to be such a major online publisher, the VCG auction has still its importance in the landscape of online ad auctions, but it seems that it stays at the moment the only big publisher interested in the VCG auction.

Finally, we are going to talk about a last type of auction that seems to be used exclusively by one big publisher: the second-price auction. In fact, the second-price auction seems to be used exclusively by one big publisher: Twitter (see [24] by *Li et al.* and [25]). It thus seems from the documentation available that Twitter exclusively uses the second-price auction by permitting a bid per type of ad, and then allocating the bids using second-price auction rules. Twitter seems thus to be the only major publisher using single-item auctions for its ad allocation.

2.3.2 Ad exchanges and single-item auctions

We have now talked in detail about the functioning of ad auctions for major publishers, we can now start to talk about the functioning of ad auctions for *Ad exchanges* (AdX). As we have seen in Figure 2.1, even by taking into account all the major publishers of online ads, there is still a big share of online ad auctions done by other companies than the major publishers. Even if some of these ads are still done by other smaller publishers, an increasingly big amount is being sold in real time. We will thus detail the functioning of these ads sold in real time and describe the mechanism used by these small publishers.

We are first going to start by introducing the concept of *Ad Exchanges* or AdX. Ad Exchanges were studied by *Balseiro et al.* in [26], by *Sayedi* in [27] and by *Muthukrishnan* in [28]. AdX were described in [26] as online services that can be seen as spot markets for the real-time sale of online ad slots. In a generic Ad Exchange, publishers post an ad slot with a reservation price, advertisers post bids, and an auction is run; this happens between the time a user visits a page and the ad is displayed. Some examples of Ad Exchanges are the Rubicon Project Exchange or DoubleClick.

The concept of AdX is thus quite useful for advertisers who want to target specific individuals. The concept of AdX is intrinsically linked to the concept of *Real Time Bidding* or RTB. This concept was best described by *Sayedi* in [27], we will do here a quick summary of his paper. Historically, online ads were sold in *reservation contracts*, pre-negotiated contracts between publishers and advertisers where the price for each impression and the number of impressions were decided in advance in the contract and no dynamical changes were made in prices during the duration of the contract. While this approach might seem interesting, it lacks flexibility: neither publishers nor advertisers could react to changes in ad slots prices during the duration of the contract. This is what lead to the emergence of RTB. As the technology only improved in efficiency, it became more and more feasible for publishers to sell their leftover inventory to advertisers through AdX platforms in real time, and for the interested advertisers to use bots to be able to bid in real time: this is what is called Real Time Bidding. The method used was to send at each customer's visit of the web page the information to an AdX, to let the AdX do an auction in real time and then publish the ad on the website for the customer. This was very interesting for publishers, which could then easily sell their leftover inventories easily without having to look for pre-negotiated contracts.

The method was then extended to multiple AdX. The methodology was summarized by *Sayedi* in [27] and by *Qin et al.* in [29]. The two methods that are used are *waterfall auctions* and *header auctions*. In waterfall auctions, the publisher calls one ad-exchange site at a time, in descending order of importance, and accept the bid if it is higher than a certain threshold. In header auctions however, the publisher calls multiple AdX at the same time. Each AdX runs an auction and returns a clearing price. If the AdX replies with an offer quick enough, the offer is accepted and the publisher compares the clearing prices of the different accepted AdX. It can then decide to publish the ad proposed if the price is high enough or to show the ad of a reservation contract if that is not the case. The difference between the two systems is illustrated in Figure 2.2.

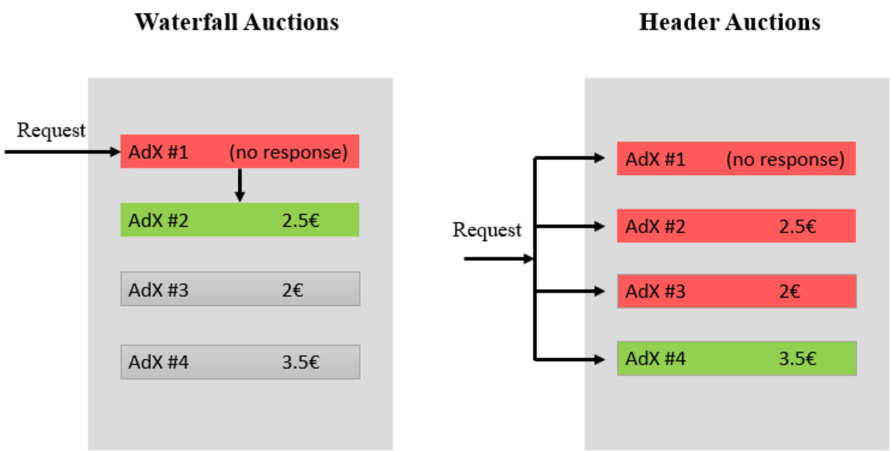

Figure 2.2: Comparaison between waterfall auctions and header auctions

Now that we have described how these exchanges work, we can talk about the type of auctions they use. As shown in [27] and in [29], AdX generally use single-item auctions. From a theoretical point of view, we can see that there are reasons why an AdX might want to use both types of single-item auctions: second-price auctions for the property of incentive compatibility and first-price auctions for their generally higher revenues. This shows us the importance of studying first- and second-price auctions in models for ad auctions.

2.4 Estimation of the valuations and of the click-through rate

While we have talked before about the valuations and click-through rates for each advertiser, we did not model in detail how these can be modelled and estimated. In fact, it is not always obvious for advertisers what exactly they should expect as CTR and how many of these clicks might lead to conversions, which make the evaluation of their valuations for an ad quite difficult. However, our algorithm requires to have at least an estimation of the different CTR and of the different valuations. This section is a quick description of the state of the art when it comes to estimation of the valuations and of the CTR.

A model of the CTR is given in [30] by *Richardson et al.*. This model was later improved by *Dave et al.* in [31] by using the fact that similar queries tend to follow similar CTR distributions. One final amelioration was made in [32] by *Kitts et al.* by using the fact that clicks decrease exponentially with position and thus that $\beta_k = \alpha e^{\gamma k}$, with α and γ parameters to be determined.

Once we have estimated the CTR, we can find the valuation using a simple formula. The general model to use is the following: let us write β_k for the average click-through rate of the ad k that we defined earlier (see Definition 2.2.18), $\overline{CVR}$ as the *conversion rate* of our website (the rate of people who buy something on our website compared to those who visit it) and $\overline{R}$ as the *average revenue* we make when one consumer buys something on our website. The model for the valuation for the ad slot k developed in [32] and by *Lee et al.* in [33] is the following:

$$\bar{v}_k = \beta_k \cdot \overline{CVR} \cdot \overline{R}$$

In the following chapters, we will use v to denote the intrinsic valuation behind an ad slot (in other words $v = \overline{CVR} \cdot \overline{R}$) and $\bar{v}_k$ to denote the real valuation for an ad slot, taking into account the CTR. In the case of a single-item auctions, $\bar{v}_1$ will simply be written as $\bar{v}$.

A method that can be used in order to find the conversion rate is described by *Lee et al.* in [33]. While the model is pretty simple, *Bax et al.* have shown in [34] that the main disadvantage of the method is the large expected error of the approximation, even for large samples. Another improved method can be found in the same paper. This method does not have an error as large as the method described earlier, but is way more costly in computational time, the choice between the two methods should be thus a compromise between the need for precision and the computational time required for the method.

3 Censored estimation of adverse bids

3.1 Behavior optimization in repeated games

The problem of behavior optimization in repeated games has been extensively studied in the literature. Some of the most used learning dynamics in repeated games are *best-response dynamics* and *no-regret dynamics*. We have already talked about the concept of best-response before (see Definition 2.2.15), we can therefore derive the concept of best-response dynamics, using the definition given in [5].

Definition 3.1.1: Best response dynamics

Best-response dynamics are procedures by which players search for an equilibrium of a game by deviating from the current strategy to an arbitrary beneficial deviation strategy until it converges to an equilibrium.

The other type of algorithms that help us to converge to an equilibrium are no-regret dynamics, using the definition in [35].

Definition 3.1.2: No-regret dynamics

No-regret dynamics are procedures by which players search for an equilibrium of a game by trying to minimize the difference between the actions made by the player and the single best action in retrospect.

These two methods of optimizing the behavior of a player in repeated games each have their advantages and problems. However, this thesis will be devoted to the development of a third method for behavior optimization in repeated games (and most specifically in repeated auctions) that we are going to call *bid estimation dynamics*. These dynamics are based on the following strategy: use probabilistic estimation techniques in order to obtain a probability distribution of the k^{th} highest adverse bid and then use these probability distributions in order to formulate an optimization problem. This is thus the method that we are going to describe in this thesis.

3.2 Hypotheses for the model

Since we want to develop a method for the estimation of the probability distributions of the highest adverse bids in repeated auctions, we are going to need to make some assumptions about the auctions we are taking part in. The assumptions we are going to make are the following:

1. Online auctions can be modelled as repeated Bayesian games
2. The highest adverse bids in the repeated auction are *stationary*: they are all realizations of a random variable with a probability distribution that *does not change over time*
3. In the case of multiple-item auctions, the highest adverse bids are either considered *independent* or we consider that we know the relation between the marginal distributions and the joint probability distribution

The assumptions we made here seem pretty restrictive and they are. We are here going to give some justifications on these hypotheses and show how they relate to the framework of Game Theory.

We start by talking about the assumption of modelling auctions as repeated Bayesian games. This is an hypothesis that is central to our model and central to Game Theory in general, however, one could argue that the Game Theoretical framework is too limited for online ad auctions. An example of its limitations is the assumption that the number of players is *constant over time* (see Definition 2.1.5). This is extremely debatable in the case of repeated ad auctions: since the entry costs are low and the number of internet users is growing over time, the number of advertisers in an auction will then logically increase with the number of consumers. The assumption is however central to Game Theory: to our knowledge, there is not any paper existing in Game Theory where the number of players is not constant. Therefore, since the goal of this thesis is to analyze ad auctions in a Game Theoretical framework, we will consider that the assumptions made in Definition 2.1.5 are satisfied and that online ad auctions can be modelled as a repeated Bayesian game.
The second assumption is the assumption about stationary probability distributions. This assumption is also pretty restrictive, but we believe it is a reasonable

assumption to make in the context of online ad auctions. In fact, as we have studied before in Chapter 2, Section 2.2, repeated online ad auctions converge to equilibrium prices in a lot of cases. As some papers have shown that this is however not always the case (see [8] for instance), we believe that a natural extension of this assumption is to assume that the highest adverse bids in the auction are a sample of an unknown random variable and thus that the equilibrium prices are also samples from random variables, since they are a direct function of these highest adverse bids. This assumption is going to be discussed later and we will explain in the conclusion of this thesis how this assumption could possibly be relaxed to a less strict assumption on the bidding behavior of the adverse advertisers.
Finally, we can talk about the last assumption, which is also pretty debatable. In fact, in the general case, it would make a lot of sense for the highest adverse bids to be correlated with each other, as the behavior of one bidder would probably be influenced by the behavior of the others. The reason behind this assumption is however quite simple: to our knowledge, while there exists multiple paper on the subject of censored estimators for dependent random variables (see [36] by *Van Keilegom*, [37] by *Huang et al.* and [38] by *Uña Alvarez et. al.*), the methods developed are not general enough for our case, where we need to find a general estimator for the joint distribution of interval-censored observations. As the main objective of this thesis was not the development of a new estimator in the general case, it seemed to make sense to restrict ourselves to the case of independent highest adverse bids.

Now that we have talked about our hypotheses, we can show a special case where both the second and third assumptions are satisfied. This is the case where the bids made by an advertiser are all realizations of a random variable. In order to show that, we are going to talk about a specific field of statistics, the field of *order statistics*. We use here the definition given by *Bapat et al.* in [39].

Definition 3.2.1: Order statistics

Let $X_1, ..., X_N$ be N independent random variables. Then we can write the order statistics of X as Y_i, with $Y_1 \leq Y_2 \leq ... \leq Y_N$. In other words, the n^{th} order statistic of $X_1, ..., X_N$ is n^{th} minimum value of $X_1, ..., X_N$.

An important theorem of order statistics is their independence for independent random variables, as proved in [39].

Theorem 3.2.1: Independence of order statistics

Let $X_1, ..., X_N$ be N independent random variables and $Y_1, ..., Y_N$ be the N associated order statistics. Then $Y_1, ..., Y_N$ are independent random variables.

This property of order statistics has some important consequences: if the bids made by all advertisers are all realizations of random variables, then all our assumptions are satisfied. We are not going to suppose that this assumption is always satisfied in the development of our algorithm; however, this case represents an example of a case where our assumptions are satisfied.

This is thus our needed hypotheses about the highest adverse bids in the auction. We also need to make hypotheses about our knowledge of the auction we are taking part in:

- We have knowledge about the type of auction used for selling the ad slots (see Chapter 2, Section 2.2)
- At each auction we have taken part in, we have knowledge of the previous bids we made, whether we won the auction or not and of the price we paid for each previous auction won
- We have an estimation of the click-through rate of each ad slot and of our valuation for a given ad slot (see Chapter 2, Section 2.4)

3.3 Censored estimators

In order to simplify the notations, we will use starting from this section some abbreviations. Let us suppose we have a random variable X. We will write the *cumulative distribution function* (cdf) of X as $F(x) = P[X \leq x]$ and the *probability distribution function* (pdf) of X as $f(x) = \frac{dF(x)}{dx}$.

As we have explained in the previous section, we will suppose in this thesis that the highest adverse bids in each auction are all realizations of the same random variables. However, we clearly do not observe directly the highest adverse bids in each auction, instead we observe a *censored realization* of a random variable. The concept of censored realization of a random variable is defined by *Hainaut* in [40] in the following way.

Definition 3.3.1: Censored data

A sample of a random variable T is said to be right-censored if instead of observing a realization T_i we observe the couple (X_i, D_i), where $X_i = \inf(T_i, C_i)$ and $D_i = \mathbb{1}\{T_i \leq C_i\}$, where C_i is a constant or a positive random variable.
A sample is right censored if we observe the couple (X_i, D_i), with $X_i = \sup(T_i, C_i)$ and $D_i = \mathbb{1}\{T_i \geq C_i\}$.
A sample is interval censored if we observe the couple (L_i, R_i), where $T_i \in [L_i, R_i]$.

This concept is thus going to be useful in the context of highest adverse bid estimation for auctions since we have in general *censored observations* of the highest adverse bids in an auction. Let us describe for instance the example of second-price auctions and the problem of the estimation of the distribution of the highest adverse bids of the auction. We remember by our hypothesis that the adverse bids in an auction are viewed as random variables. We write X for the highest adverse bid, and each auction represents a (possibly censored) observation of this random variable.

- If we win the auction, we pay the highest adverse bid. This is thus an *uncensored observation* of the highest adverse bid

- If we lose the auction, we do not know the highest adverse bid, but we know that it is higher than the bid we have made. This is thus a *right-censored observation* of the highest adverse bid

We thus have possibly censored observations of the highest adverse bid and we can then use different algorithms in order to estimate the distribution of these adverse bids. This method can also be extended to other type of auctions. Let us start by extending the concept to first-price auctions. We thus have the following observations of this random variable.

- If we win the auction, we know that the bid made by the second highest bidder is necessarily lower than our bid, but higher than 0. This is thus an *interval censored observation* of the highest adverse bid

- If we lose the auction, we do not know the highest bid made by an adverse bidder, but we know that it is higher than the bid we have made. This is thus a *right-censored observation* of the highest adverse bid

This can also be extended to multiple-item auctions. Let us suppose that we have auctions with M items. Instead of considering that we have only one random

variable X to estimate, we suppose that we have M random variables to estimate $X^{(1)}, ..., X^{(M)}$, where each random variable $X^{(k)}$ represents the k^{th} highest adverse bid. We can now talk about the observation we have of these random variables at each auctions.
Let us start by talking about the GSP auction. We have the following observations of the random variable $X^{(k)}$

- Every time we win a slot higher than the slot k, we know that the k^{th} highest adverse bid is lower than the price we paid, but higher than 0. This is thus an *interval censored observation* of the k^{th} highest adverse bid
- Every time we win the slot k, we have a realization for the k^{th} highest adverse bid. This is thus an *uncensored observation* of the k^{th} highest adverse bid
- Every time we win a slot lower or we win nothing at all, we know that the k^{th} highest adverse bid is higher. This is thus an *right-censored observation* of the k^{th} highest adverse bid

We can do the same for VCG auctions. However, thanks to the formula for the clearing price of the VCG auctions, we can have a bit of a more precise estimation. Let us use the same notations as before and let us denote the random variables denoting the M highest adverse bids as $X^{(1)}, ..., X^{(M)}$. The variable we are observing in this case is a linear combination of the highest adverse bids. Because of the definition of VCG auctions for ad slots (see Definition 2.2.19), we have that we observe the random variable $P^{(k)}$, which is equal to:

$$P^{(k)} = \sum_{j=k}^{M} X^{(j)}(\beta_j - \beta_{j+1})$$

This random variable can be interpreted as the clearing price of the slots behind the slot we have won, considering we have participated in the auction. We can then use this formula to deduce information about the random variable $P^{(k)}$. We are going to subdivide the information we have of this random variable in function of the different cases.
Let us suppose we have won the slot i in the auction. We then have an exact realization of the random variable $P^{(i)}$. We can then also use this realization in order to gain information about all the $P^{(k)}$, for $k \neq i$.
Let us start by studying the case $i < k$. By the definition of $P^{(i)}$, this is given by:

$$P^{(i)} = (\beta_i - \beta_{i+1})X^{(i)} + ... + (\beta_{k-1} - \beta_k)X^{(k-1)} + P^{(k)}$$

We clearly do not know the exact value of $X^{(i)}, ..., X^{(k-1)}$. However, we do know that $X^{(i)}, ..., X^{(k-1)} \geq b$. We can use that in the following way:

$$\begin{aligned}
P^{(k)} &= P^{(i)} - (\beta_i - \beta_{i+1})X^{(i)} - ... - (\beta_{k-1} - \beta_k)X^{(k-1)} \\
&\geq P^{(i)} - (\beta_i - \beta_{i+1})b - ... - (\beta_{k-1} - \beta_k)b \\
&\geq P^{(i)} - (\beta_i - \beta_{i+1} + \beta_{i+1} - ... + \beta_{k-1} - \beta_k)b \\
&\geq P^{(i)} - (\beta_i - \beta_k)b
\end{aligned}$$

We have as well that $P^{(k)} \leq P^{(i)}$ and we thus know that the realization of $P^{(k)}$ will be in the interval $[P^{(i)} - (\beta_i - \beta_k)b; P^{(i)}]$.
We can do the same in the case where $i > k$. By definition we have that $P^{(k)} = (\beta_k - \beta_{k+1})X^{(k)} + ... + (\beta_{i-1} - \beta_i)X^{(i-1)} + P^{(i)}$. We know as well that $X^{(k)}, ..., X^{(i-1)} \geq b$. We thus have that:

$$\begin{aligned}
P^{(k)} &= (\beta_k - \beta_{k+1})X^{(k)} + ... + (\beta_{i-1} - \beta_i)X^{(i-1)} + P^{(i)} \\
&\geq (\beta_k - \beta_{k+1})b + ... + (\beta_{i-1} - \beta_i)b + P^{(i)} \\
&\geq (\beta_k - \beta_{k+1} + ... + \beta_{i-1} - \beta_i)b + P^{(i)} \\
&\geq (\beta_k - \beta_i)b + P^{(i)}
\end{aligned}$$

We thus know in this case that $P^{(k)}$ is in the interval $[(\beta_k - \beta_i)b + P^{(i)}, +\infty[$ and we can use this information to improve our knowledge of the clearing price. Our full knowledge of $P^{(k)}$ is thus the following.

- Every time we win a slot i higher than the slot k, we know that $P^{(k)}$ is in the interval $[P^{(i)} - (\beta_i - \beta_k)b; P^{(i)}]$. This is thus an *interval censored observation* of the k^{th} highest clearing price
- Every time we win the slot k, we have a realization of $P^{(k)}$. This is thus an *uncensored observation* of the k^{th} highest clearing price
- Every time we win a slot i lower than the slot k or we win nothing at all ($P^{(i)} = 0$), we know that $P^{(k)}$ is higher than $(\beta_k - \beta_i)b + P^{(i)}$. This is thus a *right censored observation* of the k^{th} highest clearing price

Once we have managed to estimate the distributions of $P^{(1)}, ..., P^{(M)}$, we can use them to estimate distributions of $X^{(1)}, ..., X^{(M)}$. This can be done by using the following recurrence relation:

$$\begin{cases}
X^{(M)} = \frac{P^{(M)}}{\beta_M} \\
X^{(k)} = \frac{P^{(k)} - P^{(k+1)}}{\beta_k - \beta_{k+1}} \quad \forall k \in \{1, ..., M-1\}
\end{cases}$$

We have thus fully described the observations we make at each auction and how they relate to the random variables we are trying to estimate.

3.4 Parametric censored estimation

3.4.1 Maximum likelihood estimation

Now that we have described how the problem of censored estimation of a random variable relates to auction theory, we can explain different methods that can be used in order to estimate the distribution of random variables using censored data. One of the ways to do that is by using *parametric maximum likelihood estimation*. We are first going to introduce the concept of *parametric models* as defined by *Zhang* in [41].

Definition 3.4.1: Parametric models

Let us suppose we have a random variable X with a pdf $f(x)$ and a cdf $F(x)$. We say that the pdf and cdf of X follow a parametric model if there exists a vector of parameters $\theta \in \mathbb{R}^n$ such that $f(x)$ and $F(x)$ are fully determined by θ.

A good example of a random variable following a parametric model is a normally distributed random variable: the pdf and the cdf of X are fully determined by the two parameters of the normal distribution, which are the mean μ and the standard deviation σ.

Now that we have described the concept of parametric models, we can describe a common method of estimating these parameters given a set of data: *maximum likelihood estimation*. The method was summarized by *Myung* in [42].

Definition 3.4.2: Maximum likelihood estimation

Maximum likelihood estimation is a method of determining the parameters of a parametric model using samples from the distribution and knowledge about the pdf and cdf of the random variable to estimate. It consists of determining a *likelihood function* using the different samples and of maximizing this function in function of the parameter θ defining the parametric model.

The likelihood function of the distribution is dependent on the form of the samples. In the case of general censored data, we can find the likelihood function by using the definition by *Zhang* in [41].

Definition 3.4.3: General likelihood function for censored data

Let us denote by D the set of direct observations, L the set of left censored observations, R the set of right censored observations and I the set of interval censored observations (where each observation is in $[L_i, R_i] \quad \forall i \in I$). Let also suppose that the random variable studied a parametric model of parameter θ, with pdf $f_\theta(x)$, cdf $F_\theta(x)$ and survival function $S_\theta(x)$.
Then the general likelihood function for parametric estimation of censored data is given by

$$L(\theta) = \prod_{d \in D} f_\theta(x_d) \prod_{r \in R} S_\theta(x_r) \prod_{l \in L} F_\theta(x_l) \prod_{i \in I} [S_\theta(L_i) - S_\theta(R_i)]$$

We can thus see that all the specific cases of MLE estimation for the different type of auctions are all just specific cases of this formula; we can thus detail the different log-likelihood functions to maximize for each type of auction.

Let us start by talking about single-item auctions. Let us suppose we have already played in T sequential single-item auctions of the same type. We will define three variables related to our data for each auction $t \in \{1, ..., T\}$: u_t, representing whether or not we won auction t, p_t representing the price paid at auction t and b_t the bid made at auction t. We can then give the two likelihood functions for first- and second-price auctions.
For second-price auctions we have:

$$L(\theta) = \prod_{t=1}^{T} (f_\theta(p_t))^{u_t} (S_\theta(b_t))^{1-u_t}$$

For first-price auctions we have:

$$L(\theta) = \prod_{t=1}^{T} (F_\theta(p_t) - F_\theta(0))^{u_t} (S_\theta(b_t))^{1-u_t}$$

We can then do the same for multiple-item auctions. We suppose we play in T sequential multiple-item auctions of the same type with M items. We denote by $u_{t,k}$ a binary variable representing whether we won item k in auction t and $v_{t,k}$ a binary variable representing whether we won an item better than k (so one of the items in $\{1, ..., k-1\}$). We also write down as $f_\theta^{(k)}(x)$ the pdf of the random variable on the k^{th} highest slot, $F_\theta^{(k)}(x)$ the cdf of this random variable and $S_\theta^{(k)}(x)$ the survival function.

We can then define a likelihood function for each of the M items for GSP auctions. It is given by:

$$L(\theta)^{(k)} = \prod_{t=1}^{T}(f_\theta^{(k)}(p_t))^{u_{t,k}}(F_\theta^{(k)}(p_t) - F_\theta^{(k)}(0))^{v_{t,k}}(S_\theta^{(k)}(b_t))^{1-u_{t,k}-v_{t,k}} \quad k \in \{1, .., M\}$$

In the case of VCG auctions, the method is similar. The general likelihood function in this case is given by:

$$L(\theta)^{(k)} = \prod_{t=1}^{T}(f_\theta^{(k)}(p_t))^{u_{t,k}}(F_\theta^{(k)}(p_t) - F_\theta^{(k)}(p_t - (\beta_i - \beta_k)b_t))^{v_{t,k}}$$
$$(S_\theta^{(k)}((\beta_k - \beta_i)b_t + p_t))^{1-u_{t,k}-v_{t,k}}$$

for $k \in \{1, .., M\}$, where i represent the position of the ad slot won in the auction.

We have thus fully defined our algorithm for parametric censored estimation applied to auctions.

3.4.2 Choice of the parametric model

While we have now detailed the functioning of the maximum likelihood estimation for the censored parametric case, we can see that we have until now assumed that we knew in advance the parametric model that was going to be used. In reality, that is in general not the case, since we do not assume we have a priori information on the model that should be used. What we can do however is maximizing the log-likelihood for different parametric models and then derive the best parametric model to use using concepts derived from information theory in order to be able to compare the goodness-of-fit.

Generally, in order to check whether a model is a good fit using maximum likelihood estimation, we can use criteria such as the *Bayesian Information Criterion* or BIC. The BIC was introduced by *Shwarz* in [43] and is defined in the following way.

Definition 3.4.4: Bayesian Information Criterion (BIC)

Let $\theta \in \mathbb{R}^k$ be a vector of size k representing the parameter to estimate, θ^* be the optimal parameter obtained by maximizing the likelihood function, $l(x)$ be the log of the likelihood function and T the number of observations. Then the BIC of the model is equal to:

$$BIC = k \log(T) - 2l(\theta^*)$$

Shwarz has proven in [43] that the BIC has the advantage of being consistent under certain conditions, which means that as T becomes large, with probability approaching 1, choosing the model with the lowest BIC will give us the best model. However, while the consistency of the BIC is proved for non-censored models, the consistency is way more debatable for censored models. The consistency of BIC has been extensively studied for survival analysis, notably by *Volinsky* in [44] and by *Dirick et al.* in [45]. However, the conclusion drawn in these two papers cannot be directly applied in our case, since they rely on the knowledge of the size of the population studied at a given time, a parameter that cannot be directly found in our problem. The conclusion of *Volinksy* in [44] was however that even if the convergence results only holds using the size of the population studied, using T in the BIC model gives relatively good results. We are thus going to use this hypothesis in our model for the estimation of the distribution of the adverse bids. Our algorithm is thus the following:

Algorithm 1: Parametric censored estimation algorithm

Data: p,b,u

Result: Best parameter choice θ^* and best parametric model associated to it

Optimize the log-likelihood for a number of parametric distributions using censored data.

Compute the BIC for each distribution studied.

Choose the distribution with the lowest BIC and return the parameter θ^*.

3.4.3 Hypothesis testing

While we have now found a way to deduce the best fit of distribution for a parametric estimator, nothing is guaranteeing that the best parametric fit for our distribution is in fact good enough. For instance, we might have a random variable with a probability distribution that cannot be fitted using classical parametric models. What we are going to do now is to try to test whether our sample is in fact generated for a given distribution using the Anderson-Darling test, as introduced by *Anderson et al.* in [46].

Definition 3.4.5: Anderson-Darling test

Let us have the sample $Y_1 < ... < Y_n$. The Anderson-Darling test check whether the sample Y comes from the cdf F by studying the test statistic A that is defined by $A^2 = -n - S$ where S is given by

$$S = \sum_{i=1}^{n} \frac{2i-1}{n} (\log(F(Y_i)) - \log(F(Y_{n+1-i})))$$

The test statistic can then be compared to the critical value for the specific distribution studied.

We will thus use this statistic to test whether our fit is good using a significance level of 5%, since it seems to be the standard significance level for such goodness-of-fit test that was used in [46]. The logic will be the following: if the fit is good enough to reject the null hypothesis at our significance level, we consider that we have found a good enough parametric fit for our distribution. If the fit is not good enough, we continue further by using nonparametric estimators, as we are going to detail in the following section.

While this method is interesting and easy to understand, we remember that we do not have direct observations of our data but that our observations are censored. In order to circumvent this problem, we are going to use a simple method inspired by the method of *Alharpy et al.* in [47]. In order to understand the method we are going to need to introduce the concept of *truncated distributions*, as defined by *Greene* in [48].

Definition 3.4.6: Truncated distributions

Let us define a general random variable X with cdf $F(x)$ and pdf $f(x)$. Then the truncated distribution of X to $[a, b]$ is the part of the distribution of X that belongs to the interval $[a, b]$. The truncated pdf and cdf of X to $[a, b]$ are written in the following way:

$$F(x|a \leq X \leq b) = \begin{cases} 0 & \text{if } x < a \\ \frac{F(x)-F(a)}{F(b)-F(a)} & \text{if } a \leq x \leq b \\ 1 & \text{if } x > b \end{cases}$$

$$f(x|a \leq X \leq b) = \begin{cases} 0 & \text{if } x < a \\ \frac{f(x)}{F(b)-F(a)} & \text{if } a \leq x \leq b \\ 0 & \text{if } x > b \end{cases}$$

Now that we have described the concept, we can describe our method: we simulate a sample of our distribution using the distribution estimated by log-likelihood truncated by L_i and R_i. Our algorithm is thus the following:

Algorithm 2: Parametric censored hypothesis testing

Data: θ^*: distribution with lowest BIC from before
$\hat{F}$: optimal distribution associated to θ^*
L, R: vectors of size n of the left and right bounds of the samples
X: exact samples
Result: Accept/Reject parametric distribution
for $i \in \{1, ..., n\}$ **do**
  if $\hat{F}(L_i) > \hat{F}(R_i)$ **then**
    Reject parametric distribution: impossible that sample in interval
  else
    Generate Y_i from $\hat{F}$ truncated to $[L_i, R_i]$
$X = X \cup Y$
Use Anderson-Darling hypothesis test on X using $\hat{F}$

3.5 Nonparametric censored estimation

Now that we have talked about the concept of censored estimation for parametric models, we remember that the model has its limits. In fact, while we can use log-likelihood maximization techniques in order to try to best fit censored data to a specific family of parametric distribution, we have to compute one likelihood function for each type of parametric distribution. However the number of types of parametric distribution is infinite and the computation of the parametric maximum likelihood estimator can be costly in computational time. Therefore, what we are going to do in this case is to try to maximize the likelihood function for a number of common distributions and then move on to *nonparametric estimators*, estimators that do not require any knowledge about the distribution of the data. We are now going to detail the methods to be used in this context.

3.5.1 Right-censored estimation and the Kaplan-Meier estimator

It is first important to realize that we have two types of censored observations in the different types of auctions studied:

- In the case of second-price auctions we only have uncensored and right-censored observations
- In all the other cases we have a mix of uncensored, right-censored and interval-censored observations

We are thus going to study the two different cases independently as different techniques might be used for the different types of censorship. We are going to start by studying second-price auctions, where the observations of the data are either uncensored or right-censored. In that case, we can use a well-known type of estimator: the Kaplan-Meier estimator. The definition we are going to use is the definition given by *Van Keilegom* in [36], summarizing the work of *Kaplan et al.* in [49].

Definition 3.5.1: Kaplan-Meier estimator

Let us suppose we have $Y_1, ..., Y_n$ i.i.d. distributed non-negative random variables with unknown cdf $F(x)$. Let us suppose that the observations of the random variables are possibly right-censored by C_i, a constant. In other words, instead of observing realizations of Y we observe the pair (T_i, Δ_i), where $T_i = \min(Y_i, C_i)$ and $\Delta_i = \mathbb{1}\{Y_i \leq C_i\}$. Then the Kaplan-Meier estimator of the cdf of Y is given by:

$$F(t) = 1 - \left\{\Pi_{T_{(i)} \leq t}\left(1 - \frac{1}{n - i + 1}\right)^{\Delta_{(i)}}\right\}\mathbb{1}\{t \leq T_{(n)}\}$$

where $T_{(1)}, ..., T_{(n)}$ are the ordered T_i and $\Delta_{(1)}, ..., \Delta_{(n)}$ are the corresponding indicators.

This estimator can thus directly be applied in the case of non-parametric estimation of the highest adverse bid for second-price auctions. In order to do that we are going to need to use the variable m_t defined in the following way:

$$m_t = \begin{cases} p_t & \text{if we win the auction} \\ b_t & \text{otherwise} \end{cases}$$

Let us also suppose we have the m_t ranked in increasing order such that $m_{(1)} \leq m_{(2)} \leq ... \leq m_{(T)}$. Then the expression of the Kaplan-Meier estimator in the case of second-price auctions is given by:

$$F_n(p) = 1 - \left\{\Pi_{m_{(t)} \leq p}\left(1 - \frac{1}{T - p + 1}\right)^{u_t}\right\}\mathbb{1}\{p \leq m_{(T)}\}$$

An important remark made in [36] is that the Kaplan-Meier estimator is a step function and is thus not differentiable. In certain cases it can be problematic and it can be needed to have to smooth it in order to have a multiple time differentiable cdf. We are here going to present a method for smoothing the cdf obtained by the Kaplan-Meier estimator using *Kernel Density Estimation*. A method for Kernel Density Estimation for cdf was introduced by *Baszczyńska* in [50].

Definition 3.5.2: Kernel Density Estimation

Let us suppose we have a sample $X_1, ..., X_n$ of a random variable X and that we want to use this sample in order to find a smooth estimate of the cdf of X. Then a method to estimate a smooth cdf for X is to write the cdf of X as:

$$F(x) = \frac{1}{N} \sum_{n=1}^{N} W(\frac{x - X_i}{h})$$

where h is a smoothing parameter and W is called a kernel function.

This technique is thus interesting because it can be applied to the Kaplan-Meier estimator in order to smooth the cdf. The method we are going to use is the following:

- Use the Kaplan-Meier estimator to have an estimation of the cdf of the random variable
- Use this estimator to compute the value of the cdf at certain chosen points
- Smooth the cdf using Kernel Density Estimation methods (if needed)

An important property of Kernel Density Estimation methods is that the form of the cdf is heavily dependent on the type of kernel used. The most frequently used kernel function is the Gaussian kernel (see [50]), this is thus also the kernel that we are going to use in this case. This also means that our smoothing procedure will be way more precise for "normal-like" distributions: if other types of distributions are needed for smoothing, the kernel function should be adapted in consequence. We have thus obtained a full procedure in order to obtain a smooth estimation of the cdf of the highest adverse bid.

3.5.2 Interval-censored estimation and the Turnbull estimator

Now that we have described the functioning of the Kaplan-Meier estimator, we remember that it is only applicable to right-censored observations. However, as we have described before, in some type of auctions we have interval-censored observations and the Kaplan-Meier estimator cannot be used anymore. For this type of cases, we are going to need to describe a new type of estimator. This estimator can be seen as the application of the *Expectation Maximization algorithm* (an algorithm used for parametric estimation) to interval censored data and was first described by *Turnbull*, in two papers on the subjects [51][52]. For this reason, we will call this estimator the *Turnbull estimator*. We are now going to briefly describe the algorithm behind this estimator.

The first idea needed to describe Turnbull's algorithm is the concept of *Turnbull's intervals*. In [53], *Fay* summarized the work done by *Turnbull* in [51] and defined Turnbull's intervals in the following way.

> **Definition 3.5.3: Turnbull's intervals**
>
> Let us define as $(L_i, R_i]$ the observed censored interval for the i^{th} auction. Let us define as well $I = \cup_{i=1}^{n} L_i \cup_{i=1}^{n} R_i$ and m as the number of elements in I. Then Turnbull's intervals are the intervals $(t_0, t_1], (t_1, t_2], ..., (t_m, t_{m+1}]$ such that $t_0 = 0$, $t_{m+1} = +\infty$ and $\forall i, t_i \in I$, with $t_i \leq t_{i+1}$

Now that we have defined the concept of Turnbull's intervals, we can describe how these intervals can be useful for an estimator. In [51], *Turnbull* has proved that the boundaries of Turnbull's intervals are the only points where the density function is non equal to 0. As such, we only have to find the probabilities associated to the $m+1$ boundary points of the interval. In other words, we are trying to find a *vector* of probabilities $p = [p(t_1), p(t_2), ..., p(t_{m+1})]$. Turnbull's algorithm was best described by *Fay* in [53] in the following way.

Algorithm 3: Turnbull's algorithm

Data: n: the numbers of observations
$(L_i, R_i]$: the censored intervals for each observation
I: the set of Turnbull intervals
$t_0, ..., t_{m+1}$: all the points in I
Result: $\hat{p}^*$
$\hat{p} = [\frac{1}{m+1}, ..., \frac{1}{m+1}]$
while *Not converged* **do**
  for $j \in \{1, ..., m+1\}$ **do**
    for $i \in \{1, ..., n\}$ **do**
      $\hat{p}_i(t_j) = \frac{\hat{p}(t_j)\mathbb{1}\{t_j \in (L_i, R_i]\}}{\sum_{t_k \in (L_i, R_i]} \hat{p}(t_k)}$
    $\hat{p}(t_j) = \frac{\sum_{i=1}^{n} \hat{p}_i(t_j)}{n}$
$\hat{p}^* = \hat{p}$

This is thus an algorithm for an estimation of the pdf of X at certain given points $t_0, ..., t_m$. The cdf can then be obtained by simply integrating this pdf(and then obtaining a step function). *Gentleman et al.* have then proven in [54] that the empirical distribution function obtained using this algorithm is *strongly consistent*

and converges to the true distribution of the observations. This algorithm was later implemented by *Fay et al.* in R in the interval package: this is the implementation we used for our simulations (see [55]).
Now that we have described Turnbull's algorithm, we are going to describe how we use it in our case. The method is the same as the one described in the previous section: describe the left and right bounds for each observation using our knowledge about the auction and then run Turnbull's algorithm using these left and right bounds.
Once we have obtained this estimator, we can see that it is once again not smooth. If needed, depending on the application, it can be smoothed using Kernel Density Estimation.

3.6 Log-concave interval-censored estimation and the logconcens algorithm

We have now covered all the consistent nonparametric estimation methods that could be used in order to estimate the distributions of a random variable using censored estimation. In certain cases, however, it might be useful not to obtain the most precise estimator but to obtain a slightly less precise estimator that satisfies a certain property, in order to be able to formulate a problem having more interesting properties. In our case, an interesting property that might be used is the property of *log-concavity*. We are first going to introduce this property and then later detail how it might be used in the context of *probabilistic optimization*.

We can first start by defining the concept of a log-concave function, as defined by *Wellner* in [56].

Definition 3.6.1: Log-concavity of a function

A function $f : \mathbb{R} \to \mathbb{R}^+$ is log-concave if $\forall t \in [0, 1], \forall x \in dom(f)$, we have:

$$f(tx + (1 - t)y) \geq f(x)^t f(y)^{1-t}$$

Related to this concept, we can introduce the concept of a *log-concave random variable*, using the definition in [56].

Definition 3.6.2: Log-concave random variable

A log-concave random variable is a random variable where the pdf is a log-concave function.

A direct consequence of the log-concavity of the pdf of a random variable, as shown in [56], is that the a log-concave random variable also has a log-concave cdf, since the primitive of a log-concave function is also log-concave.

One of the reasons why log-concave random variables are an interesting concept is their applications in *probabilistic optimization* as a consequence of *Prékopa's Theorem*. In order to introduce this theorem, we will first introduce a related concept: the concept of quasi-concave functions, as defined by *Stoer et al.* in [57].

Definition 3.6.3: Quasi-concave functions

A function $f : C \to \mathbb{R}$ with C convex is quasi-concave if

$$\forall x, y \in C, \forall z \in [x, y] : f(z) \geq \min(f(x), f(y))$$

We can now introduce Prékopa's Theorem, as proved by *Prékopa* in [58].

Theorem 3.6.1: Prékopa's Theorem

Let $X_1, ..., X_n$ be continuous random variables with log-concave distributions, $g : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ be a quasi-concave function and p be a known probability. Then the set C defined by $P[g(X_1, ..., X_n, y) \geq 0] \geq p \quad \forall y \in C$ is a closed convex set.

As we are going to talk about later in this thesis, the property of being closed and convex for a set is quite useful in the context of optimization. We will come back to this property later.

Now that we have described the concept of log-concave random variables and their properties, we are going to study the problem of censored estimation of the distribution of a random using log-concave pdf and cdf. The algorithm that we will hereby describe to use an algorithm using nonparametric maximum likelihood estimation, and was introduced by *Dümbgen et al.* in [59]. Since this algorithm was later implemented in R using the name *logconcens* by *Schumacher et al.* (see [60]), this is the name we are going to use in order to describe this algorithm.

To understand the mechanism behind the algorithm, we will first need to introduce a likelihood function for nonparametric estimation, using the definitions given in [59]. If the observations were not censored and we suppose that there is a probability q that the random variable X_i is equal to $+\infty$, an appropriate log-likelihood function is given by: [59]

$$l(\phi, q) = \frac{1}{n}\sum_{i=1}^{n}(\mathbb{1}\{X_i < \infty\}\phi(X_i) + (\mathbb{1}\{X_i = \infty\}\log(q))$$

where $\phi(x)$ is such that the pdf of X is given by $f(x) = e^{\phi(x)}$.
Using that, we can derive an appropriate log-likelihood for censored observations:[59]

$$l(\phi, q) = \frac{1}{n}\sum_{i=1}^{n}(\mathbb{1}\{L_i = R_i\}\phi(X_i) + \mathbb{1}\{L_i < R_i\}\log(P_{\phi,q}([L_i, R_i])))$$

The goal is thus going to be to maximize this function on the sets of concave functions. In order to simplify this complex set, the algorithm only consider functions linear by parts. The last thing we are going to need for this algorithm is the concept of augmented log-likelihood $\Lambda(\phi, q)$ such that:[59]

$$\Lambda(\phi, q) = l(\phi, q) - \int e^{\phi(x)}\, dx - q + 1$$

We can then define the algorithm for the logconcens estimator using these concepts. We use the description of the algorithm as given in [59] and in [60].

Algorithm 4: Logconcens algorithm

Data: L_i: the left bound of each interval
R_i: the right bound of each interval
N: length of L_i and of R_i
Result: $\phi(x)$: concave function linear by parts
Find q such that $\frac{1}{N}\sum_{i=1}^{N}\frac{\mathbb{1}\{R_i=\infty\}}{\int_{L_i}^{\infty} e^{\phi(x)}\, dx+q} = 1$
Use this parameter to define the augmented log-likelihood
Define $\{\tau_1, ..., \tau_{m+1}\} = \{L_1, ..., L_n\} \cup \{R_1, ..., R_n\} \cup \{\infty\}$ with $-\infty < \tau_1 < ... < \tau_m < \tau_{m+1} = \infty$
Maximize the augmented log-likelihood using the expectation-maximization algorithm on the set of function linear by parts on $[\tau_i, \tau_{i+1}]$.

In practice the logconcens estimator returns us the evaluation of $F(x)$ at the point $t_1, ..., t_{m+1}$. We can then use an interpolation method in order to obtain a continuous function. As the algorithm also returns $F(x)$ evaluated at the t_i, we are going to approximate the cdf by using these points, then using cubic splines in order to interpolate the missing points. The method we are going to use is the *Fritsch-Carlson method*, which has the interesting property of conserving monoticity of the interpolant and obtaining a differentiable curve. Since $F(x)$ must be monotone and differentiable, these are obviously two interesting properties. The *Fritsch-Carlson* algorithm for monotone interpolation was introduced by *Fritsch et al.* in [61] and can be described in the following way:

Algorithm 5: Fritsch-Carlson algorithm

Data: x: vector of the points where the function is known
y: vector of the evaluation of the function
N: length of x
Result: $c_i(x)$: the cubic splines interpolating the data
for $i \in \{1, ..., N\}$ **do**
  $h_i = x_{i+1} - x_i$
  $\delta_i = \frac{y_{i+1} - y_i}{h_i}$
  $m_i = \frac{\delta_{i-1} + \delta_i}{2}$
  $\alpha_i = \frac{m_i}{\delta_i}$
  $\beta_i = \frac{m_{i+1}}{\delta_i}$
  $c_i(x) = \frac{\delta_i}{h_i^2}[(\alpha_i + \beta_i - 2)(x - x_i)^3 - (2\alpha_i + \beta_i - 3)h_i(x - x_i)^2 + \alpha_i h_i^2 (x - x_i)] + y_i$

Since the Fritsch-Carlson method is a method using cubic splines to interpolate a monotone function, the function obtained is twice differentiable and we can obtain the pdf of the random variable efficiently by simply differentiating each cubic spline individually.

3.7 Implementation of the method and different results

3.7.1 Introduction and methodology

Now that we have described our general algorithm, we will detail the different results obtained and the precision obtained by our methods. We start by describing our general methodology. In order to do that, we are going to introduce two hypothesis tests for comparing multiple cdf: the *Cramer-Von Mises test* and the *Kolmogorov–Smirnov test*. The two tests were best summarized by *Darling* in [62], we use here his definitions.

Definition 3.7.1: Cramer-Von Mises test

Let us have two cdf $F_X(x)$ and $F_Y(x)$ derived from a sample $x_1, ..., x_N$ of size N and a sample $y_1, ..., y_M$ of size M. The Cramer-Von Mises test tests the hypothesis $H_0 : F_X(x) = F_Y(x)$ against the hypothesis $H_1 : F_X(x) \neq F_Y(x)$ by studying the statistic U, where U is given by:

$$U = N \sum_{i=1}^{N} (r_i - i)^2 + M \sum_{j=1}^{M} (s_j - j)^2$$

where $r_1, ..., r_N$ are the ranks of the x in the combined sample and $s_1, ..., s_M$ are the ranks of the y is the combined sample.

Definition 3.7.2: Kolmogorov–Smirnov test

Let us have two cdf $F_X(x)$ and $F_Y(x)$. The Kolmogorov–Smirnov test tests the hypothesis $H_0 : F_X(x) = F_Y(x)$ against the hypothesis $H_1 : F_X(x) \neq F_Y(x)$ by studying the statistic D, where D is given by:

$$D = \sup_x |F_X(x) - F_Y(x)|$$

These two tests will be used together in order to assess the quality of the estimation of the cdf. What we will do is use the two tests together: first we will use the Cramer-Von Mises test in order to assess whether or not our fit is good enough to reject the hypothesis of an alternative distribution, then we will use the Kolmogorov-Smirnov test in order to compute the average distance between the original and the approximated cdf.
Our methodology will thus be the following:

- Generate the distribution(s) of the adverse bids using some chosen probability distribution(s)
- Compute the M highest bids of each sample
- Simulate a bidding pattern of an advertiser by using a uniform bid distribution
- Estimate the distribution(s) of the adverse bid using the censored observations
- Test whether our estimator is close enough to the real distribution using the Cramer-Von Mises and the Kolmogorov-Smirnov test

In order to be able to apply this methodology, we are going to need to compute the probability distributions of the *order statistics* of the adverse bids. This can be done using a powerful theorem, the *Bapat–Beg Theorem*, as proved by *Bapat et al.* in [39]. In order to understand this theorem, we are going to introduce a concept used in matrix theory, the concept of the *permanent* of a matrix. We use here the definition of *Glynn* in [63].

Definition 3.7.3: Permanent of a Matrix

Let $A = (a_{ij})$ be a $m \times m$ matrix over a field F and S_m the symmetric group, i.e. the set of all permutations in $\{1, ..., m\}$. Then the permanent of the matrix A is written $per(A)$ and is given by:

$$per(A) = \sum_{\gamma \in S_m} \prod_{i=1}^{m} a_{i,\gamma(i)}$$

The permanent of a matrix is closely linked to the determinant of the matrix: if we take for example the case of a two by two matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ then $per(A) = ad + bc$. This is thus extremely close to the determinant of A, given by $det(A) = ad - bc$.
Now that we have explained the concept of permanent, we can introduce the *Bapat-Beg Theorem*, as described in [39].

Theorem 3.7.1: Bapat-Beg Theorem

Let $X_1, ..., X_n$ be independent random variables with respective distribution functions $F_1(x), ..., F_n(x)$. Then the distribution functions of the order statistics of $X_1, .., X_n$ written Y_r, with $1 \leq r \leq n$ is given by:

$$P[Y_r \leq y] = \sum_{i=r}^{n} \frac{1}{i!(n-i)!} per \begin{pmatrix} F_1(y) & 1 - F_1(y) \\ \vdots & \vdots \\ \underbrace{F_n(y)}_{i} & \underbrace{1 - F_n(y)}_{n-i} \end{pmatrix}$$

The underbrace under the column represent the number of columns of the respective block matrix.

This theorem can thus be used to compute the probability distributions of the order statistics that are useful for the different statistical tests. We have also talked before that we would generate our adverse bids by using multiple types of distributions for adverse bids.

The first type of distributions we will use to generate adverse bids are simple normal distributions, with a mean and a standard deviation that will be uniformly distributed on a chosen interval. The choice of this type of distributions is due to the particularities of our algorithm: since our algorithm will smooth the distribution using normal kernels and since the normal distribution is a log-concave distribution, testing our algorithm will be useful to assess the quality of our approximation in a "good" case.

The second type of distributions we are going to use are Pareto distributions. We decided to use Pareto distributions because they have two types of interesting properties: first they are not log-concave, second they have what we call *heavy tails*. We use here the definition given by *Foss et al.* in [64].

Definition 3.7.4: Heavy-tailed distribution

Let us define a random variable X with a cdf $F(x)$ and a survival function $S(x) = 1 - F(x)$. Then X has a heavy-tailed distribution if

$$\lim_{x \to \infty} \sup_{x} S(x) e^{\lambda x} = \infty \quad \forall \lambda > 0$$

Heavy tails distribution have quite important consequences in probabilistic optimization, as they may cause problems for constraints that must be satisfied with high probability. As this is one of the applications of our algorithm, we are going to need to test the behavior of our algorithm in those specific cases. Since Pareto distributions are as well not log-concave, we can then test our algorithm efficiently using this type of distributions.

Finally, the last type of distributions we are going to use are *mixture distributions*, as defined by *Huff* in [65].

Definition 3.7.5: Mixture distribution

Let $X_1, ..., X_N$ be N random variables with cdf $F_1(x), ..., F_N(x)$. Then the distribution of a random variable Y is a mixture of the distributions of $X_1, ..., X_N$ if $\exists w_i \geq 0, i = 1, ..., N$, $\sum_{i=1}^{N} w_i = 1$ such that $F_Y(x) = \sum_{i=1}^{N} w_i F_i(x)$

The main reason why we are testing our estimation algorithm using this type of distributions is because they represent a "worst-case scenario" (under our assumptions) for our algorithm: since we are using some random variables that are not log-concave and with heavy-tails in the mixture, this type of distributions have all the possible difficulties combined, and will thus be the most difficult to estimate.

3.7.2 Results

Now that we have detailed the methodology, we can start to talk about the results. In order to understand these results, we are going to present the methodology used. The tables are all written using the distributions described earlier. The adverse bids are all realizations of these distributions while we bid using a uniform distribution on a selected interval. We suppose we participate in the auctions T times and we use our observations to estimate the distributions. We then use an hypothesis test to check if the fit is good enough and we compute the distance between the original and estimated distribution. We then repeat this procedure N times: we then compute the percentage of the time where we managed to reject any alternative distribution (under the column %), the average distance between the two distributions (under the column Dist) and the average distance in the case where the hypothesis test has failed (column DistFail). We then comment the results in order to explain the differences.

We will start by talking about the results for our algorithm for the second-price auction, that are given in Table 3.1. We can thus see that the results are extremely dependent on the type of distribution studied and on the method used. As we can see using the percentage of success rejecting the alternative hypothesis, the estimation using the Kaplan-Meier estimator was precise enough that the hypothesis test was able to reject the alternative hypothesis in approximately 90% of the cases, whatever the type of underlying distribution! This proves the efficiency of our algorithm for the general case. However, the smoothing and the estimation by a log-concave distribution cause a big loss in the precision of the algorithm. In fact, the estimation by a log-concave distribution is only precise enough for hypothesis testing in the case where the underlying distribution is itself log-concave. We can see as well that the smoothing causes a big difference in the precision of the approximation, even for normal distributions: the cdf obtained should thus only be smoothed if necessary. We will discuss in Chapter 4 how the cdf should however be smoothed for a number of cases.

	Original			Smooth			Logcon		
Distrib	%	Dist	DistFail	%	Dist	DistFail	%	Dist	DistFail
Normal	88	0.0296	0.0563	83	0.0733	0.2993	88	0.0975	0.2226
Pareto	80	0.0844	0.2703	0	0.45	0.45	51	0.0789	0.1119
Mixture	90	0.0518	0.1313	5	0.2086	0.2165	2	0.2	0.2029

Table 3.1: Comparison of the methods for SP auctions (N=100, T=500)

Now that we have talked about the results of our algorithm for second-price auctions, we can talk about first-price auctions. The results are given in Table 3.2. We can thus see that the algorithm is way less precise than in the case of second-price auctions. This is due to the less precise estimation in the case of first-price auctions: in fact, since we never have exact observations but only intervals, the precision is not as good as for the second-price auctions, where we have direct observations. We can see as well that the log-concave estimator is the most precise one in the case where the underlying distribution is in itself log-concave: this confirms the intuition that having a basic idea about the form of the underlying distribution can lead to more precise results. Finally we can see that even the most efficient algorithm in the general case (no smoothing, not log-concave) still gives an average Kolmogorov-Smirnov statistic of almost 0.2, which shows that the error can be as high as 0.2 in the estimation of the distribution: this is a large error. We will see in Chapter 4 how this relatively large error still gives relatively good results for the method used.

	Original			Smooth			Logcon		
Distrib	%	Dist	DistFail	%	Dist	DistFail	%	Dist	DistFail
Normal	0	0.2253	0.2253	5	0.1447	0.1488	22	0.1128	0.1298
Pareto	0	0.2179	0.2179	0	0.2457	0.2457	8	0.1528	0.1595
Mixture	0	0.1999	0.1999	0	0.2924	0.2924	1	0.2016	0.2029

Table 3.2: Comparaison of the methods for FP auctions (N=100, T=500)

Now that we have finished talking about single-item auctions, we can start talking about multiple-item auctions. We will start by talking about our method for the estimation of the distributions for GSP auctions. The results for the GSP auction are given in Table 3.3. As we can see , the results are better than for the FP auction but still not as good as for the SP auction. This is due to the number of direct observations to the objective function: our method has more direct observations than for the FP auction but less than for the SP auction, due to the higher number of slots. We can see as well that the log-concave estimation gives good results when the underlying distributions are log-concave but that the results are not as good when the underlying distribution is not log-concave, as predicted by the theory. Finally, we can see that the average distance between the underlying distribution and the estimated one is equal to approximately 0.2, which shows that there still exists a big approximation error, whatever the method might be.

	Original			Smooth			Logcon		
Distrib	%	Dist	DistFail	%	Dist	DistFail	%	Dist	DistFail
Normal	16	0.1527	0.1732	14	0.1524	0.1706	15	0.1496	0.1689
Pareto	8	0.2107	0.2252	3.67	0.2148	0.2211	13.33	0.1945	0.2177
Mixed	4.33	0.1867	0.1929	1.67	0.228	0.231	1.67	0.2638	0.2674

Table 3.3: Comparison of the methods for GSP auctions (N=100, T=500)

We can finally talk about the VCG auction. The results of the different tests can be found in Table 3.4. Since the knowledge of the interval for the clearing price in VCG auctions is better than for GSP auctions, our estimation is better than for GSP auctions. We can see as well that the log-concave estimator gives a surprisingly precise approximation of the real distribution in most cases: this reinforces the idea that log-concave estimation methods should be used if the application requires it.

	Original			Smooth			Logcon		
Distrib	%	Dist	DistFail	%	Dist	DistFail	%	Dist	DistFail
Normal	20.33	0.2404	0.2917	20	0.2216	0.2656	32	0.212	0.2944
Pareto	9.33	0.2419	0.2622	2.67	0.2402	0.2454	16	0.1921	0.2196
Mixture	4.33	0.1951	0.2017	4.67	0.18	0.1862	8.67	0.179	0.191

Table 3.4: Comparison of the methods for VCG auctions (N=100, T=500)

4 | Formulation as a probabilistic optimization problem

Now that we have talked about the methods used in order to estimate the distribution of the highest adverse bids, we can use these methods in order to model an optimization problem with complex constraints. We are going to give a general formulation for the more classical constraints that can exist for the problem and try to give insight on how the problem obtained can then be solved efficiently using different methods. The problem we are going to study will be the problem of trying to maximize the expectation of profit of the advertiser, subject to probabilistic constraints: constraints that only have to be satisfied with a certain probability. This is what we are going to call a problem of *probabilistic optimization*.
Our goal behind this model is to try to group all the constraints found in the literature in order to study the behavior of a player in repeated auctions. As we are in the framework of Game Theory, we will assume that the player is rational and will thus try to maximize its expected revenue while trying to satisfy the different constraints. We will show in this chapter that under these assumptions, the behavior of the player will tend to an equilibrium behavior and we will explain how to compute this equilibrium behavior.

4.1 General formulation of the problem

4.1.1 Single-item auctions

The work done in this section is an extension of the work done by *Balseiro et al.* in [35].

In order to formally describe the problem, we will need to describe two random variables, X and P. Let us denote by X the random variable representing the highest adverse bid that we have estimated in Chapter 3 and P as the (possibly deterministic) random variable representing the price paid if we win the auction. Let us suppose we take part in T sequential auctions and let us write as X_t and P_t the t^{th} realization of the random variables X and P defined earlier. Finally, let us suppose we have a constant valuation $\bar{v}$ for the ad slots being sold, as developed in Chapter 2, Section 2.4.

In the framework of Game Theory, a rational player will have one objective, which is to try to maximize its expected revenue. The random variable representing the revenue obtained in auction t is equal to $\mathbb{1}\{b_t \geq X_t\}(\bar{v} - P_t)$; the goal of the advertiser is thus to maximize the quantity $\sum_{t=1}^{T} \mathbb{E}[\mathbb{1}\{b_t \geq X_t\}(\bar{v} - P_t)]$.

Now that we have described the objective function of this optimization problem, we will describe the different constraints that may be relevant for the advertiser.

The first constraint we are going to introduce is simply a budget constraint. Budget constraints in repeated auctions are one of the most studied type of constraints: they have notably been studied by *Amin et al.* in [1], by *Balseiro et al.* in [35] and by *Agarwal et al.* in [21]. Let us suppose we have a budget B for the T sequential auctions. The constraint of the satisfaction of budget can be formulated by writing $\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}P_t \leq B$. We can thus see that we have a random variable that has to be smaller than a scalar. In the general case, this is not a deterministic constraint. Instead of seeing one as such, we will see it as a probabilistic constraint: we will consider that the constraint does not always need to be satisfied, but that it only need to be satisfied in a given percentage of the cases, that will in general be close to 100%. If we write this percentage as θ, the constraint can be written as $P[\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}P_t \leq B] \geq \theta$. This is thus a condition on the random variable $\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}P_t$.

Another constraint that has been studied before in the literature is the *Revenue on Investment* (ROI) constraint. Such a constraint has notably studied by *Fisher* in [66]. The idea of this constraint is to divide the revenue by the costs and considering that this quantity is higher than a wanted percentage. This constraint can be equivalently written as $\frac{\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}\bar{v}}{\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}P_t} \geq R_{min}$, with R_{min} the minimal revenue on investment needed. Finally, this can be rewritten in order to avoid having a fraction, this is $\sum_{t=1}^{T} R_{min}\mathbb{1}\{b_t \geq X_t\}P_t - \bar{v}\mathbb{1}\{b_t \geq X_t\} \leq 0$. As this is a probabilistic constraint, it can be rewritten as $P[\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}(R_{min}P_t - \bar{v}) \leq 0] \geq \gamma$, with γ a given parameter.

Finally, the last constraint we will study is a constraint on an objective price. Such a constraint has notably been studied by *Zhang et al.* in [67]. Let us suppose we have an objective price P and that we want that at each auction that we win to pay a price very close to the price P. If we suppose that we want to have an error smaller than ϵ on average, this is equivalent to wanting that $\sum_{t=1}^{T} |\mathbb{1}\{b_t \geq X_t\}P_t - P| \leq T\epsilon$ in total. Since we still have a probabilistic constraint, we can formulate this constraint by defining a parameter μ and having that $P[\sum_{t=1}^{T} |\mathbb{1}\{b_t \geq X_t\}P_t - P| \leq T\epsilon] \geq \mu$.

We can thus write the full problem for probabilistic optimization for single-item auctions.

Definition 4.1.1: Probabilistic optimization for single-item auctions

A general formulation of the problem of probabilistic optimization of the revenue with multiple constraints can be written in the following way:

$$\min_{b_t} \mathbb{E}[\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}(P_t - \bar{v})]$$
$$P[\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}P_t \leq B] \geq \theta$$
$$P[\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}(R_{min}P_t - \bar{v}) \leq 0] \geq \gamma$$
$$P[\sum_{t=1}^{T} |\mathbb{1}\{b_t \geq X_t\}P_t - P| \leq T\epsilon] \geq \mu$$

where we consider that:

- X_t is a random variable with a known probability distribution. We suppose that all X_t are i.i.d. random variables. X_t represents the highest adverse bid in the auction.
- P_t is a variable that is dependent on X_t, on the bid made b_t and on the type of auction.
- θ, γ, μ and ϵ are all known parameters such that $\theta, \gamma, \mu \to 1$ and $\epsilon \to 0$

As we did not find any other relevant constraints on the subject, we are going to limit ourselves to these 3 constraints, but the framework can be extended if other relevant constraints arise.

4.1.2 Multiple-item auctions

We have now talked in detail on the formulation of our problem as a probabilistic optimization problem for single-item auctions, the concept can be extended to multiple-item auctions in order to have an algorithm valid in the general case. We will thus detail how we can formulate the problem for multiple-item auctions and then detail optimization techniques that can be used in this case.

To start, in order to be able to define the problem, we are first going to need to define an important concept, the concept of *First-Order Stochastic Dominance* or FSD. We use the definition of the concept as given by *Quirk et al.* in [68].

Definition 4.1.2: First-Order Stochastic Dominance

The random variable A has First-Order Stochastic Dominance over the random variable B if the two following conditions are satisfied:

$$\begin{aligned} \forall x, P[A \geq x] &\geq P[B \geq x] \\ \exists x \text{ s.t. } P[A \geq x] &> P[B \geq x] \end{aligned}$$

This is equivalent to $F_A(x) \leq F_B(x)$ $\forall x$ and $F_A(x) < F_B(x)$ for some x.

As a consequence of Definition 4.1.2 and of Theorem 3.7.1 (Bapat-Bag theorem), we can prove an interesting result about FSD and order statistics.

Lemma 4.1.1: FSD and order statistics

Let $X_1, ..., X_n$ be independent random variables and $Y_1, ..., Y_n$ be the order statistics of $X_1, .., X_n$. If there exists $y \in \mathbb{R}$ such that $F_i(y) \in (0, 1)$ $\forall i \in \{1, ..., n\}$, then Y_i is FSD over Y_j, $\forall i < j$.

Proof. We can prove this relation by recurrence. By Definition 4.1.2, we know that if A is FSD over B and B is FSD over C, then A is FSD over C. Therefore, the lemma is proved if we prove that Y_r is FSD over Y_{r+1}.
Let us write $F_r(y)$ as the cdf of Y_r and $F_{r+1}(y)$ as the cdf of Y_{r+1}. By the definition of FSD we have that we have to prove that $F_{r+1}(y) \leq F_r(y)$ $\forall y$. By Theorem 3.7.1

(Bapat-Bag theorem), we can show that:

$$F_{r+1}(y) \leq F_r(y)$$

$$\sum_{i=r+1}^{n} \frac{1}{i!(n-i)!} per \begin{pmatrix} F_1(y) & 1-F_1(y) \\ \vdots & \vdots \\ \underbrace{F_n(y)}_{i} & \underbrace{1-F_n(y)}_{n-i} \end{pmatrix} \leq \sum_{i=r}^{n} \frac{1}{i!(n-i)!} per \begin{pmatrix} F_1(y) & 1-F_1(y) \\ \vdots & \vdots \\ \underbrace{F_n(y)}_{i} & \underbrace{1-F_n(y)}_{n-i} \end{pmatrix}$$

$$0 \leq \frac{1}{r!(n-r)!} per \begin{pmatrix} F_1(y) & 1-F_1(y) \\ \vdots & \vdots \\ \underbrace{F_n(y)}_{r} & \underbrace{1-F_n(y)}_{n-r} \end{pmatrix}$$

This condition is always satisfied since the matrix is non-negative and the permanent of a non-negative matrix is always non-negative. If we suppose as well that there exists at least one y such that all the $F_i(y)$ are in the interval $(0, 1)$, then the quantity on the right-hand side is strictly positive for at least one y and Y_r is FSD over Y_{r+1}. □

The reason why this lemma is interesting is because of the specific case where we considered that the bids made by the advertisers were all realizations of independent random variables. In that specific case, we have that the highest adverse bids are always FSD over each other by definition.

Now that we have defined the concept of FSD between two random variables and proved our lemma, we can formally define the problem for multiple-item auctions. In order to do so, we need to make some assumptions about our knowledge in the auction.

- We take part in T sequential auctions. In each auction, the publisher is selling M ad slots
- We have an estimation of our valuation $\bar{v}_1, ..., \bar{v}_M$ and of the click-through rates $\beta_1, ..., \beta_M$ for each slot. We suppose as well that there exist a $\bar{v}$ s.t. $\bar{v}_k = \beta_k \bar{v}$ (see Chapter 2, Section 2.4)
- We have the (estimated) distribution of the cdf of $X^{(k)}$ (which represent the kth highest bid in an auction) and of the cdf of $P^{(k)}$ (which represent the clearing price of the kth highest slot). The $X^{(k)}$ are independent between each other.

- The r.v. $X^{(k)}$ and $P^{(k)}$ are such that $X^{(k)}$ is FSD over $X^{(k+1)}$, $k = 1, ..., M-1$ and $P^{(k)}$ is FSD over $P^{(k+1)}$, $k = 1, ..., M-1$. This is always true if $X^{(k)}$ are the order statistics of other random variables (see Lemma 4.1.1).

Now that we have defined all these quantities, we can define our problem in the following way.

Definition 4.1.3: Probabilistic optimization for multiple-item auctions

A general formulation of the problem of optimization of the revenue with multiple constraints can be written in the following way:

$$\min_{b_t} \mathbb{E}[\sum_{t=1}^{T}(\mathbb{1}\{b_t \geq X_t^{(1)}\}(P_t^{(1)} - \bar{v}_1) + \sum_{k=2}^{M} \mathbb{1}\{X_t^{(k-1)} \geq b_t \geq X_t^{(k)}\}(P_t^{(k)} - \bar{v}_k))]$$

$$P[\sum_{t=1}^{T}(\mathbb{1}\{b_t \geq X_t^{(1)}\}P_t^{(1)} + \sum_{k=2}^{M} \mathbb{1}\{X_t^{(k-1)} \geq b_t \geq X_t^{(k)}\}P_t^{(k)}) \leq B] \geq \theta$$

$$P[\sum_{t=1}^{T}(\mathbb{1}\{b_t \geq X_t^{(1)}\}(R_{min}P_t^{(1)} - \bar{v}_1) + \sum_{k=2}^{M}\mathbb{1}\{X_t^{(k-1)} \geq b_t \geq X_t^{(k)}\}(R_{min}P_t^{(k)} - \bar{v}_k)) \leq 0] \geq \gamma$$

$$P[\sum_{t=1}^{T}|(\mathbb{1}\{b_t \geq X_t^{(1)}\}P_t^{(1)} + \sum_{k=2}^{M}\mathbb{1}\{X_t^{(k-1)} \geq b_t \geq X_t^{(k)}\}P_t^{(k)}) - P| \leq T\epsilon] \geq \mu$$

where we consider that:

- θ, γ, μ and ϵ are all known parameters such that $\theta, \gamma, \mu \to 1$ and $\epsilon \to 0$

We have thus fully defined the problem of probabilistic optimization in the case of repeated online auctions.

4.2 Supposition of equal bids

While we have now properly defined our problem, we can try to simplify it. In fact, while we have managed to find methods to estimate the distribution of the X_t, our constraints have expressions that depend on the *sum* of random variables dependent on X_t and b_t. However, this is in general not easy to find. However, there is an interesting property that we can use in order to simplify the problem: that all the b_t are equal in the optimal solution. We are going to prove that property in the two types of problem.

We start by introducing the concepts of convex sets and closed convex set, as defined by *Boyd et al.* in [75].

Definition 4.2.1: Convex sets

A set C is a convex set if $\forall x, y \in C$ and $\forall t \in [0, 1]$:

$$tx + (1 - t)y \in C$$

We introduce as well the following definitions related to topology as defined by *Babinec et al.* in [70].

Definition 4.2.2: Limit point

Let $B(x, r)$ be the open ball centered in x of radius r and C be a closed set. Then x is a limit point of C if:

$$\forall r > 0, \exists y \in B(x, r) \quad s.t. \quad y \notin C$$

Definition 4.2.3: Closed convex set

A set C is a closed convex set if it contains all its limit points.

Definition 4.2.4: Minimum and minimizers

Let us define C a closed convex set, $f : C \to \mathbb{R}$ a function. Then f^* is a minimum of f on C if $f^* \leq f(x)\ \forall x \in C$ and $\exists x^* \in C$ s.t. $f(x^*) = f^*$. x^* is then called the minimizer of f on C.

Now that we have defined these concepts, we can prove an important result.

Proposition 4.2.1: Minimizer of a sum of functions

Let us have $C \subseteq \mathbb{R}$ a closed convex set, $h : C \to \mathbb{R}$ a function defined on C.
Then there exists a minimizer $(x_1^*, ..., x_n^*)$ of the function $g(x_1, ..., x_n) = \sum_{i=1}^{n} h(x_i)$ such that $x_1^* = ... = x_n^* = \arg\min_{x \in C} h(x)$

Proof. Since C is a closed convex set, we know by the extreme value theorem that $\exists x^* \in C$ such that $h(x^*) = \min_{x \in C} h(x)$.
We can then prove the rest of the result by induction on n. We start with $n = 2$. We have that $g(x_1, x_2) = h(x_1) + h(x_2)$. We search for $\min_{(x_1,x_2) \in C^2} g(x_1, x_2)$. Since we are searching for a minimum of a function in two dimensions and that each function is minimized on the same point, we have that $\min_{(x_1,x_2) \in C^2} g(x_1, x_2) = \min_{x_1 \in C} h(x_1) + \min_{x_2 \in C} h(x_2) = 2h(x^*)$ and we proved that we have a minimizer $x = (x_1^*, x_2^*)$ such that $x_1^* = x_2^*$.
We then suppose that the result is true for n and we prove it is also true for $n + 1$. Let us have $g(x_1, ..., x_n, x_{n+1}) = \sum_{i=1}^{n+1} h(x_i) = \sum_{i=1}^{n} h(x_i) + h(x_{n+1})$. Since we suppose that the result is true for n, we have that $\sum_{i=1}^{n} h(x_i)$ is minimized for $x_1 = ... = x_n = x^*$ and since x^* is also the minimizer of $h(x)$, we have that $g(x_1, ..., x_{n+1})$ is minimized for $x_1 = ... = x_{n+1} = x^*$ which concludes the proof. □

This proposition has interesting consequences for our problem.

Corollary 4.2.1: Equal bids for single-item auctions

Let $C \subseteq \mathbb{R}$ be the space of feasible solutions and let us define $g(b_1, ..., b_T) = \mathbb{E}[\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}(P_t - \bar{v})]$ with each $b_t \in C$.
If we suppose that C is closed and convex, then $b^* = \arg\min_{b \in C} \mathbb{E}[\mathbb{1}\{b_t \geq X_t\}(P_t - \bar{v})]$ is such that $g(b^*, ..., b^*)$ is minimal.

Proof. By the linearity of the expectation we have that $g(b_1, ..., b_T) = \mathbb{E}[\sum_{t=1}^{T} \mathbb{1}\{b_t \geq X_t\}(P_t - \bar{v})] = \sum_{t=1}^{T} \mathbb{E}[\mathbb{1}\{b_t \geq X_t\}(P_t - \bar{v})]$. We write as $h(b_t) = \mathbb{E}[\mathbb{1}\{b_t \geq X_t\}(P_t - \bar{v})]$ and we have therefore that $g(b_1, ..., b_T) = \sum_{t=1}^{T} h(b_t)$. By Proposition 4.2.1, we know that there exists $b^* = \arg\min_{b \in C} h(b)$ and that $g(b^*, ..., b^*)$ is minimal. □

Corollary 4.2.2: Equal bids for multiple-item auctions

Let $C \subseteq \mathbb{R}$ be the space of feasible solutions and let us define $g(b_1, ..., b_T) = \mathbb{E}[\sum_{t=1}^{T}(\mathbb{1}\{b_t \geq X_t^{(1)}\}(P_t^{(1)} - \bar{v}_1) + \sum_{k=2}^{M} \mathbb{1}\{X_t^{(k-1)} \geq b_t \geq X_t^{(k)}\}(P_t^{(k)} - \bar{v}_k))]$ with each $b_t \in C$.
If we suppose that C is closed and convex, then $b^* = \arg\min_{b \in C} \mathbb{E}[(\mathbb{1}\{b \geq X_t^{(1)}\}(P_t^{(1)} - \bar{v}_1) + \sum_{k=2}^{M} \mathbb{1}\{X_t^{(k-1)} \geq b \geq X_t^{(k)}\}(P_t^{(k)} - \bar{v}_k))]$ is such that $g(b^*, ..., b^*)$ is minimal.

Proof. The proof follows from Corollary 4.2.1. □

This two corollaries are really interesting in our cases since they show that if we can prove that the set defined by the probabilistic constraints is closed and convex, then the problem can be reduced to one dimension. We are thus going to prove in Section 4.4 of this chapter that the problem is in fact defined on a closed and convex set.

4.3 Simplification of the constraints

Once we have made this supposition, we can try to simplify our problem. In fact, while we have defined our problem using probabilistic constraints, these constraints are in general quite difficult to interpret and generally require Monte Carlo simulations (see [71] by *Henrion* for more details). This is however not the case for our problem, as it has specific properties. One of the most interesting properties of our problem is the fact that the constraints are always constraints on the sum of i.i.d. variables. We can thus use this property in order to simplify our constraints. For that, we are going to introduce a well-known theorem: the Central Limit Theorem or CLT. We use the definition given in [72] by *Billingsley*.

Theorem 4.3.1: Central Limit Theorem (CLT)

Let $X_1, ..., X_T$ be a random sample of size T, that is a sample of i.i.d. random variables, with $E[X_1] = \mu$ and $E[(X_1 - \mu)^2] = \sigma^2 > 0$. Let us also define as $S_T = \sum_{t=1}^{T} X_t$. Then, as T approaches infinity, we have:

$$\sqrt{T}(S_T - \mu) \to N(0, \sigma^2)$$
$$P[\frac{S_T - T\mu}{\sqrt{T}\sigma} \leq z] \to \Phi(z)$$

The convergence is uniform in z.

This theorem is thus useful for our algorithm, as it will make it possible for us to approximate the constraint efficiently. In order to be able to simplify the constraints, we first see that the constraints for all types of auctions can all be expressed in the following way:

$$P[\sum_{t=1}^{T} Y_t \leq B] \geq \theta$$
$$P[\sum_{t=1}^{T} R_t \leq 0] \geq \gamma$$
$$P[\sum_{t=1}^{T} |Y_t - P| \leq T\epsilon] \geq \mu$$

In these constraints, Y_t and R_t are defined differently in the case of the auction that is being used and are dependent on X_t, b_t and the parameters of the problem. The explicit expression of the distribution of these random variables can be found using the definitions of the different type of auctions. We will now explain how we can use the CLT in order to simplify the constraints to our problem.

Corollary 4.3.1: Approximation of the budget constraint

Let us reuse the condition $P[\sum_{t=1}^{T} Y_t \leq B] \geq \theta$ defined earlier and define the mean and standard deviation of Y_t as $\mu_Y(b)$ and $\sigma_Y(b)$. Then we have that the condition

$$\lim_{T\to\infty} P[\sum_{t=1}^{T} Y_t \leq B] \approx \Phi(\frac{B - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}})$$

For a large T, this means that the condition $P[\sum_{t=1}^{T} Y_t \leq B] \geq \theta$ can be approximated by the equivalent condition

$$T\mu_Y(b) + \sigma_Y(b)\sqrt{T}\Phi^{-1}(\theta) \leq B$$

Proof. The condition $P[\sum_{t=1}^{T} Y_t \leq B] \geq \theta$ can be reformulated as follows:

$$P[\sum_{t=1}^{T} Y_t \leq B] \geq \theta$$
$$P[\frac{\sum_{t=1}^{T} Y_t - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}} \leq \frac{B - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}}] \geq \theta$$

Since the Y_t are i.i.d. distributed for a given b, we can use the CLT:

$$P[\frac{\sum_{t=1}^{T} Y_t - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}} \leq \frac{B - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}}] \geq \theta$$
$$\Phi(\frac{B - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}}) \geq \theta$$

Since $\Phi(x)$ is a monotonically increasing function in function of x, we obtain:

$$\Phi(\frac{B - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}}) \geq \theta$$
$$\frac{B - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}} \geq \Phi^{-1}(\theta)$$
$$T\mu_Y(b) + \sigma_Y(b)\sqrt{T}\Phi^{-1}(\theta) \leq B$$

□

Similarly, we can approximate the second and third conditions for a large T using the CLT.

Corollary 4.3.2: Approximation of the ROI constraint

Let us use the ROI constraint $P[\sum_{t=1}^{T} R_t \leq 0] \geq \gamma$ and define the mean and standard deviation of R_t as $\mu_R(b)$ and $\sigma_R(b)$. Then we have that

$$\lim_{T\to\infty} P[\sum_{t=1}^{T} R_t \leq 0] \approx \Phi(\frac{-T\mu_R(b)}{\sigma_R(b)\sqrt{T}})$$

For a large T, this means that the condition can be approximated by the equivalent condition

$$\sqrt{T}\mu_R(b) + \sigma_R(b)\Phi^{-1}(\gamma) \leq 0$$

Proof. The proof is similar to the proof of Corollary 4.3.1. □

While we can obviously use the same logic to be able to simplify the third constraint, we are going to have a problem if we use the same method as before, since the mean and standard deviation of the random variable defined in the constraint is way more difficult to compute. In order to avoid this problem, we are going to simplify the constraint using two-sided intervals. What we will do is consider that there should be as much probability of being over the objective than under it. This simplification of our problem seemed coherent: in fact, if an advertiser has an objective price decided, he would not want the bid obtained to be "biased" in one way or another.

This transformation can thus be written in the following way: we transform the constraint $P[\sum_{t=1}^{T} |Y_t - P| \leq T\epsilon] \geq \mu$ into the two constraints $P[\sum_{t=1}^{T}(Y_t - P) \geq T\epsilon] \leq \frac{1-\mu}{2}$ and $P[\sum_{t=1}^{T}(P - Y_t) \geq T\epsilon] \leq \frac{1-\mu}{2}$. In order to have a similar form as before, these constraints can be rewritten as

$$P[\sum_{t=1}^{T} Y_t \leq T(P+\epsilon)] \geq \frac{1+\mu}{2}$$

$$P[\sum_{t=1}^{T} Y_t \leq T(P-\epsilon)] \leq \frac{1-\mu}{2}$$

We can then simplify these constraints using the Central Limit Theorem.

Corollary 4.3.3: Approximation of the objective price constraint

Let us reuse the conditions $P[\sum_{t=1}^{T} Y_t \leq T(P+\epsilon)] \geq \frac{1+\mu}{2}$ and $P[\sum_{t=1}^{T} Y_t \leq T(P-\epsilon)] \leq \frac{1-\mu}{2}$ defined earlier and let us define the mean and standard deviation of Y_t as $\mu_Y(b)$ and $\sigma_Y(b)$. Then we have that:

$$\lim_{T\to\infty} P[\sum_{t=1}^{T} Y_t \leq T(P+\epsilon)] \approx \Phi(\frac{T(P+\epsilon) - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}})$$

$$\lim_{T\to\infty} P[\sum_{t=1}^{T} Y_t \leq T(P-\epsilon)] \approx \Phi(\frac{T(P-\epsilon) - T\mu_Y(b)}{\sigma_Y(b)\sqrt{T}})$$

For a large T, this means that the two conditions can be approximated by the equivalent conditions

$$\sqrt{T}\mu_Y(b) + \sigma_Y(b)\Phi^{-1}(\frac{1+\mu}{2}) \leq (P+\epsilon)\sqrt{T}$$

$$\sqrt{T}\mu_Y(b) + \sigma_Y(b)\Phi^{-1}(\frac{1-\mu}{2}) \leq (P-\epsilon)\sqrt{T}$$

Proof. The proof is similar to the proof of Corollary 4.3.1. □

Now that we have fully simplified our domain, we can come back to one of our assumptions, the assumption of the independence of the highest adverse bids. While all the results we have proved in this thesis are also valid if this assumption is not satisfied, the computation of the domain here requires the assumption of independence or the knowledge of the joint pdf. In fact, if the computation of the standard deviation of the different quantities is possible if the highest adverse bids is independent, it is impossible to do if we do not have any knowledge of the joint pdf between the highest adverse bids. This is thus the reason why this hypothesis is needed in the general case.

4.4 Convexity of the domain

In this section, we will prove that our domain may be further simplified, using the fact that is convex. In fact, since a convex domain on $\mathbb{R}$ is an interval, we can simplify the domain by finding the roots of the inequalities defined in the previous section. In order to do so, we will prove that each constraint individually defines a convex set. We are going to start by proving a first lemma that will be used later.

Lemma 4.4.1: Monotonicity of a sum of random variables

Let us suppose we have some random variables $Y_1(b), ..., Y_T(b)$ dependent on a parameter b with a cdf $F_{Y(b)}(x)$ differentiable in x and in b. If the function $g(b, x) = P[Y_1(b) \leq x]$ is decreasing in b, then the function $h(b, x) = P[\sum_{t=1}^{T} Y_t(b) \leq x]$ is decreasing is b.

Proof. Let us prove it for $T = 2$ and go on by induction. We use the fact that $g(b, a) = F_{Y(b)}(a)$. By the definition of the sum of random variables, we have that $P[Y_1(b) + Y_2(b) \leq x] = F_{Y(b)}(x) * f_{Y(b)}(x)$, where $*$ denotes a convolution. Since $F_{Y(b)}(x)$ is differentiable in x and in b, we have that $\frac{\partial F_{Y(b)}(x)}{\partial b} \leq 0$. We differentiate the cdf of the sum of random variables in function of b. By the properties of convolution, we obtain:

$$\begin{aligned}\frac{\partial P[Y_1(b) + Y_2(b) \leq x]}{\partial b} &= \frac{\partial F_{Y(b)}(x)}{\partial b} * f_{Y(b)}(x) \\ &= \int_{-\infty}^{\infty} \frac{\partial F_{Y(b)}(z)}{\partial b} f_{Y(b)}(x - z)\, dz \\ &\leq 0\end{aligned}$$

We then suppose it is true for T and prove it for $T + 1$. We thus want to show that $P[\sum_{t=1}^{T+1} Y_t(b) \leq x]$ is decreasing in b. Since $P[\sum_{t=1}^{T+1} Y_t(b) \leq x] = P[\sum_{t=1}^{T} Y_t(b) + Y_{T+1}(b) \leq x]$, we can use once again the properties of the derivative of a convolution and the lemma is proved. □

This lemma can then be used in order to prove that the set defined by the constraints is convex. We start by the budget constraint.

Proposition 4.4.1: Convexity of the budget constraint

Let us define the random variable $Y_t(b)$ as the price paid in auction t when bidding a constant valuation b and let us define a budget constraint written as $P[\sum_{t=1}^{T} Y_t(b) \leq B] \geq \theta$. Then the set of b such that the budget constraint is satisfied is a convex set.

Proof. Let us suppose we bid $b = 0$ to each auction. Since the budget B is positive and the price paid $Y_t(b)$ is equal to 0 with probability 1, this constraint is always satisfied.
Let us now suppose we bid $b > 0$ to each auction. Regardless of the billing method used (in the auctions considered in this thesis), when we increase the bid to a series of auctions the total price paid after T auctions weakly goes up. In other words, if we define as $S(b) = \sum_{t=1}^{T} Y_t(b)$, then we have that if $b_1 < b_2$ then $S(b_2)$ is FSD over $S(b_1)$. By the property of FSD, this means that $F_{S(b_2)}(x) \leq F_{S(b_1)}(x) \quad \forall x$. Therefore, if we come back to our constraint this means that $F_{S(b_2)}(B) \leq F_{S(b_1)}(B)$. This thus means that the probability of the constraint being satisfied decrease in the amount of the bid made.
Let us now divide our problem into two cases. Since we have shown before that the constraint is always satisfied when $b = 0$, we can evaluate our constraint at $b \to \infty$. If we have that $\lim_{b\to\infty} P[\sum_{t=1}^{T} Y_t(b) \leq B] \geq \theta$, this means that our constraint is always satisfied and that it is equivalent to say that $b \in \mathbb{R}^+$, which is a convex set. If we have the case that $\lim_{b\to\infty} P[\sum_{t=1}^{T} Y_t(b) \leq B] < \theta$, we have by the intermediate value theorem that there exists a b^* such that $P[\sum_{t=1}^{T} Y_t(b^*) \leq B] = \theta$ since the cdf of the distribution is supposed to be continuous. The set of the values satisfying the constraint is thus $b \in [0, b^*]$, which is a convex set. $\square$

We can do the same for the second constraint.

Proposition 4.4.2: Convexity of the ROI constraint for single-item auctions

Let us define the random variable $R_t(b) = \mathbb{1}\{b_t \geq X_t\}(R_{min}P_t - \bar{v})$ and let us define our ROI constraint written as $P[\sum_{t=1}^{T} R_t(b) \leq 0] \geq \gamma$. Then the set of b such that the budget constraint is satisfied is a convex set.

Proof. This will be done first for second-price auctions and then for first-price auctions.
Let us start by second-price auctions. We define $Z_t(b) = \mathbb{1}\{b \geq X_t\}$ and $Y_t(b) =$

$Z_t(b)X_t$. We thus have that $R_t(b) = R_{min}Z_t(b)X_t - \bar{v}Z_t(b) = Z_t(b)(R_{min}X_t - \bar{v})$. In order to study the constraint, we are going to compute the probability distribution of the random variable. We can do that by using the law of total probability.

$$\begin{aligned} P[R_t(b) \leq x] &= P[\mathbb{1}\{X_t \leq b\}(R_{min}X_t - \bar{v}) \leq x] \\ &= P[R_{min}X_t - \bar{v} \leq x, X_t \leq b] + P[0 \leq x, X_t > b] \end{aligned}$$

We can then evaluate the probability distribution in $x = 0$.

$$P[R_t(b) \leq 0] = P[X_t \leq \frac{\bar{v}}{R_{min}}, X_t \leq b] + P[X_t > b]$$

If we consider the case that $b \leq \frac{\bar{v}}{R_{min}}$, we have that $P[R_t(b) \leq 0] = P[X_t \leq b] + P[X_t > b] = 1$ and the constraint is thus always satisfied. If we take the case where $b > \frac{\bar{v}}{R_{min}}$, the probability becomes $P[R_t(b) \leq 0] = F(\frac{\bar{v}}{R_{min}}) + 1 - F(b)$. This is thus a decreasing function of b. By Lemma 4.4.1, we know that the probability that the sum is also smaller than 0 is also a decreasing function of b and by using the same argument than in the proof of Proposition 4.4.1, we know that the constraint defines a convex set.
Let us now talk about first-price auctions. We define $Z_t(b) = \mathbb{1}\{b \geq X_t\}$ and $Y_t(b) = Z_t(b)X_t$. The cdf of $R_t(b)$ can be written as follows.

$$\begin{aligned} P[R_t(b) \leq x] &= P[\mathbb{1}\{X_t \leq b\}(R_{min}b - \bar{v}) \leq x] \\ &= P[R_{min}b - \bar{v} \leq x, X_t \leq b] + P[0 \leq x, X_t > b] \end{aligned}$$

We can then evaluate the probability distribution in $x = 0$.

$$P[R_t(b) \leq 0] = P[b \leq \frac{\bar{v}}{R_{min}}, X_t \leq b] + P[X_t > b]$$

If we consider the case where $b \leq \frac{\bar{v}}{R_{min}}$, we have that $P[R_t(b) \leq 0] = P[X_t \leq b] + P[X_t > b] = 1$ and the constraint is always satisfied. If $b > \frac{\bar{v}}{R_{min}}$, we have that $P[R_t(b) \leq 0] = 1 - F(b)$. This is a decreasing function of b. By Lemma 4.4.1, we know that the probability that the sum is also smaller than 0 is also a decreasing function of b and by using the same argument than in the proof of Proposition 4.4.1, we know that the constraint defines a convex set. □

This property can be expanded to multiple-item auctions but the proofs are a bit more complicated.

Proposition 4.4.3: Convexity of the ROI constraint for multiple-item auctions

Let us define the random variable $R_t(b) = \mathbb{1}\{b \geq X_t^{(1)}\}(R_{min}P_t^{(1)} - \bar{v}_1) + \sum_{k=2}^{M} \mathbb{1}\{X_t^{(k-1)} \geq b \geq X_t^{(k)}\}(R_{min}P_t^{(k)} - \bar{v}_k)$ and let us define a ROI constraint written as $P[\sum_{t=1}^{T} R_t(b) \leq 0] \geq \gamma$. Then the set of b such that the budget constraint is satisfied is a convex set.

Proof. This will be done first for GSP auctions and then for VCG auctions.
We thus start by GSP auctions. By definition of these auctions we have that $P^{(k)} = X^{(k)}$. By Lemma 4.4.1, we know that we can simplify the problem by studying the constraint $P[R_t(b) \leq 0]$. We start by simplifying the constraint by using the law of total probability.

$$\begin{aligned}
P[R(b) \leq 0] &= P[\mathbb{1}\{b \geq X^{(1)}\}(R_{min}X^{(1)} - \bar{v}_1) \\
&+ \sum_{k=2}^{M} \mathbb{1}\{X^{(k-1)} \geq b \geq X^{(k)}\}(R_{min}X^{(k)} - \bar{v}_k)] \\
&= P[X^{(1)} \leq b, R_{min}X^{(1)} - \bar{v}_1 \leq 0] \\
&+ P[X^{(1)} \geq b, X^{(2)} \geq b, R_{min}X^{(2)} - \bar{v}_2 \leq 0] \\
&+ ... + P[X^{(1)} \geq b, ..., X^{(M-1)} \geq b, X^{(M)} \leq b, R_{min}X^{(M)} - \bar{v}_M \leq 0] \\
&+ P[X^{(1)} \geq b, ..., X^{(M)} \geq b]
\end{aligned}$$

We are then going to subdivide this condition in function of b. We start by studying the case $b \leq \frac{\bar{v}_M}{R_{min}}$. The constraint then becomes:

$$\begin{aligned}
P[R(b) \leq 0] &= P[X^{(1)} \leq b] + P[X^{(1)} \geq b, X^{(2)} \geq b] \\
&+ ... + P[X^{(1)} \geq b, ..., X^{(M-1)} \geq b, X^{(M)} \leq b] \\
&+ P[X^{(1)} \geq b, ..., X^{(M)} \geq b] \\
&= 1
\end{aligned}$$

In order to see that this quantity is in fact equal to 1, we can simply use the fact that $P[X^{(1)} \geq b, ..., X^{(M-1)} \geq b, X^{(M)} \leq b] + P[X^{(1)} \geq b, ..., X^{(M)} \geq b] = P[X^{(1)} \geq b, ..., X^{(M-1)} \geq b]$ and go by recursion from there.

We can then study the case $\frac{\bar{v}_M}{R_{min}} \leq b \leq \frac{\bar{v}_{M-1}}{R_{min}}$. The constraint is then equal to:

$$\begin{aligned}P[R(b) \leq 0] &= P[X^{(1)} \leq b] + P[X^{(1)} \geq b, X^{(2)} \geq b] \\ &+ ... + P[X^{(1)} \geq b, ..., X^{(M-1)} \geq b, X^{(M)} \leq \frac{\overline{v_M}}{R_{min}}] \\ &+ P[X^{(1)} \geq b, ..., X^{(M)} \geq b] \\ &< 1\end{aligned}$$

We can then continue this reasoning for other values of b and see that the probability of the constraint being satisfied in monotonically decreasing in function of b. By Lemma 4.4.1, we know that the probability that the sum being also smaller than 0 is a decreasing function of b and by using the same argument than in the proof of Proposition 4.4.1, we know that the constraint defines a convex set.
For VCG auctions we use the same reasoning but using the formula of $P^{(k)} = \sum_{j=k}^{M}(\beta_j - \beta_{j+1})X^{(j)}$. We can do as before and study the constraint $P[R_t(b) \leq 0]$.

$$\begin{aligned}P[R(b) \leq 0] &= P[\mathbb{1}\{b \geq X^{(1)}\}(R_{min}P^{(1)} - \bar{v}_1) \\ &+ \sum_{k=2}^{M} \mathbb{1}\{X^{(k-1)} \geq b \geq X^{(k)}\}(R_{min}P^{(k)} - \bar{v}_k)] \\ &= P[X^{(1)} \leq b, R_{min}P^{(1)} \leq \bar{v}_1] + P[X^{(1)} \geq b, X^{(2)} \geq b, R_{min}P^{(2)} \leq \bar{v}_2] \\ &+ ... + P[X^{(1)} \geq b, ..., X^{(M-1)} \geq b, X^{(M)} \leq b, R_{min}P^{(M)} \leq \bar{v}_M] \\ &+ P[X^{(1)} \geq b, ..., X^{(M)} \geq b]\end{aligned}$$

We thus subdivide this condition in function of b. We start by the case $b \leq \frac{\bar{v}_M}{R_{min}(\beta_M - \beta_{M+1})} = \frac{\bar{v}}{R_{min}}$. If this constraint is satisfied we have that if $X^{(M)} \leq b$ implies that $P^{(M)} \leq \frac{\bar{v}_M}{R_{min}(\beta_M - \beta_{M+1})}(\beta_M - \beta_{M+1}) = \frac{\bar{v}_M}{R_{min}}$. Recursively, we have as well that $P^{(M-1)} = (\beta_{M-1} - \beta_M)X^{(M-1)} + P^{(M)} \leq (\beta_{M-1} - \beta_M)X^{(M-1)} + \frac{\bar{v}_M}{R_{min}}$ that has to be $\leq \frac{\bar{v}_{M-1}}{R_{min}}$. This is equivalent to $X^{(M-1)} \leq \frac{\bar{v}_{M-1} - \bar{v}_M}{R_{min}(\beta_{M-1} - \beta_M)} = \frac{\bar{v}}{R_{min}}$. We thus have that for $b \leq \frac{\bar{v}}{R_{min}}$ the constraint is always satisfied and for $b > \frac{\bar{v}}{R_{min}}$ the constraint monotonically decreases with b. By Lemma 4.4.1, we know that the probability that the sum being also smaller than 0 is a decreasing function of b and by using the same argument than in the proof of Proposition 4.4.1, we know that the constraint defines a convex set. $\square$

We can finish by proving the same for the third constraint.

> **Proposition 4.4.4: Convexity of the objective price constraint**
>
> Let us define the random variable $Y_t(b)$ as the price paid in auction t when bidding a constant valuation b and let us define objective price constraints written as $P[\sum_{t=1}^{T} Y_t \leq T(P+\epsilon)] \geq \frac{1+\mu}{2}$ and $P[\sum_{t=1}^{T} Y_t \leq T(P-\epsilon)] \leq \frac{1-\mu}{2}$. Then the set of b such that the constraint is satisfied is a convex set.

Proof. The first constraint is exactly the same constraint as the budget constraint with a budget equal to $T(P+\epsilon)$ and the proof is thus similar to the one of Proposition 4.4.1. The second constraint use a similar argument: as the price paid can only go down when we bid less, the constraint is equivalent to writing $b \geq b_{min}$, which defines a convex set. □

These propositions combined thus brings us to our main theorem.

> **Theorem 4.4.1: Convexity of the domain**
>
> The domain C of the points such that the probabilistic constraints are satisfied is a closed convex set.

Proof. By Propositions 4.4.1, 4.4.2, 4.4.3 and 4.4.4, we know that each probabilistic constraint defines a closed convex set and the intersection of closed convex sets is a closed convex set. □

This thus shows that the domain is closed and convex and that it can be reduced in one dimension by Corollaries 4.2.1 and 4.2.2. In the following sections, we will use the notation C to denote the set of the points satisfying the constraints.

4.5 Convexity of the objective function

An interesting property of optimization problems is whether there are convex or not. Let us first define a few notions. We use the definition of a convex function as defined by *Glineur et al.* in [73].

Definition 4.5.1: Convex function

Let us define I an interval. A function $f : I \to \mathbb{R}$ is convex if $\forall x, y \in I$ and $\forall t \in [0, 1]$:

$$f(tx + (1-t)y) \leq tf(x) + (1-t)f(y)$$

If we want to explain this definition intuitively, this means that if we draw a line connecting two points of the objective function, every point on the line is above the function; this definition will be used again later. We can then introduce another concept using the definition given in [73].

Definition 4.5.2: Convex optimization problem

A convex optimization problem is an optimization problem of the form $\min_{x \in C} f(x)$ where both $f(x)$ and C are convex.

Once we have defined these two concepts, we can talk about an important property of convex optimization problems, as summarized in [73].

Theorem 4.5.1: Local and global minimum in a convex problem

In a convex optimization problem, every local minimum is a global minimum.

This explains thus all the interest for convex optimization problems. Sadly, in our case the optimization problem is *not* convex. One example of that is by taking the second-price auction and supposing a normal distribution of the adverse bid. The objective function of such a distribution is given in Figure 4.1. It is possible to see that this function is not convex: in fact, it is possible to draw a straight line crossing the function in three points, something that is impossible for convex functions.

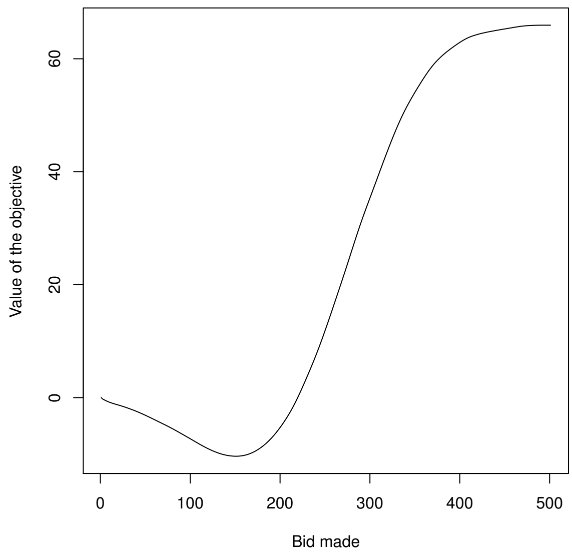

Figure 4.1: Convexity of the objective for second-price auctions

However, while the problem is not convex, we can prove that the problem satisfies another property: the property of quasi-convexity. Let us first introduce some concepts by using the definitions given by *Stoer et al.* in [57].

Definition 4.5.3: Quasi-convex function

Let us define I an interval. A function $f : I \to \mathbb{R}$ is quasi-convex if $\forall x, y \in I$ and $\forall t \in [0, 1]$:

$$f(tx + (1-t)y) \leq \max(f(x), f(y))$$

If the inequality is strict, the function is strictly quasi-convex.

Definition 4.5.4: Quasi-convex optimization problem

Let C be a convex set and $f : C \to \mathbb{R}$ a (strictly) quasi-convex function. Then the problem $\min_{x \in C} f(x)$ is a (strictly) quasi-convex optimization problem.

Quasi-convex functions do have some interesting properties in optimization. We use one of the theorem given by *Greenberg et al.* in [74].

Theorem 4.5.2: Quasi-convexity and minimizers

Let $f : C \to \mathbb{R}$ a quasi-convex function defined on a closed and convex set C. If f has a local minimum in a and if there exists a neighbourhood of a such that $f'(x) \neq 0$ everywhere in that neighbourhood (except in a), then a is a global minimum.

We will use this property later. We are first going to prove that our different problems have some interesting quasi-convex properties. In order to do that, we can use a property of quasi-convex functions given by *Stoer et. al.* in [57].

Theorem 4.5.3: Quasi-convexity and partitions

A function $f : I \to \mathbb{R}$ is quasi-convex if and only if there exists a sequential partition of I into two possibly empty intervals $I_1 = I^{\geq}$ and $I_2 = I^{\leq}$ such that f does not increase on $I^{\geq}$ and not decrease on $I^{\leq}$.

This property can be used in order to prove the quasi-convexity of the objective for single-item auctions.

Theorem 4.5.4: Quasi-convexity of the problem for second-price auctions

Let us assume that X has a continuous pdf. Then the problem of probabilistic optimization for second-price auctions is quasi-convex and admits a minimum in $\bar{v}$.
If we have as well that the pdf of X is strictly positive in a neighbourhood around $\bar{v}$, then that minimum is unique.

Proof. It can be proved that the objective function of the optimization problem for second-price auction is $O(b) = \int_0^b xf(x)\,dx - \bar{v}F(b)$, with f the pdf of X and F its cdf. Since we consider that the pdf of the distribution is continuous, we can differentiate the objective function. The derivative is then given by $O'(b) = (b - \bar{v})f(b)$. Since by definition the pdf is always non-negative, we have that the objective is decreasing for $b < \bar{v}$ and that it is increasing for $b \geq \bar{v}$. $\bar{v}$ is thus a local minimizer of the objective. By Theorem 4.5.3, we thus have that the objective function is quasi-convex.
If the pdf of X is strictly positive in a neighbourhood around $\bar{v}$, then the derivative of the objective is necessarily different than 0 in this neighbourhood. By Theorem 4.5.2, we know that $\bar{v}$ is then a unique global minimizer. $\square$

Theorem 4.5.5: Quasi-convexity of the problem for first-price auctions

Let us assume that X has a continuous pdf and that $\bar{v}$ is high enough i.e. it is possible to win an auction by bidding strictly less than $\bar{v}$. Then the problem of probabilistic optimization for first-price auctions admits a local minimum in $]0, \bar{v}[$ where the bidder makes a profit.
If we have as well that the pdf of X is log-concave, then the problem is quasi-convex and the objective function has a global minimum with a unique minimizer.

Proof. It can be proved that the objective function for first-price auctions is equal to $O(b) = F(b)(b - \bar{v})$. By the properties of the first-price auction, we know that the minimizer is necessarily in $[0, \bar{v}]$, since $O(b) \geq 0$ for $b > \bar{v}$ and $O(\bar{v}) = 0$. Since we suppose that the valuation is high enough, this means that $\exists a \in]0, \bar{v}[$ s.t. $O(a) < 0$. Since $O(b)$ is continuous and $O(0) = O(\bar{v}) = 0$, this means $\exists c \in]0, a], d \in [a, \bar{v}[$ s.t. $O'(c) < 0$ and $O'(d) > 0$. Since the pdf of X is continuous, so is $O'(b)$ and we can use the intermediate value theorem: $\exists e \in [c, d]$ s.t. $O'(e) = (e - \bar{v})f(e) + F(e) = 0$ and e is a local minimizer of $O(b)$ in $]0, \bar{v}[$.
If we suppose that the pdf of X is log-concave, we can prove more results. Since we have that $O'(e) = 0$, we have $f(e) = \frac{F(e)}{\bar{v}-e} > 0$. Since the pdf is supposed to be continuous, $\exists \epsilon > 0$ s.t. $f(e - \epsilon) > 0$ and $f(e + \epsilon) > 0$. Since the pdf of X is supposed to be log-concave, we have that $\forall t \in [0, 1]$:

$$f(t(e - \epsilon) + (1 - t)(e + \epsilon)) \geq f(e - \epsilon)^t f(e + \epsilon)^{1-t} > 0$$

We now want to use this to prove that $\exists I \subseteq [e - \epsilon, e + \epsilon]$, $e \in I$ s.t. $O'(x) \neq 0$ $\forall x \in I \backslash \{e\}$. The only case where this is not possible is if there exists $J \subseteq [e - \epsilon, e + \epsilon]$ with $e \in J$ s.t. $O'(x) = 0$ $\forall x \in J$. This means that, $\forall x \in J$, $F(x) + f(x)(x - \bar{v}) = 0$. Since $f(x) > 0$ for all points in I, we can see this as a differential equation. By resolving it, we obtain that it is satisfied if $F(x) = \frac{K}{\bar{v}-x}$, $K > 0$. As this is not a log-concave function, this is contradictory with our hypothesis and thus $\exists I \subseteq [e - \epsilon, e + \epsilon]$ s.t. $e \in I$ and $O'(x) \neq 0$ $\forall x \in I \setminus \{e\}$. By Theorem 4.5.2, this means that e is the unique minimizer of $O(b)$ on the domain.
Finally, to prove that $O(b)$ is quasi-convex, we can equivalently prove that $-O(b)$ is quasi-concave. Since the cdf of the distribution is also log-concave, we have that $-O(b)$ is the multiplication of two log-concave functions since $\bar{v} - b$ is log-concave on $[0, \bar{v}]$. Since the multiplication of log-concave functions is log-concave, we have that $-O(b)$ is log-concave. Since log-concavity implies quasi-concavity, we have that $-O(b)$ is quasi-concave and thus that $O(b)$ is quasi-convex. $\square$

While we have these interesting properties for single-item auctions, we sadly cannot extend these properties for all types of multiple-item auctions.

Theorem 4.5.6: Quasi-convexity of the problem for GSP auctions

Let us assume that $X^{(1)}, ..., X^{(M)}$ have continuous distributions. Then the objective function in the problem of probabilistic optimization for GSP auctions is a sum of quasi-convex functions.
The objective function has also a global minimum, and the minimizer is in $[0, \bar{v}]$.

Proof. Let us start by decomposing the objective function of the problem. By definition, we have that $\mathbb{1}\{X_t^{(k-1)} \geq b \geq X_t^{(k)}\} = \mathbb{1}\{b \geq X_t^{(k)}\} - \mathbb{1}\{b \geq X_t^{(k-1)}\}$. We can then compute the objective which is equal to :

$$O(b) = \sum_{k=1}^{M-1} (\mu_Y^{(k)}(b) - \bar{v}_k F^{(k)}(b) + F^{(k)}(b)\bar{v}_{k+1} - F^{(k)}(b)\mu_X^{(k+1)}) + \mu_Y^{(M)}(b) - \bar{v}_M F^{(M)}(b)$$

where $\mu_X^{(k)} = \mathbb{E}[X^{(k)}]$ and $\mu_Y^{(k)}(b) = \mathbb{E}[X^{(k)}\mathbb{1}\{b \geq X^{(k)}\}]$.
We can decompose this function into a sum of functions. We take:

$$\begin{aligned} d^{(k)}(b) &= \mu_Y^{(k)}(b) - \bar{v}_k F^{(k)}(b) + \bar{v}_{k+1} F^{(k)}(b) - F^{(k)}(b)\mu_X^{(k+1)} \quad k = 1, ..., M-1 \\ d^{(M)}(b) &= \mu_Y^{(M)}(b) - \bar{v}_M F^{(M)}(b) \end{aligned}$$

We thus have that $O(b) = \sum_{k=1}^{M} d^{(k)}(b)$. We can prove that each $d^{(k)}$ is quasi-convex by using Theorem 4.5.3. In fact, we can compute the derivative of each $d^{(k)}(b)$ individually. The derivative is the following:

$$\begin{aligned} d^{(k)}(b)' &= f^{(k)}(b)(b - \bar{v}_k + \bar{v}_{k+1} - \mu_X^{(k+1)}) \quad k = 1, ..., M-1 \\ d^{(M)}(b)' &= f^{(M)}(b)(b - \bar{v}_M) \end{aligned}$$

For $k < M$ we have that $d^{(k)}(b)$ is decreasing on $[0, \bar{v}_k - \bar{v}_{k+1} + \mu_X^{(k+1)}]$ and increasing on $[\bar{v}_k - \bar{v}_{k+1} + \mu_X^{(k+1)}, +\infty)$, by Theorem 4.5.3 it is thus quasi-convex. For $d^{(M)}(b)$ we have that the function is decreasing on $[0, \bar{v}_M]$ and increasing on $[\bar{v}_M, +\infty)$, it is thus quasi-convex as well. Therefore, $O(b)$ is a sum of quasi-convex functions.
For the second part, we use the fact that overbidding is a weakly-dominated strategy in GSP auctions compared to bidding the valuation[16]: this thus means that $\forall b > \bar{v}, O(b) \geq O(\bar{v})$ and that $O(b)$ has a global minimum, with the minimizer in $[0, \bar{v}]$. By the extreme value theorem, we know that this minimizer is attained and the proof is completed. $\square$

Theorem 4.5.7: Quasi-convexity of the problem for VCG auctions

Let us assume that $X^{(1)}, ..., X^{(M)}$ have continuous distributions. Then the problem of probabilistic optimization for VCG auctions is quasi-convex and admits a minimum in $\bar{v}$.
If we have as well that any of the pdfs of $X^{(1)}, ..., X^{(M)}$ are strictly positive in a neighbourhood around $\bar{v}$, then that minimum is unique.

Proof. By computing the expectation, we can show that the objective function is given by:

$$O(b) = \sum_{k=1}^{M} (\beta_k - \beta_{k+1}) \mu_Y^{(k)}(b) + F^{(k)}(b)(\bar{v}_{k+1} - \bar{v}_k)$$

Using the same decomposition than in the proof of Theorem 4.5.6, we can decompose this function into a sum of functions. We take:

$$d^{(k)}(b) = (\beta_k - \beta_{k+1}) \mu_Y^{(k)}(b) + F^{(k)}(b)(\bar{v}_{k+1} - \bar{v}_k)$$

By deriving each $d^{(k)}(b)$ we have:

$$d^{(k)}(b)' = f^{(k)}(b)((\beta_k - \beta_{k+1})b + \bar{v}_{k+1} - \bar{v}_k)$$

We thus have that each $d^{(k)}(b)$ is decreasing on: $[0, \frac{\bar{v}_k - \bar{v}_{k+1}}{\beta_k - \beta_{k+1}}] = [0, \bar{v}]$ and increasing $[\bar{v}, +\infty)$. Therefore, by Theorem 4.5.3, $O(b)$ is quasi-convex and $\bar{v}$ is a local minimum.
For the second part, by Theorem 4.5.2, if any $d^{(k)}(b)$ is strictly positive in a neighbourhood around $\bar{v}$ then the derivative of the objective is different than 0 in a neighbourhood of $\bar{v}$ and $\bar{v}$ is the unique minimizer. This is the case if any of the $f^{(k)}(b)$ is strictly positive around $\bar{v}$. $\square$

Now that we have proved these important results about the objective functions of the different optimization problems described until now, we can now describe different methods on how to resolve them.

4.6 Optimization method

Since we have proven in the previous section that our methods were either quasi-convex or the sum or quasi-convex functions, we are going to try to use this result in order to find methods to use. We can first start by introducing a well-known optimization method, the projected gradient method. We use the definition of the method given by *Glineur et al.* in [73].

Algorithm 6: Projected gradient method

Data: I: an interval
$f : I \to \mathbb{R}$ a smooth function
x_0: a starting point
P: the projection to I
Result: x^*: a point that minimizes $f(x)$
$k = 0$
$f_{min} = \infty$
while *Not converged* **do**
  Compute $g_k = f'(x_k)$
  Choose step size α_k
  $x_{k+1} = P[x_k - \alpha_k g_k]$
  $f_{min} = f(x_k)$
  $k = k + 1$

The gradient method is one of the most popular methods for convex optimization, thanks to its relatively simple implementation and its good convergence results. The method can also be applied to quasi-convex optimization under certain conditions on the step size. We can therefore introduce one well-known type of step size, the Polyak step size. We use the definition given by *Boyd et al.* in [69].

Definition 4.6.1: Polyak step size

Let us suppose we have a differentiable function $f : C \to \mathbb{R}$ to minimize, x_k a point in the gradient iteration, $\hat{f}^*$ an estimation of the optimal value and $g = f'(x_k)$, the derivative of the function f at x_k. Then the Polyak step size is given by:

$$\alpha_k = \frac{f(x_k) - \hat{f}^*}{g^2}$$

The Polyak step size is widely used for the gradient method for convex problems because of its good convergence results[69]. The interesting property for our problem is the convergence of this step size choice for quasi-convex problems.

In fact, using the Polyak step size, it has proved by *Wang et al.* in [76] that the projected gradient method converges to the minimum of the objective function. This is thus a method that can be used in order to minimize the problem in the case of single-item auctions and for VCG auctions.
For GSP auctions, we cannot use the same results than for the other types of auctions, as the sum of quasi-convex functions is in general not quasi-convex. We can however introduce a new method for this type of problems: the randomized incremental subgradient method. This method was introduced by *Hu et al.* in [77].

Algorithm 7: Randomized incremental subgradient method

Data: I an interval
$f_i : I \to \mathbb{R}$, $i = 1, .., M$: a sequence of smooth functions
x_0: a starting point
v_k: a step-size sequence
Result: x^*: a point that minimizes $f(x) = \sum_{i=1}^{M} f_i(x)$
$k = 0$
while *Not converged* **do**
  Let $I_k = \{i \in \{1, ..., M\} : f_i(x_k) > f_i^*\}$
  Pick up equiprobably a random variable ω_k from the set I_k
  Compute $g = f'_{\omega_k}(x_k)$
  $x_{k+1} = P[x_k - v_k g]$
  $k = k + 1$

The randomized incremental subgradient method has some interesting convergence property for functions that are the sum of quasi-convex functions. For a diminishing step size rule, it was proven in [77] that the algorithm converges to the global minimizer of the function with probability 1, if a global minimizer does exists. Since it is the case for GSP auctions, we can use this method in order to find the global minimizer of the objective of the problem for GSP auctions. The results obtained with this method will be detailed later in the chapter.

4.7 Convergence to the optimal bid

The last thing we can study in the problem of probabilistic optimization is the convergence of the algorithm developed in this thesis. We can show this important result.

Theorem 4.7.1: Convergence of the algorithm

Let us take any type of auction defined in this thesis and let us suppose our different assumptions about the distributions of the highest adverse bids are satisfied. Let us describe our algorithm in the following way:

- We bid in the auction (for the first few auctions we bid $\bar{v}$)
- We use our result in the auction to estimate the probability distributions of the highest adverse bids using the appropriate estimators estimators
- We use the estimated distribution to compute our domain C and to find the maximizer of the expected profit using the appropriate method

Then there exists an optimal bid maximizing the expected profit and our algorithm converges to a maximizer.

Proof. Regardless of the type of auction used, the optimal bid for the expected profit is in $[0, \bar{v}]$, as bidding over $\bar{v}$ is a dominated strategy in any auction[5]. We know as well that the estimators converge to the real distributions on the considered interval $[0, \bar{v}]$ by the property of these estimators (see [41] for parametric estimators, [49] for Kaplan-Meier estimator, [51] for Turnbull estimator and [59] for logconcens estimator). Since we know as well that the optimal bid is in $[0, \bar{v}]$, we know by the extreme value theorem that the objective function attains its maximum in the interval and thus that there exists a global maximizer of the expected profit in $[0, \bar{v}]$. Finally, as both the incremental subgradient method and the projected gradient method converge to the maximizer of the objective with probability 1 (see [76] for the projected gradient [77] and for the incremental subgradient), our algorithm converges to a maximizer of the expected profit in C. $\square$

This thus conclude the theoretical development this thesis: as we prove here that our algorithm converges to the optimal bid, we have managed to obtain our objective of developing an algorithm for behavior optimization using censored estimators.
This is thus going to be the last important theorem of this thesis: we will now apply this algorithm using simulations.

4.8 Implementation and discussion of the results

Now that we have described the methods used and proved all the necessary theorems, we can start to talk about their implementation. We are going to talk about all the different results obtained and how they compare to some basic strategies. As the results can be significantly different for different types of auctions, we are going to discuss them for each auction independently. We will start by talking about single-item auctions.

4.8.1 Second-price auctions

We can start by talking about the second-price auctions. Since we have shown before that our algorithm was extremely efficient for the estimation of the distributions, it is therefore also quite precise for the estimation of the objective function. A comparison of an exact and an estimated objective function is given in Figure 4.2.

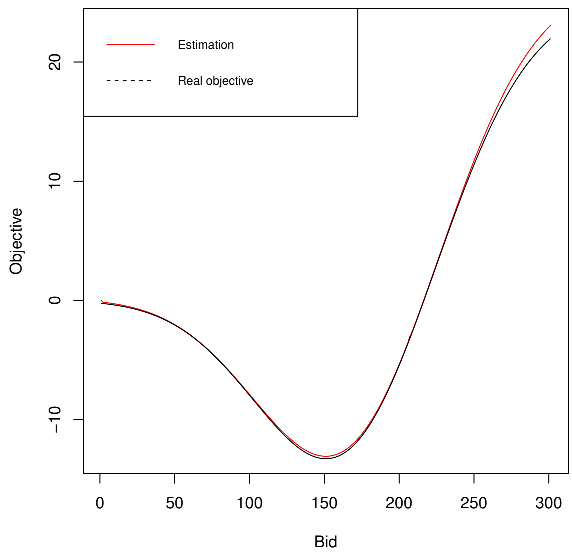

Figure 4.2: Comparison of the real and estimated objective (SP)

We can see that the difference of the two objective functions is extremely small: as such the difference in the optimal bid is extremely small. However, the estimation of the objective function in the unconstrained case is not as interesting as for the other types of auctions, since the second-price auctions are DSIC as we have shown at Theorem 2.2.1. Therefore, it is always a dominant strategy to bid its true valuation and the estimation of the objective function is not as important as it is in the other type of auctions.

4.8.2 First-price auctions

We can now talk about the second type of auctions, the first-price auction. As we have shown before, the estimation of the distribution for first-price auctions is not as efficient as it is for second-price auctions. As such, the estimation of the objective function is also not as precise. However, the algorithm still converges to an optimal bid with an approximation error that is quite small. An example of a real and estimated distribution of the highest adverse bid is given in Figure 4.3.

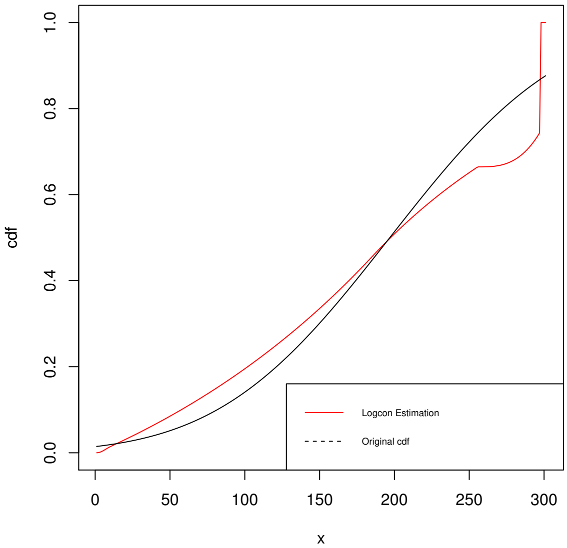

Figure 4.3: Comparison of the real and estimated cdf (FP)

We can see as well that there exists a significant difference between the two distributions, even if the difference in the probabilities is mostly centered around the lower tails of the distribution. However, the difference in the objective function is quite small, as we can see in Figure 4.4. This thus shows the quality of our approximation for the bid optimization.

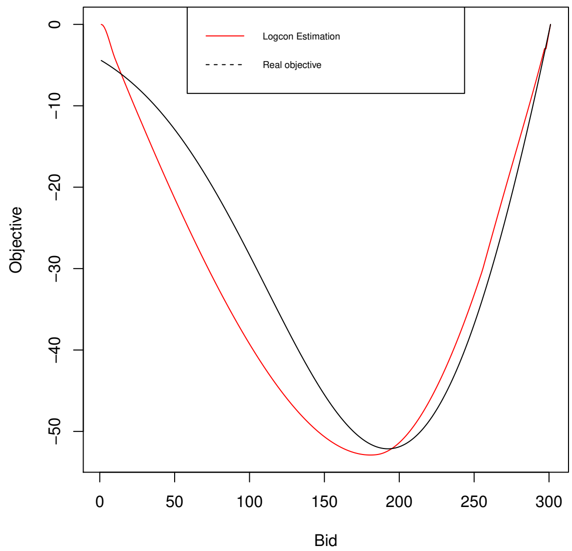

Figure 4.4: Comparison of the real and estimated objective function (FP)

Finally, we can talk about the convergence of the estimated optimized bid to the real optimal bid. The convergence to the optimal bid can be seen in Figure 4.5. We can see that the bid obtained by our algorithm converges to the bid that maximizes the expected profit, with a small error due to the approximation errors.

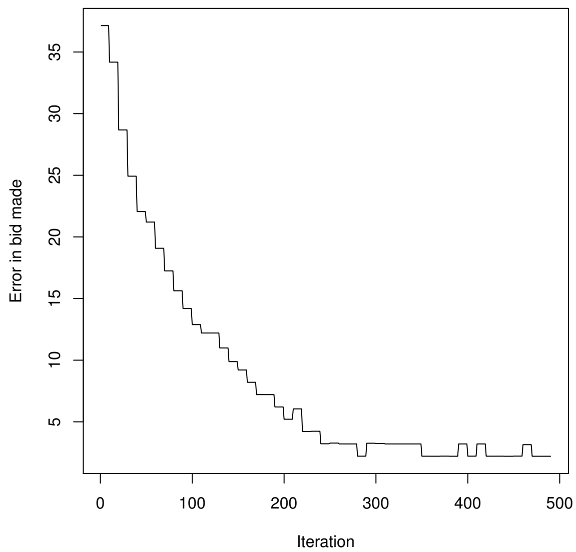

Figure 4.5: Convergence to the optimal bid (FP)

4.8.3 Generalized second-price auctions

Now that we have talked about single-item auctions, we can start to talk about multiple-item auctions. As we have shown before (see Theorem 2.2.3), the GSP auction is not DSIC and therefore we need to estimate the distributions of the different highest adverse bids in order to optimize our behavior. Once we have estimated the distributions, we can compute the estimated objective function. An example of such an objective function is given in Figure 4.6. As we can see, the estimation of the objective function is quite good, and the objective function obtained is quite close to the real objective function.

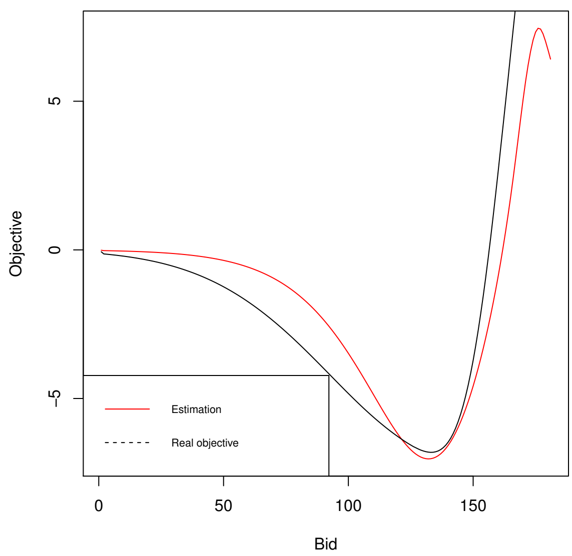

Figure 4.6: Comparison of the real and estimated objective function (GSP)

We can finally talk about the different bids obtained by the algorithm. The convergence of the estimated bid to the optimal bid can be seen in Figure 4.7. We might think that in this case the bids do not converge to the real one; however, it is important to note that the scale of the graphic is extremely small: the bids obtained by the algorithm are close to the optimal one from the beginning. We can still see that the algorithm has an approximation error, that has to be taken into account.

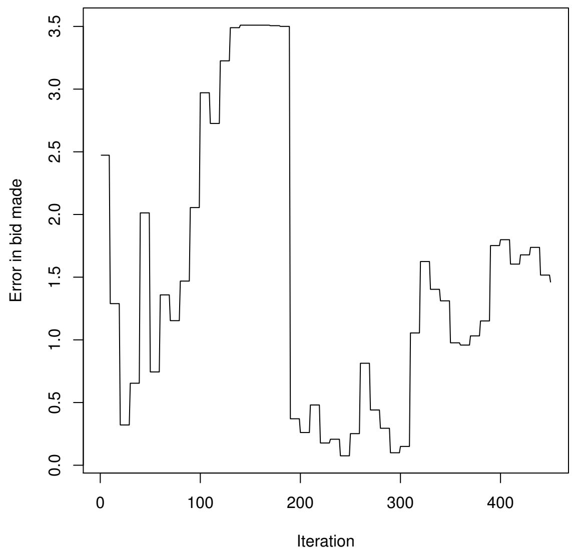

Figure 4.7: Convergence of to the optimal bid (GSP)

4.8.4 Vickrey-Clarke-Groves auctions

Now that we have talked about our algorithm for all the other type of auctions, we can introduce the results obtained by our algorithm for VCG auctions. As we have talked before, the VCG auctions are DSIC (see Corollary 2.2.4) and it is therefore the best strategy possible to bid its true valuation (note here that we are talking about the intrinsic valuation $\bar{v}$). We can still show that our algorithm converges to this predicted result by showing that the bids obtained by our optimization problem do in fact converge to the true valuation.

As we can see in Figure 4.8, the estimated and real objective functions are still quite close and they do have the same minimum. This minimum is attained for the real valuation of the bidder (the objective function in Figure 4.8 was plotted with a valuation of 300) and our results are thus coherent with the theory.

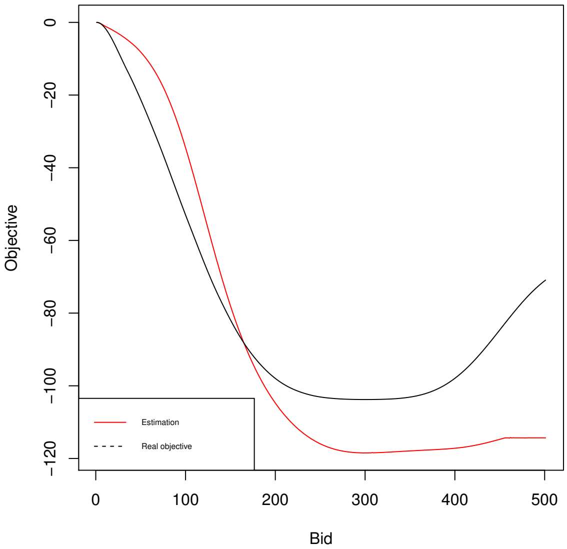

Figure 4.8: Comparison of the real and estimated objective function (VCG)

We can finally talk about the convergence of our optimized bid to our valuation. The evolution of the error is given in Figure 4.9. As we can see, our algorithm converges quite quickly to the valuation: there still exists an error in the optimal bid but it is mainly due to numerical imprecisions in the estimation of the distribution.

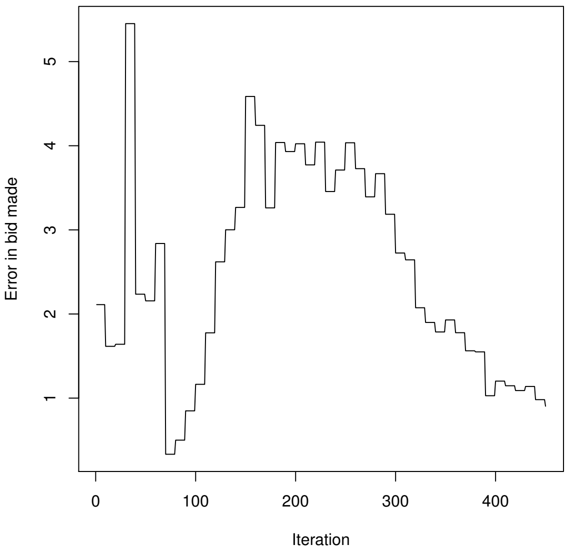

Figure 4.9: Convergence to the optimal bid (VCG)

5 | Conclusion and possible future work

5.1 Summary

In this thesis, we studied the use of censored estimators in the context of Game Theory and showed how these estimators can be used in order to find an optimal bid for advertisers participating in repeated online ad auctions.

We started our thesis in Chapter 2 by introducing the problem and by giving a short state of the art. We first defined some concepts of Game Theory then explained how these concepts are relevant to the design of online ad auctions. We then continued by introducing the types of auctions that are relevant to this thesis: the auctions used in online advertising. We explained their mechanism, modelled them as an infinitely repeated Bayesian game, then introduced their different properties. After that, we explained the mechanism relevant to online ad auctions in detail by explaining which type of auctions was used in which context. Finally, we talked about the methods that could be used by advertisers in order to estimate their valuation for an ad and the Click-Through Rate related to an ad slot, two quantities necessary for our model.

The second part of this thesis was focused on our contributions to the literature. We started by presenting our algorithm for adverse bid estimation in Chapter 3, which is one of our main contributions to the literature, extending the work done by multiple authors before us. We explained the hypotheses that needed to be made in order for us to use this algorithm. We continued by talking about censored estimators in detail: we talked about the different types of censored estimators, then went on to explain how these estimators could be used in our problem in order to approximate the distribution of the highest adverse bids. We finished Chapter 3 by explaining how we managed to implement our method and by giving the different results we obtained using simulations.

We finished this thesis in Chapter 4 by introducing our second main contribution to the literature: a new model for behavior optimization in repeated auctions using a probabilistic optimization model with various constraints. We introduced different properties and proved how our optimization problem satisfied these properties. Finally, we introduced algorithms that could be used in order to solve this problem in the general case, implemented these algorithms, made some simulations and discussed our results. During this discussion, we discussed the possible issues of our algorithm, but showed that it managed to converge to the optimal bid, which was the behavior we were interested in.

We can thus conclude this thesis by saying that we have managed to attain our objective, which was to study in detail the possibilities of using censored regression for behavior optimization in repeated auctions. By developing an algorithm that could be used in the general case and by using methods of probabilistic optimization, we have managed to provide a general framework that could be used by advertisers in order to improve their behavior in these repeated auctions. As we have proved that our algorithm converges to the optimal bid, we have managed to show that, under our hypotheses, our algorithm maximizes the expected profit of the advertisers.

5.2 Future work and possible improvements

While we have made some noticeable contributions to the current literature on the subject of bid optimization, we firmly believe that there exists possible improvements to our algorithm. We are here going to list here some of these possible improvements.

The main improvements that could be made in our opinion relates to the strictness of our assumptions. In fact, the assumptions we had to make were quite strong; however, there seems to be possibilities to improve our algorithm by relaxing these assumptions.
The most obvious assumption to relax is the assumption of independence of the highest adverse bids. This however requires further work on the theory of censored estimators in order to be possible, but as there seems to be significant research on the subject, it seems plausible that some new estimators will appear in the literature, making it possible to improve our current algorithm.

The second way we could relax our assumptions is by permitting the distributions of the adverse bids to be non-stationary, with probability distributions that change over time. Once again, this seems to be possible by supposing these probability distributions change continuously over time, and by adapting our algorithm in order to give more weight to more recent observations, permitting us to improve our estimation of the probability distributions.
A final example of an improvement that could be made is to improve our model of the valuations. For instance, a possible improvement of our valuation model would be to permit a valuation function with decreasing marginal values for clicks (see [5] p. 711). This extension seems to be compatible with our framework: we could in fact imagine to have a valuation that is a function of the expected number of auctions won, something that seems possible to do in our model.

We can thus see that our model allows for improvements that should be further investigated.

Bibliography

[1] Kareem Amin, Michael Kearns, Peter Key, Anton Schwaighofer. *Budget Optimization for Sponsored Search: Censored Learning in MDPs*. October 2012.

[2] Wush Chi-Hsuan Wu, Mi-Yen Yeh, Ming-Syan Chen. *Predicting Winning Price in Real Time Bidding with Censored Data*. KDD'15, August 10-13, 2015, Sydney, NSW, Australia.

[3] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, July 1991.

[4] Raphaël Jungers, Benoît Legat, and Matthew Philippe. *Game Theory: Rational decisions among individuals*. Course notes, February 2017.

[5] Noam Nisan , Tim Roughgarden , Eva Tardos , Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, 2007.

[6] Francoise Forges, Antoine Salomon. *Bayesian Repeated Games*. 2013. ffhal-00803919v2.

[7] Ben Polak. *Game Theory: Lecture 1 Transcript*. ECON 159, 5 September 2007, Open Yale Courses.

[8] Tim Roughgarden. *CS364B: Frontiers in Mechanism Design, Lecture #16: The Price of Anarchy in First-Price Auctions*. February 2014.

[9] Christian Grece. *The online advertising market in the EU - Update 2015 and Focus on programmatic advertising*. European Audiovisual Observatory, Strasbourg, 2016.

[10] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. *Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords*. American Economic Review, American Economic Association, vol. 97(1), pages 242-259, March 2007.

[11] Yi Zhu, Kenneth C. Wilbur. *Hybrid Advertising Auctions.* Marketing Science, Vol. 30, No. 2, March–April 2011, pp. 249–273.

[12] Ron Holzman, and Dov Monderer. *Characterization of ex post equilibrium in the VCG combinatorial auctions.* Games and Economic Behavior 47 (2004) 87–103.

[13] Roger B. Myerson. *Optimal Auction Design.* Mathematics of Operations Research II, Vol. 5, No 1, February 1981.

[14] Lawrence M. Ausubel , Paul Milgrom. *The Lovely but Lonely Vickrey Auction.* Combinatorial Auctions, chapter 1, 2006.

[15] K. McLaughlin, D. Friedman. *Online ad auctions: An experiment.* WZB Discussion Paper, No. SP II 2016-501, Wissenschaftszentrum Berlin für Sozialforschung (WZB), Berlin (2016).

[16] Brendan Lucier, Renato Paes Leme, Éva Tardos. *On Revenue in the Generalized Second Price Auction.* January 2011.

[17] Weiran Shen, Binghui Peng, Hanpeng Liu, Michael Zhang, Ruohan Qian, Yan Hong, Zhi Guo, Zongyao Ding, Pengjun Lu, and Pingzhong Tang. *Reinforcement Mechanism Design, with Applications to Dynamic Pricing in Sponsored Search Auctions.* Nov. 2017.

[18] Gang Bai, Zhihui Xie, Liang Wang. *Practical Constrained Optimization of Auction Mechanisms in E-Commerce Sponsored Search Advertising.* Jul. 2018.

[19] Yi Wang, Xuemin Zhao, Zhenlong Sun, Hao Yan, Lifeng Wang, Zhihui Jin, Liubin Wang, Yang Gao, Ching Law, Jia Zeng. *Peacock: Learning Long-Tail Topic Features for Industrial Applications.* ACM Transactions on Intelligent Systems and Technology, Vol. 9, No. 4, Article 39, December 2014.

[20] Matthew Cary, Aparna Das, Benjamin Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin, Claire Mathieu, and Michael Schwarz. *On Best-Response Bidding in GSP Auctions.* NBER Working Paper No. 13788, February 2008.

[21] Deepak Agarwal, Souvik Ghosh, Kai Wei, Siyu You. *Budget Pacing for Targeted Online Advertisements at LinkedIn.* KDD'14, August 24–27, 2014, New York, NY, USA.

[22] Hal R. Varian and C. Harris. *The VCG Auction in Theory and Practice.* The American Economic Review 104, no. 5 (2014): 442-45. http://www.jstor.org/stable/42920977.

[23] Nikolaj Bjorner, Sanjiva Prasad, Laxmi Parida. *Distributed Computing and Internet Technology*. 12th International Conference, ICDCIT 2016, Bhubaneswar, India, January 15-18, 2016, Proceedings.

[24] Cheng Li, Yue Lu, Qiaozhu Mei, Dong Wang, Sandeep Pandey. *Click-through Prediction for Advertising in Twitter Timeline*. KDD 15, August 10-13, 2015, Sydney, NSW, Australia.

[25] Twitter. *Auctions*.
Retrieved from https://business.twitter.com/fr/help/troubleshooting/bidding-and-auctions-faqs.html.

[26] Santiago R. Balseiro, Jon Feldman, Vahab Mirrokni, S. Muthukrishnan. *Yield Optimization of Display Advertising with Ad Exchange*. Management Science, 2014.

[27] Amin Sayedi. *Real-Time Bidding in Online Display Advertising*. Marketing Science, Vol. 37, No. 4, July–August 2018, pp. 553–568.

[28] Muthukrishnan S. *Ad exchanges: Research issues*. Internet and Network Economics, Lecture Notes in Computer Science. 1–12, 2009.

[29] Rui Qin, Yong Yuan, Fei-Yue Wang. *Optimizing the Revenue for Ad Exchanges in Header Bidding Advertising Markets*. 2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC).

[30] Matthew Richardson, Ewa Dominowska, and Robert Ragno. *Predicting Clicks: Estimating the Click-Through Rate for New Ads*. Proceedings of the 16th International World Wide Web Conference(WWW-2007), January 2007.

[31] Kushal Dave, Vasudeva Varma. *Learning the Click-Through Rate for Rare/New Ads from Similar Ads*. 33rd Annual ACM SIGIR Conference, July 2010.

[32] Brendan Kitts, Parameshvyas Laxminarayan, Benjamin LeBlanc, Ryan Meech. *A Formal Analysis of Search Auctions Including Predictions on Click Fraud and Bidding Tactics*. Proceedings of the First Workshop on Sponsored Search, Sixth ACM Conference on E-Commerce, June 2005.

[33] Kuang-chih Lee, Burkay Orten, Ali Dasdan, Wentong L. *Estimating Conversion Rate in Display Advertising from Past Performance Data*. KDD'12, August 12–16, 2012, Beijing, China.

[34] Eric Bax, Anand Kuratti, Preston Mcafee, Julian Romero. *Comparing predicted prices in auctions for online advertising*. International Journal of Industrial Organization 30 (2012) 80–88.

[35] Santiago R. Balseiro, Yonatan Gur. *Learning in Repeated Auctions with Budgets : Regret Minimization and Equilibrium.* June 2018.

[36] Ingrid Van Keilegom. *Nonparametric Estimation of the Conditional Distribution in Regression with Censored Data* (PhD thesis). Limburgs Universitair Centrum, 1998.

[37] Yijian Huang and Thomas A. Louis. *Nonparametric Estimation of the Joint Distribution of Survival Time and Mark Variables.* Biometrika, Vol. 85, No. 4 (Dec., 1998), pp. 785-798.

[38] Jacobo de Uña Alvarez and Luis F. Meira Machad. *A Simple Estimator of the Bivariate Distribution Function for Censored Gap Times.* July 2006.

[39] R. B. Bapat and M. I. Beg. *Order Statistics for Nonidentically Distributed Variables and Permanents.* Sankhyā: The Indian Journal of Statistics, Series A (1961-2002), Vol. 51, No. 1 (Feb.,1989), pp. 79-93.

[40] Donatien Hainaut. *LACTU 2060 : Life Insurance II.* 2018-2019.

[41] Daowen Zhang. *Likelihood and Censored (or Truncated) Survival Data.* Chapter 3, ST 745.

[42] In Jae Myung. *Tutorial on maximum likelihood estimation.* Journal of Mathematical Psychology 47 (2003) 90–100.

[43] Gideon E. Schwarz. *Estimating the dimension of a model.* Annals of Statistics, vol. 6, no 2, 1978, pp. 461-464.

[44] Chris T. Volinsky, Adrian E. Raftery. *Bayesian Information Criterion for Censored Survival Models.* Office of Naval Research Grant, July 1998.

[45] L. Dirick, G. Claeskens, B. Baesens. *An Akaike information criterion for multiple event mixture cure models.* May 2014.

[46] T. W. Anderson, D. A. Darling. *A Test of Goodness of Fit.* Journal of the American Statistical Association, Vol. 49, No. 268. (Dec., 1954), pp. 765-769.

[47] Azzah Mohammad Alharpy, Noor Akma Ibrahim. *Parametric Tests for Partly Interval-censored Failure Time Data under Weibull Distribution via Multiple Imputation.* Journal of Applied Sciences, 13: 621-626, 2013.

[48] William H. Greene. *Econometric Analysis (6th ed.).* Prentice Hall.

[49] E. L. Kaplan and Paul Meier. *Nonparametric Estimation from Incomplete Observations*. Journal of the American Statistical Association, Vol. 53, No. 282 (Jun., 1958), pp. 457-481.

[50] Aleksandra Baszczyńska. *Kernel Estimation of Cumulative Distribution Function of a random variable with bounded support*. Statistics in Transition new series, September 2016, Vol. 17, No. 3, pp. 541–556.

[51] Bruce W. Turnbull. *Nonparametric Estimation of a Survivorship Function with Doubly Censored Data*. Journal of the American Statistical Association, March 1974, Volume 69, No 345, Theory and Methods Section.

[52] Bruce W. Turnbull. *The Empirical Distribution Function with Arbitrarily Grouped, Censored, and Truncated Data*. July 1976.

[53] Michael P. Fay. *Interval Censored Data Analysis*. Tutorial:Interval Censored Data Analysis, July 2010.

[54] Robert Gentleman, Charles J. Geyer. *Maximum Likelihood for Interval Censored Data: Consistency and Computation*. Biometrika, Vol. 81, No. 3. (Aug., 1994), pp. 618-623.

[55] Michael P. Fay, Pamela A. Shaw. *Exact and Asymptotic Weighted Logrank Tests for Interval Censored Data: The interval R package*. April 2009.

[56] Jon A. Wellner. *Log-concave distributions: definitions, properties, and consequences*. January 2012.

[57] Josef Stoer, Christoph Witzgall. *Convexity and Optimization in Finite Dimensions I*. Grundlehren der mathematischen Wissenschaften, no 163, 1970, p.172, th. 4.9.11.

[58] András Prékopa. *On Probabilistic Constrained Programming*. Proceedings of the Princeton Symposium on Mathematical Programming, Princeton University Press, Princeton, NJ 1970. pp. 113–138.

[59] Lutz Dümbgen, Kaspar Rufibach. *Maximum likelihood estimation of a log-concave density and its distribution function: Basic properties and uniform consistency*. Bernoulli 15(1), 2009, 40–68, DOI: 10.3150/08-BEJ141.

[60] D. Schuhmacher, K. Rufibach, L. Dümbgen. *Logconcens: maximum likelihood estimation of a log-concave density based on censored data*. Retrieved from http://CRAN.R-project.org/package=logconcens (2013).

[61] F.N. Fritsch, R.E. Carlson. *Piecewise Cubic Interpolation Methods*. November 1978.

[62] D.A. Darling. *The Kolmogorov-Smirnov, Cramer-Von Mises Tests*. Ann. Math. Statist., Volume 28, Number 4 (1957), 823-83.

[63] David G. Glynn. *The permanent of a square matrix*. European Journal of Combinatorics 31 (2010) 1887–1891.

[64] Sergey Foss, Dmitry Korshunov and Stan Zachary. *An Introduction to Heavy-tailed and Subexponential Distributions*. Mathematisches Forschungsinstitut Oberwolfach gGmbH Oberwolfach Preprints (OWP), OWP 2009 - 13.

[65] Barthel W. Huff. *Mixtures and Order Statistics*. The American Mathematical Monthly, Vol. 99, No. 3 (Mar., 1992), pp. 239-242.

[66] Tia Fisher. *ROI in social media: A look at the arguments*. Journal of Database Marketing & Customer Strategy Management, September 2009, Volume 16, Issue 3, pp 189–195.

[67] Weinan Zhang, Yifei Rong, Jun Wang, Tianchi Zhu, Xiaofan Wang. *Feedback Control of Real-Time Display Advertising*. WSDM 2016, February 22-25, 2016, San Francisco, CA, USA.

[68] James P. Quirk and Rubin Saposnik. *Admissibility and Measurable Utility Functions*. The Review of Economic Studies, Vol. 29, No. 2 (Feb., 1962), p. 141.

[69] Stephen Boyd, Jaehyun Park. *Subgradient Methods*. Notes for EE364b, Stanford University, Spring 2013–14.

[70] Babinec et. al. *Introduction to Topology*. Mathematics 490 – Introduction to Topology, Lecture notes.

[71] René Henrion. *Optimization Problems with Probabilistic Constraints*. 10th International Conference on Stochastic Programming.

[72] Patrick Billingsley. *Probability and Measure (3rd ed.)*. John Wiley & Sons, p. 357, 1995.

[73] François Glineur et al.. *LINMA2471: Optimization models and methods II, Course notes*. January 2010.

[74] Harvey J. Greenberg, William P. Pierskalla. *A Review of Quasi-Convex functions*. Operations Research, Vol. 19, No. 7, N

[75] Stephen P. Boyd, Lieven Vandenberghe. *Convex Optimization.* EE364a: Convex Optimization I, Lecture 2.

[76] Changyu Wang, Naihua Xiu. *Convergence of the Gradient Projection Method for Generalized Convex Minimization.* Computational Optimization and Applications, 16, 111–120, 2000.

[77] Yaohua Hu, Carisa Kwok Wai Yu, Xiaoqi Yang. *Incremental Subgradient Methods for Minimizing The Sum of Quasi-convex Functions.* October 2017.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain
Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | **www.uclouvain.be/epl**