

École polytechnique de Louvain

Reserve price optimization in sponsored search auctions

A multi-stage stochastic programming approach

Author: **João Gabriel SOARES ACCORSI**
Supervisors: **Raphaël JUNGERS, Andre WIBISONO**
Readers: **François GLINEUR, Débora PRETTI RONCONI**
Academic year 2021–2022
Master [120] in Mathematical Engineering

Abstract

Sponsored search advertising, a technique that displays one or more ads in search engines such as Google or Facebook whenever someone searches for the products or services of an advertiser, has a remarkable and increasing economic importance. In sponsored search ads, auction mechanisms are usually used for selling the ad slots. When observing the problem from the viewpoint of the search engines (also known as publishers), setting a reserve price is the main mechanism through which they can influence their revenues in the auctions. While there is a relevant body of literature related to the best way to choose reserve prices in sponsored search advertising, there are practical constraints that make the problem more complex.

This thesis proposes a model to optimize the reserve price in single-item and generalized second-price auctions, two of the main types of auctions used in sponsored search advertising, while taking the advertisers' budgets into account. Assuming that publishers are able to learn bid distributions of the advertisers participating in the auctions, the model is formulated as a multi-stage stochastic mixed-integer linear program (MS-SMILP), whose objective is to maximize the publisher's revenue in a set of subsequent auctions. It is proposed that, for the case of a single-item second-price auction, the revenue resulting from this model is always greater or equal to the revenue when using four other approaches to choose the reserve price. Moreover, still for the single-item case, the performance of the model is analyzed and compared to the same four approaches by a numerical example and simulations, in which the solution of the proposed model led to significantly higher revenues.

Acknowledgements

I would like to take this opportunity to thank all the people who have contributed to this thesis.

First, I am extremely grateful to my supervisor Raphaël Jungers for his invaluable patience and feedback. His constructive comments have helped me a lot and I could not have undertaken this journey without his assistance and support.

I would also like to express my deepest gratitude to Andre Wibisono, who was also my supervisor. Accepting me to work with him, he contributed to an unforgettable experience and generously provided me with knowledge and expertise.

I would like to extend my sincere thanks to François Glineur and Débora Pretti Ronconi for agreeing to be part of my thesis' jury despite their many other academic activities.

Finally, I would like to thank my friends and family, especially my parents. Their continued support and belief in me have kept my spirits high during this process.

Contents

1	Introduction	1
1.1	Context	1
1.2	State of the art	4
1.3	Contributions	9
1.4	Structure	10
2	Mathematical background	11
2.1	Game Theory	11
2.1.1	Basic concepts	12
2.1.2	Auctions as repeated Bayesian games	17
2.1.3	Auctions with reserve price	23
2.2	Stochastic programming	27
2.2.1	Two-stage stochastic programming	27
2.2.2	Multi-stage stochastic programming	33
3	A new multi-stage stochastic MILP approach	41
3.1	Auction rules and details	41
3.2	Model formulation	45
3.2.1	Second-price auction	45
3.2.2	GSP auction	50
3.3	Analysis of the model	52
4	Numerical example and simulations	57
4.1	Numerical example	57
4.1.1	Characteristics	58
4.1.2	Performance	59
4.1.3	Comparisons	64
4.2	Simulations	66
4.2.1	Performance	67
4.2.2	Comparisons	68

5	Conclusion	71
5.1	Summary	71
5.2	Future directions	72
	Appendices	75
A	Nomenclature of the optimization model	77
A.1	Second-price auction	77
A.2	GSP auction	78
B	Results of the numerical example for each scenario	79
C	Generation of parameters for simulations	81

List of Figures

1.1	Sponsored ads for keyword "sunglasses" on Google.	2
1.2	Net digital ad revenues of the biggest publishers.	3
1.3	Agents involved in sponsored search advertising and their relationships.	4
1.4	Timeline for reserve price optimization in sponsored search auctions.	8
2.1	Sequence of events in two-stage stochastic linear programs.	28
2.2	Relations between WS , SP , and EEV	33
2.3	Graphical representations of a scenario tree and a lattice.	35
3.1	Scenario tree of a (single-item) second-price auction.	42
3.2	Scenario tree of a GSP auction.	44
4.1	Scenario tree with the chosen parameters (numerical example).	58
4.2	Optimal value of each solution (numerical example).	63
4.3	Number of simplex iterations for each solution (numerical example).	63
4.4	Performance of the numerical example.	64
4.5	Average revenue for each approach (numerical example).	66
4.6	Average number of simplex iterations (simulations).	67
4.7	Average solutions (simulations).	67
4.8	Average performance (simulations).	68
4.9	Average revenue for each approach (simulations).	69
B.1	Revenue for each approach (scenario 4.1).	79
B.2	Revenue for each approach (scenario 4.2).	79
B.3	Revenue for each approach (scenario 4.3).	79
B.4	Revenue for each approach (scenario 4.4).	79
B.5	Revenue for each approach (scenario 4.5).	80
B.6	Revenue for each approach (scenario 4.6).	80
B.7	Revenue for each approach (scenario 4.7).	80
B.8	Revenue for each approach (scenario 4.8).	80

List of Tables

4.1	Budgets and possible bids for advertisers A and B.	58
4.2	SP solution and number of simplex iterations (numerical example). . .	59
4.3	Optimal reserve prices of the SP solution (numerical example). . . .	59
4.4	WS solution and number of simplex iterations (numerical example). . .	60
4.5	EV solution and number of simplex iterations (numerical example). . .	63
4.6	EEV_4 solution and number of simplex iterations (numerical example). .	63
4.7	Reserve prices for the second approach (numerical example).	65
4.8	Reserve prices for the third approach (numerical example).	65
4.9	Reserve prices for the fourth approach (numerical example).	65
C.1	Parameters generation (2 scenarios).	81
C.2	Parameters generation (4 scenarios).	81
C.3	Parameters generation (8 scenarios).	82

Acronyms

CTR Click-Through Rate

DSIC Dominant Strategy Incentive Compatible

EEV Expected value of using the EV solution

EV Expected Value

EVPI Expected Value of Perfect Information

FMFE Fluid Mean-Field Equilibrium

GSP Generalized Second-Price

MILP Mixed-Integer Linear Programming

MIP Mixed-Integer Programming

MPEC Mathematical Problem with Equilibrium Constraints

MS-SLP Multi-Stage Stochastic Linear Programming

MS-SMILP Multi-Stage Stochastic Mixed-Integer Linear Programming

SLP Stochastic Linear Programming

SMILP Stochastic Mixed-Integer Linear Programming

SP Stochastic Program (Here-and-Now)

VCG Vickrey-Clarke-Groves

VSS Value of the Stochastic Solution

WS Wait-and-See

1 | Introduction

This chapter's goal is to introduce the problems tackled in the course of this thesis. First, the general context of sponsored search ads is introduced. It shows the main players of this market and how economically important it is, as well as which types of auctions are the most used in this context and the agents involved.

Then, the state of the art is presented. Different problems studied in the literature from the viewpoints of both the advertisers and the publishers are shown. Focusing on the problem of optimizing the reserve price (from the point of view of the publishers), the main recent developments are mentioned and summarized in a timeline.

After that, the main contributions of this thesis to the state of the art are listed. Finally, the general structure of the thesis is shown to the reader, with the content of each of the subsequent chapters.

1.1 Context

According to the Interactive Advertising Bureau (IAB) report, internet advertisement revenues in the United States totaled \$139.8 billion in 2020. Even with the impact of the COVID-19 pandemic, it represented a double-digit growth of 12.2% in overall internet advertising revenues between 2019 and 2020. From this amount, search advertisements were responsible for \$59.0 billion, up from the \$54.7 billion reported in fiscal year 2019. [1] Also in 2020, a total of \$378.16 billion was spent in digital advertising worldwide. [2]

Given the notable and increasing economic importance of search advertisements, the literature on the topic is vast. Different types of search advertisements are studied,

in which the problem can be seen from the perspective of each player involved in the search ad ecosystem. Among the various topics studied, there is a relevant body of literature on sponsored search advertising, a technique that displays one or more ads in search engines such as Google or Facebook whenever someone searches for the products or services of an advertiser, matching the searched query (see the screenshot created on March 1st, 2022, for searching the keyword "sunglasses" on Google in Figure 1.1). It is also present in e-commerce platforms, such as Amazon and eBay, where the order in which the products are displayed for the consumers is a critical decision for them. In sponsored search ads, auction mechanisms are usually used for selling ad slots, therefore making the topic closely related to Game Theory. Sponsored search auctions are then the primary way through which companies like Google and Facebook monetize their search engines.




Figure 1.1: Sponsored ads for keyword "sunglasses" on Google.

There are several types of auctions used in sponsored search advertisements. Each auction differ from the others by how the bids are ranked and what the winners pay, for example. Among these different types, a tendency for the use of some of them by the big players is observed. The auction type and pricing scheme is generally not kept as secret by the publishers and a great deal has been written about it. In Figure 1.2, the biggest players are ranked by their net digital ad revenues in 2021 and it also shows their projections for the subsequent years. It is clear that Google and Facebook dominate the market. However, publishers like Alibaba, Amazon, and Tencent play an important role in this market and must be also taken into account.

Net Digital Ad Revenues Worldwide, by Company, 2021-2023			
<i>billions</i>			
	2021	2022	2023
Google	\$146.12	\$171.34	\$186.27
—YouTube	\$14.61	\$18.32	\$21.71
Facebook	\$114.32	\$135.14	\$155.25
—Instagram	\$46.42	\$60.45	\$74.91
Alibaba*	\$39.71	\$47.69	\$55.56
Amazon	\$31.54	\$43.05	\$55.96
Tencent	\$13.46	\$16.18	\$19.10
Microsoft	\$10.06	\$11.64	\$13.03
—LinkedIn	\$4.66	\$5.78	\$6.88
Baidu	\$9.80	\$9.87	\$9.94
JD.com	\$6.48	\$7.76	\$9.19
Kuaishou	\$6.18	\$8.65	\$11.14
Yahoo	\$4.84	\$5.03	\$5.15
Meituan	\$4.39	\$5.92	\$7.62
Twitter	\$4.35	\$5.28	\$6.16
Snapchat	\$3.34	\$4.65	\$6.45
Pinterest	\$2.62	\$3.42	\$4.34
Sina	\$2.12	\$2.17	\$2.20
Pandora	\$1.43	\$1.44	\$1.46
Spotify	\$1.25	\$1.50	\$1.81
Yelp	\$1.00	\$1.10	\$1.16
IAC	\$0.79	\$0.83	\$0.88
Reddit**	\$0.32	\$0.46	\$0.57
Other	\$87.60	\$88.07	\$94.92
Digital ad revenues	\$491.70	\$571.16	\$648.15

*Note: includes advertising that appears on desktop and laptop computers as well as mobile phones, tablets, and other internet-connected devices, and includes all the various formats of advertising on those platforms; net ad revenues after company pays traffic acquisition costs (TAC) to partner sites; *includes ad revenues from Alibaba's core commerce operations and Youku Tudou; **excludes nonadvertising revenues (e.g., Reddit Premium, Reddit Coins)*
Source: eMarketer, October 2021

T11651 eMarketer | InsiderIntelligence.com

Figure 1.2: Net digital ad revenues of the biggest publishers. [3]

The main mechanism used by the biggest players seems to be the generalized second-price (GSP) auction. There are indications that the GSP auction is used by Google (see [4] and [5]), Alibaba (see [6]), Tencent (see [7]), Microsoft (see [5]), Baidu (see [8]), Yahoo! (see [9]), and LinkedIn (see [10]). The GSP auction is a generalization of the second-price auction for the case where multiple items are auctioned. The single-item second-price auction is used by Twitter (see [11]). Moreover, there are evidences that Amazon uses a modified version of the GSP auction (see [12]).

Another mechanism highlighted here is the Vickrey-Clarke-Groves (VCG) auction, which seems to be used by the second main player of the market: Facebook (see [13]). However, the GSP is still the most widely used auction mechanism in sponsored search advertising and is then broadly studied in the literature. The focus of this thesis will be on generalized second-price auctions and its single-item version.

The sponsored search ad market (independently of its mechanism) is a two-side market.

Advertisers procure ad impressions on publishers' websites, aiming to reach potential consumers. On the other side, the publishers provide ad slots to the advertisers with the highest valuation for these slots. Moreover, one can see a third (indirect) agent: the users, who search combinations of words in the publishers' platforms and for whom the ads are shown. The relationships between these agents are illustrated in Figure 1.3.



Figure 1.3: Agents involved in sponsored search advertising and their relationships.

In this thesis, the focus is set to the viewpoint of the publishers. In order to maximize their revenues in the auctions, publishers often set up the so-called *reserve price*, the minimum price for a bid to be considered in an auction. Indeed, the reserve price is the main mechanism through which the auction revenue can be influenced by a publisher (see [14]) and is used by the biggest player of the digital ad market, Google. [15] This way, this thesis' objective is to contribute to the body of literature on reserve price optimization in (single-item and generalized) second-price auctions, proposing and analyzing a new approach to find optimal reserve prices while maximizing the publishers' revenues in the auctions.

1.2 State of the art

Given the importance of sponsored search ads nowadays, it is not surprising that several studies related to this topic have been done. These studies involve a very diverse range of fields, such as Information Systems, Marketing, Economics, Operations Research, and Computer Science. When observing the literature related specifically to the auctions behind sponsored search ads, an increasing importance has been given to combining Machine Learning and Data Mining techniques with game-theoretic approaches.

As we have just seen, the advertisers and the publishers are the two main players involved in sponsored search auctions. Most studies on the subject focus on the viewpoint of one of these players. When observing the literature on the advertisers' viewpoint, most part of it is related to optimizing their bidding strategies. Other studies involve, for example, the choice of buying ad slots via guaranteed contracts versus non-guaranteed selling channels (see [16]) and how privacy policies should be set when observing the consumers' behaviour in consideration with their welfare (see [17]). When analyzing the advertisers' optimal bidding strategy in second-price auctions, in theory, the weakly dominant strategy for an advertiser is bidding truthfully.¹ However, in real situations, there are practical constraints that make the bid calculation more complex:

- The advertisers have budget constraints for a given campaign in repeated auctions; [18]
- Advertisers may learn their own and/or others' true valuations² over time; [19]
- Advertisers may set a number of impressions to attain; [20]
- Publishers often offer the possibility of setting pacing options so that the budget is spent smoothly. [21]

This way, as the theoretical dominant strategies usually do not cover these practical constraints, Machine Learning techniques (involving, for example, Deep Learning and Reinforcement Learning) have recently gained a lot of importance when trying to maximize the advertisers' utilities, to find out their optimal pacing strategies, and to learn other advertisers' bids.

When analyzing sponsored search auctions from the viewpoint of the other players directly involved on them, the publishers, there are other problems which have been studied in the literature:

- What information should be shared with the advertisers; [22]
- Keyword suggestion; [23]
- Query expansion and ad selection; [24]
- Prediction of the ads' clicks; [25]
- Choosing the auction mechanism; [26]
- Choosing the optimal reserve price.

When trying to tackle these problems, studies involve a variety of fields including Artificial Intelligence, Machine Learning, Data Mining, and Game Theory. As stated in

¹See Definition 2.1.6 of weakly dominant strategy. See also Subsection 2.1.2 for a further discussion on why bidding truthfully is a weakly dominant strategy in second-price auctions.

²See Definition 2.1.2 of a bidder's true valuation.

Section 1.1, the focus of this thesis will be on the last of these problems: optimizing the reserve price.

Since Yahoo! had no longer fixed its minimum bid for all its sponsored search ads as \$0.10 in early 2008 and used variable reserve prices for some sets of keywords (see [15]), the optimization of the reserve price became target of different studies. Edelman and Schwarz, in [27], stated and proved that a generalized English auction³ with reserve price is an optimal mechanism if its reserve price r^* solves $r^* - (1 - \frac{F(r^*)}{f(r^*)})$, where $F(\cdot)$ and $f(\cdot)$ refer to the probability distribution and density functions of the bidders' willingness to pay, respectively. This result means that an English auction is independent of the number of bidders, number of slots being auctioned, and the rate of decline of the click-through rate (CTR)⁴ from a position to another. Limitations for the use of this simple formula and drawbacks when trying to use it in more complex contexts were shown in [28]: it is assumed that advertisers bid exactly at their (predicted) private values and they do not consider practical constraints for achieving the formula. These assumptions are strong and not observed in real-life auctions. When considering more complex situations, there is no simple formula describing the optimal reserve price in English auctions (nor in second-price auctions).

Cesa-Bianchi, Gentile, and Mansour state that, at the best of their knowledge, they were the first ones to derive formal and concrete convergence rates for approximating optimal reserve prices in a second-price auction in [29]. They focus on the eBay model, where the publisher only stores the price paid by the winning advertiser (the maximum between the second-highest bid and the reserve price). They assume a fixed but unknown prior distribution and that bidders' utilities are independent and identically distributed, proposing learning algorithms for both known and unknown numbers of bidders. Their algorithms work in stages $s = 1, \dots, S$, each one with T_s time steps, and aim an efficient exploration of the outcomes of each auction. The extension of these algorithms for GSP auctions, however, is not straightforward, being established as one of the authors' suggestions for future works.

Other authors have also used learning algorithms to optimize the reserve price in (single-item) second-price auctions. Mohri and Medina, in [14], achieve some meaningful advances compared to [29], for example considering that companies do have access to the past highest bids, and not only to the price paid (which in most cases, including the Google case, it is indeed more realistic). Assuming that the advertisers bid their true valuations of the ad, the proposed method aims to learn the revenue of a subsequent auction given the outcomes of a set of previous auctions. The stated problem, however, involves non-convex optimization problems in higher dimensions and its algorithm only guarantees convergence to a local minimum. Moreover, the authors could not test their

³An English auction can be seen as a second-price auction with private values.

⁴See Definition 2.1.15 of click-through rate.

algorithm with real-world data.

Other studies tried to extend the problem to generalized second-price auctions. Mohri and Medina were, to the best of their knowledge, the first ones to attempt to apply learning algorithms to the problem of reserve price optimization in GSP auctions in [30]. They presented two algorithms: one estimating the density distribution of the bids and other learning the publisher's loss function, and both of them depend on advertisers playing on equilibrium to guarantee convergence to optimality.

A learning algorithm to find the optimal reserve price in GSP auctions was also proposed in [31]. The authors assume that the advertisers know the CTR of each ad position to propose two algorithms. The first one, which estimates a unique reserve price (closer to reality), depends on the bidders to follow the locally envy-free equilibrium bidding strategy. The second one estimates different reserve prices for each slot position (further from reality), and does not depend on advertisers following the locally envy-free equilibrium bidding strategy. These algorithms have a lot in common with the ones proposed in [29] and [14] for second-price auctions with one item. A drawback of these algorithms is that their convergence to optimality can be slow depending on the number of time periods of learning.

One of the limitations of all the mentioned studies is that they do not take into account some advertisers' practical constraints. The most common of these constraints is the advertisers' budget constraint. Even when observing the problem from the publisher's viewpoint, these constraints play an important role. First, the budget may change the behaviour of the advertisers. This problem is often tackled by introducing a family of adaptive pacing strategies, in which advertisers adjust the pace at which they spend their budget according to their expenditures (see [32] and [33]). Secondly, setting the reserve price such that advertisers spend the most in one auction can be a nonoptimal strategy, as they may not have budget left to spend in subsequent auctions (see [30] and [34]).

In [18], Balseiro, Besbes, and Weintraub consider that the advertisers' budget constraints affect the auctions over time. Therefore, traditional equilibrium and revenue optimization analysis for static auctions do not apply in this case. To approximate the strategic interactions among budget-constrained bidders, they use the concept of fluid mean-field equilibrium (FMFE). In this work, the authors prove that there exists an FMFE equilibrium for the advertisers' bids. This equilibrium is used to formulate a mathematical program with equilibrium constraints (MPEC) for the publishers' problem of finding a reserve price that maximizes their revenue. For homogeneous advertisers, it finds a closed-form solution for the reserve price. However, this is not possible if the advertisers are heterogeneous.

In [20], one relevant practical constraint for the advertisers is taken into account: the minimum impression constraint for them to bid. It indeed shows that the behaviour

of the advertisers' strategies are consistent with this constraint. It uses the FMFE framework developed in [18] to recover advertisers' valuations, and includes the minimum impression constraint in the model. It also extends the results to the case where budget constraints are considered in a similar way.

In [35], some mechanisms for controlling the expenditure (due to the advertisers' budgets) are analyzed and compared from both the points of view of the publishers and the advertisers. From the publishers' perspective, it was shown that the best mechanism was to exclude advertisers from some auctions and, in addition, set reserve prices. Doing that, however, drastically changes the rules of the auction, which then cannot be considered a GSP auction anymore. Allowing the publisher to exclude advertisers from some auctions even if they match all the requirements to take part in it would make its rules and fairness questionable.

As we have just seen, a thorough investigation into the body of literature yielded several studies on optimizing the reserve price in second-price auctions, involving different fields such as Machine Learning, Statistical Inference, and Game Theory. However, none of these studies provided a method for this optimization when the advertisers are heterogeneous and their budget constraints are taken into account, while straightforwardly extending it for GSP auctions and maintaining the main characteristics of the mechanism. A non-exhaustive timeline of the studies on optimizing the reserve price is illustrated in Figure 1.4.




Figure 1.4: Timeline for reserve price optimization in sponsored search auctions. To the best of our knowledge, this thesis is the first work to propose a model for GSP auctions while considering budget constraints and heterogeneous advertisers.

There are also other open issues when optimizing the reserve price in sponsored search ads. For instance, solving the problem while considering competition with other publishers (whose inventory can be seen as an imperfect substitute for the advertiser)

and developing algorithms that implement the reserve price at scale into the auctions (see [36]) are other problems to be considered.

1.3 Contributions

This thesis is devoted to modelling the optimization of the reserve price in (single-item and generalized) second-price auctions from the publishers' perspective, while taking the budget constraints of (heterogeneous) advertisers into account. To do that, a multi-stage stochastic mixed-integer linear programming (MS-SMILP) problem is formulated and applied when solving it for an example and simulations.

The main contributions of this thesis are:

- Providing a model for optimizing the reserve price in single-item second-price auctions while considering the budget constraints of heterogeneous advertisers;
- Extending this model for GSP auctions;
- To the best of our knowledge, coming up with the first multi-stage stochastic programming model for optimizing the reserve price in sponsored search auctions;
- Proposing the first MIP (and MILP, consequently) formulation for this specific problem, to the best of our knowledge;
- By presenting four novel propositions, the benefits of using these models are translated into inequalities showing that they always lead to greater or equal revenues when compared to four other approaches to choose the reserve price;
- Comparing the performance of the proposed model with other approaches to choose the reserve price by detailing a numerical example and simulations' results.

It is important to highlight that, albeit not used for optimizing the reserve price in second-price auctions, MIP and multi-stage stochastic programming models are seen in the literature on sponsored search auctions when tackling other problems. In [33], an MIP formulation is provided for finding advertisers' adaptive pacing strategies for both finding a social-welfare-maximizing equilibrium and a revenue-maximizing equilibrium. Two-stage (and multi-stage) optimization models are used, for example, for display-ad allocation problems (see [37]), and widely applied for modelling auctions in other markets (see [38], [39], [40], [41], [42], and [43]).

1.4 Structure

Apart from this introduction, this thesis is divided in four other chapters. Chapter 2 is devoted to presenting the mathematical background that is relevant for the understanding of the subsequent chapters, including concepts and notions of Game Theory and Stochastic Programming. In Chapter 3, the models proposed in this thesis, as well as their main assumptions, are introduced and analyzed. Still in this chapter, four propositions compare the model for the single-item case with other benchmarks. In Chapter 4, a numerical example is presented in details and the main results of simulations of the single-item model are shown. The results of the example and the simulations are analyzed and compared to other approaches to choose the reserve price. Finally, this work is concluded in Chapter 5.

2 | Mathematical background

This chapter presents the main mathematical background necessary for understanding the models proposed in this thesis. It is divided in two sections: the first one about Game Theory and the second one regarding Stochastic Programming.

The first section focuses on detailing the game-theoretic models of the auctions considered in this thesis. It starts with basic concepts which allow the reader to understand the subsequent contents. Then, it shows how second-price and GSP auctions are formulated as repeated Bayesian games, a common approach to model these auctions in the sponsored search advertising context. Finally, it details how the addition of reserve prices changes the mechanism of these auctions.

The second section of this chapter concerns concepts of Stochastic Programming which enable the reader to comprehend the models presented in this thesis. An overview of two-stage stochastic programs is presented, followed by its extension to the multi-stage case. For both cases, their formulations are detailed, as well as how their performances are evaluated.

2.1 Game Theory

The auctions analyzed in this thesis are based on game-theoretic auction models. To fully understand the models considered in the next chapters, some concepts of Game Theory must be presented and explained. This section is devoted to understanding the main principles and concepts of Game Theory involved in sponsored search auctions and showing how some of the main types of auctions in this framework are described as repeated Bayesian games. The awarded book "*Game Theory: Analysis of Conflict*" from Myerson in [44], course notes from Jungers et al. in [45], the work from Nisan et

al. in [46], and the lectures from Roughgarden in [47] were the main references for the definitions shown in the course of this section.

2.1.1 Basic concepts

First of all, it is important to understand the concept of auction from a game-theoretic point of view. The following definition was given in [45].

Definition 2.1.1 *Auction*

An auction is a structured framework for negotiation. Each such negotiation has certain rules which must be specified:

- (i) bidding rules (how are offers made?);
- (ii) clearing rules (when does the trade occur or what are these trades, i.e. which player gets the good, and what are the prices for the players?);
- (iii) information rules (who knows what about the state of negotiation?).

The auctions considered in this thesis are those in which a single *seller* auctions items to different *buyers* (also referred to as *bidders*). Those are the *agents* or *players* of the auction. The underlying assumption made when modelling auctions is that each bidder has her own *true value* for each item auctioned, whose definition from [45] is shown as follows.

Definition 2.1.2 *Bidder's true value (true valuation)*

The bidder's true value is the bidder's intrinsic value for the item being auctioned. She is willing to purchase the item for a price up to this value, but not for any higher price.

Each player has a utility function that depends on her true valuation of the item and on other factors of the auction. Regarding the utility function, the games considered from now on are those with *independent private values*, defined as in [46].

Definition 2.1.3 *Game with independent private values*

A game has independent private values if the utility of its players depends fully on their private information and not on any information of others as it is independent from their own information.

An important concept is the one of *Bayesian equilibrium*, as auctions will be formalized as *Bayesian games*. To present these concepts, it is necessary to first define a game with strict incomplete information, which is introduced as done in [46].

Definition 2.1.4 *Game with (independent private values and) strict incomplete information*

A game with (independent private values and) strict incomplete information for a set of n players is given by:

- (i) A set of *actions* X_i for every player i ;
- (ii) A set of *types* T_i for every player i . A value $t_i \in T_i$ is the private information that i has;
- (iii) A *utility function* (or *payoff function*) $u_i : T_i \times X_1 \times \cdots \times X_n \rightarrow \mathcal{R}$, where $u_i(t_i, x_1, \dots, x_n)$ is the utility achieved by player i , if her type (private information) is t_i , and the profile of actions taken by all players is $x_1, \dots, x_n$.

The main idea of this definition is that each player i must choose her action x_i when knowing t_i but not the other t_j 's. For all possible settings of an auction, the total behaviour of a player is represented by her *strategy*, which is defined as in [46].

Definition 2.1.5 *Strategy of a player*

A strategy of a player is a function

$$s_i : T_i \rightarrow X_i,$$

which specifies which action x_i is taken for every possible type t_i .

An important concept regarding the strategy of the players is the one of domination.

The definition of *weakly dominant strategies* shown as follows originates from [46].

Definition 2.1.6 *Weakly dominant strategy*

A strategy s_i is a weakly dominant strategy if, for all t_i , all x_{-i} , and all x'_i :

$$u_i(t_i, s_i(t_i), x_{-i}) \geq u_i(t_i, x'_i, x_{-i}).$$

Similarly, *strongly dominant strategies* can be defined.

Definition 2.1.7 *Strongly dominant strategy*

A strategy s_i is a strongly dominant strategy if, for all t_i , all x_{-i} , and all x'_i :

$$u_i(t_i, s_i(t_i), x_{-i}) > u_i(t_i, x'_i, x_{-i}).$$

Games with strict incomplete information consider that each player has no information at all about the private information of other players (not even a prior distribution). However, the usual working definition in economic theory uses a Bayesian approach, assuming some commonly known prior distribution. This leads to the concept of Bayesian games, whose definition is shown as follows and comes from [45].

Definition 2.1.8 *Bayesian game*

A Bayesian game is a tuple

$$\Gamma^b = (N, X, T, u, s, p),$$

where

- N is the non-empty set of players;
- $X = (X_i)_{i \in N}$ is the set of all actions for all players;
- $T = (T_i)_{i \in N}$ is the set of all player types;
- $u = (u_i)_{i \in N}$ is the set of all utility functions (see Definition 2.1.4);
- $s = (s_i)_{i \in N}$ is the set of all strategies for all players;
- $p = (p_i)_{i \in N}$ are *belief functions*, $p_i : T_i \rightarrow \Delta(T_{-i})$.

A Bayesian game is finite if N , C , and T are finite. As mentioned earlier in this section, sponsored search auctions are often modelled as repeated Bayesian games, whose definition is shown as follows and is adapted from [48].

Definition 2.1.9 *Repeated Bayesian game*

A repeated Bayesian game is a tuple

$$\Gamma^r = (N, \Theta, X, S, T, u, p, s, q),$$

where

- N , X , T , u , s , and p are defined as before (see Definition 2.1.8);
- Θ is the set of *states of the world*;
- $S = (S_i)_{i \in N}$ is the set of all signals that the players can receive;
- q is the *transition function* between each distribution for each player.

When dealing with a Bayesian game setting, it is possible to observe the use of the term *mechanism*. The mechanism of a Bayesian game is a crucial concept, and its definition shown in [45] is presented.

Definition 2.1.10 *Mechanism (for a Bayesian game setting)*

For a Bayesian game setting, a mechanism is a triple

$$(X, \chi, \rho),$$

where

- $X = (X_i)_{i \in N}$ is the set of all actions for all players;
- $\chi : X \rightarrow \Delta(\mathcal{A})$ maps each action profile to a distribution over choices;
- $\rho : X \rightarrow \mathbb{R}^n$ maps each action profile to a payment for each agent.

An important property of a mechanism is whether it is *Dominant Strategy Incentive Compatible* (DSIC) or not. The definition of a DSIC mechanism from [46] is shown as follows.

Definition 2.1.11 *Dominant Strategy Incentive Compatible (DSIC) mechanism*

A mechanism is called Dominant Strategy Incentive Compatible (or *truthful*, or *strategyproof*) if truth-telling is a weakly dominant strategy.

In other words, a mechanism is DSIC if a player fares best or at least not worse by bidding her true valuation of the item, regardless of the other players' strategies. Besides the domination, another concept is important when analyzing the strategies of an auction: the equilibrium. Within this context, the *Nash equilibrium* plays an important role and is indeed one of the most important concepts in Game Theory, considering the situation where the actions of all the players are consistent with the fact that they are rational and intelligent.

Definition 2.1.12 *Nash equilibrium*

Given any strategic-form game $\Gamma = (N, X, u)$, a strategy profile $s = (s_i)_{i \in N}$ is a Nash equilibrium of Γ if and only if

$$u_i(s) \geq u_i(s_{-i}, \tau_i) \quad \forall i \in N, \quad \forall \tau_i \in \Delta(X_i),$$

where $u_i(s_{-i}, \tau_i)$ is the strategy profile where the i^{th} action is τ_i and all the other actions are as in s .

Definition 2.1.12 states that s is a Nash equilibrium if and only if no player could increase her expected payoff by unilaterally deviating from the prediction of the randomized-strategy profile. Another concept of equilibrium relevant for sponsored search auctions is the one of *ex-post Nash equilibrium*. Before defining it, it is important to present the concept of *best response*, defined as in [46].

Definition 2.1.13 *Best response*

Considering a strategy profile s and a player i , a change from strategy s_i to s'_i is a best response if s'_i maximizes the player's utility $\max_{s'_i \in S_i} u_i(s'_i, s_{-i})$.

With the notion of best response, the concept of ex-post Nash equilibrium is then

defined as done in [46].

Definition 2.1.14 *Ex-post Nash equilibrium*

Given a Bayesian game $\Gamma^b = (N, X, T, u, p)$, a strategy profile $s = (s_i)_{i \in N}$ is an ex-post Nash equilibrium of Γ if, for all i , all $t_1, \dots, t_n$, and all x'_i ,

$$u_i(t_i, s_i(t_i), s_{-i}(t_{-i})) \geq u_i(t_i, x'_i, s_{-i}(t_{-i})),$$

i.e. $s_i(t_i)$ is a best response to any $s_i(t_{-i})$ for every possible value of t_{-i} .

Even if at first sight the concept of ex-post Nash equilibrium seems too good to be true, it will be shown that it is important for the most relevant type of auction used in sponsored search ads: the GSP auction.

2.1.2 Auctions as repeated Bayesian games

After formalizing some of the main concepts of Game Theory that are relevant for auctions, it is possible to focus on the theory of the auctions used in sponsored search advertisements. As explained in Section 1.1, these auctions have two main agents: the *advertisers*, who procure ad impressions on publishers' websites, aiming to reach potential consumers; and the *publishers*, who provide ad slots to the advertisers with the highest valuation for these slots. The *users*, for whom the ads are shown, are indirect agents. The value given for an ad slot depends on the users' occasional clicks on the ads as potential consumers for the advertisers.

In sponsored search ads, the auctioned items are ad slots for a set of words combinations on the publishers' platforms. For example, an item can be an ad slot in the Google search platform when a user searches "Electric car", "Electric vehicle", "Electric automobile", and other similar words. Different companies in the electric car market that are interested in showing their ads on Google will bid for it. In this context, a single item or multiple items can be auctioned. In single-item auctions, there is only one ad slot to be auctioned. Multi-item auctions consist in auctioning different items at the same time. In the sponsored search ads case, it means different slots placed differently in the publisher's platform. The value that the advertisers give for each of these items can vary, and the most common reason for that is that the *click-through rate* varies for each item. The definition of click-through rate from [49] is shown as follows.

Definition 2.1.15 *Click-through rate (CTR)*

The click-through rate is the number of times a click is made on the advertisement divided by the total impressions (the number of times an advertisement was served). It measures the proportion of visitors who initiated an advertisement that redirected them to another page where they might purchase an item or learn more about a product or service.

An ad shown at the top of the page in the publisher's platform tends to have a higher CTR than an ad shown below it, for example. In this subsection, the focus is to formalize the most relevant types of auctions in the context of sponsored search ads as repeated Bayesian games (see Definition 2.1.9). As explained in Section 1.1, this thesis focuses on single-item second-price auctions and GSP auctions, two of the most used types of auctions by the biggest publishers. Three main differences between these auctions are highlighted: whether they are single-item or multi-item auctions; if their mechanisms are DSIC or not; and which types of equilibria exist for them in the infinitely repeated case.

In Definition 2.1.1, it was shown that one of the aspects that define an auction is its clearing rules. This concept can be divided in two main concepts, *allocation rule* and *payment rule*, which are important to distinguish each auction type. Starting with the allocation rule, the following definition originates from [46].

Definition 2.1.16 *Allocation rule*

Consider a set N of agents, a set X of alternatives (actions), and for each agent $i \in N$ a set of potential preference relations $\mathcal{R}_i$ over the actions X . An allocation rule is a function

$$f^a : \times \mathcal{R}_i \rightarrow A,$$

mapping the preferences of the agents into the set of outcomes A .

A simple (and probably the most known) example of allocation rule for single-item auctions is that the highest bid wins. The allocation rule is not enough to define the clearing rules of an auction: the definition of a payment rule is shown as follows.

Definition 2.1.17 *Payment rule*

Again, considering a set N of agents, a set X of alternatives, and a set of potential preference relations $\mathcal{R}_i$ over the actions X for each agent $i \in N$, a payment rule is a function

$$f^p : \times \mathcal{R}_i \rightarrow \mathbb{R}^N,$$

mapping the preferences of the agents into their payments.

With these concepts presented, the two types of auctions mentioned before are formalized and analyzed, starting with single-item second-price auctions, or simply *second-price auctions*.

Second-price auctions

The first auction described is the second-price auction. The second-price auction is a single-item auction and, as mentioned in Section 1.1, it is used by Twitter (see [11]) and can be seen as a specific case of the generalized second-price auction. The definition from [45] of second-price auction is shown as follows.

Definition 2.1.18 *Second-price auction*

In a second-price (or *Vickrey*) auction, there are N bidders who simultaneously submit a real-valued positive bid. The bidder with the highest bid wins the good and pays the value b_2 of the second-highest bid, while the other bidders get and pay nothing (in case of a tie, the winner is determined by a flip of a coin).

To formalize its definition within the concept of repeated Bayesian game, it is important to remind that it can be described as a tuple (see Definition 2.1.9). In the case of second-price auctions, some of the elements of this tuple can be specified, namely:

- The set $\Theta = \{Beginning, winner(1), \dots, winner(i), \dots, winner(n), Finished\}$ of states of the world, where *Beginning* means that the game just started and there is still no winner, *winner(i)* means i is the current winner, and *Finished* means that the game has already finished;
- The set of actions X , which represents the set of all bids for all players;

- The set of all signals that the players can receive $S = (S_i)_{i \in N}$, where $S_i = \{Beginning, winner, not_winner, Finished\}$;
- The utility function for agent i : $u_i = \begin{cases} x_i - \max_{j \neq i} x_j & \text{if } x_i \geq \max_{j \neq i} x_j \\ 0 & \text{if } x_i < \max_{j \neq i} x_j \end{cases}$.

The other elements are straightforwardly extracted from the definition. The goal now is to understand whether a second-price auction is DSIC or not. For that, the concepts of *implementable* and *monotone* allocation rules are needed and are defined as done in [47].

Definition 2.1.19 *Implementable allocation rule*

An allocation rule f^a for a single-dimensional (single-item) environment is implementable if there is a payment rule f^p such that the sealed-bid auction (f^a, f^p) is DSIC.

Definition 2.1.20 *Monotone allocation rule*

An allocation rule f^a for a single-dimensional environment is monotone if for every bidder i and bids b_{-i} from other bidders, the allocation $f^a(z, b_{-i})$ to i is non-decreasing in i 's bid z .

To continue the study of whether second-price auctions are DSIC or not, Myerson's lemma is introduced as follows as done in [47].

Lemma 2.1.1 *Myerson's lemma*

Fixing a single-parameter environment:

- (i) An allocation rule f^a is implementable if and only if it is monotone;
- (ii) If f^a is monotone, then there is a unique payment rule such that the sealed-bid mechanism (f^a, f^p) is DSIC;
- (iii) The payment rule f^p is given by an explicit formula.

Finally, using the parts (i) and (ii) from Myerson's lemma and considering the (quasilinear) utility function previously presented, an important corollary from [47] is

shown as follows.

Corollary 2.1.1

A second-price auction with quasilinear utility functions and private values for its agents has an implementable allocation rule and a payment rule such that it is DSIC.

From the fact that bidding the true valuation of the item is a dominant strategy for all the bidders and their utilities cannot be improved by bidding another value, one can imagine that a bidder's true value is a Nash equilibrium. When considering infinitely repeated auctions, this is shown by the *Folk Theorem*. Before stating this theorem, two concepts have to be presented: the ones of *feasible* and *enforceable* utility profiles, whose definitions originate from [50].

Definition 2.1.21 *Feasible payoff profile*

A payoff profile $\bar{u} = (\bar{u}_i)_{i \in N}$ is feasible if it can be achieved, i.e., $\forall i \in N$, $\exists \alpha_x \in \mathbb{Q}^+$ such that $\sum_{x \in X} \alpha_x u_i(x) = \bar{u}_i$ and $\sum_{x \in X} \alpha_x = 1$.

Definition 2.1.22 *Enforceable payoff profile*

An enforceable payoff profile is a payoff that can be enforced by any mechanism to be achieved, which is, a feasible payoff profile in which the payoff of every player is larger than or equal to zero.

Knowing these concepts, the Folk theorem is stated as done is [45].

Theorem 2.1.1 *Folk theorem*

Every payoff profile that is both enforceable and feasible is a Nash equilibrium in an infinitely repeated game.

From the theorem, it is possible to see that there are possibly infinitely many Nash equilibria for second-price auctions, which explains the use of optimization methods for maximizing the utilities of the agents.

Generalized second-price auctions

The generalized second-price (GSP) auction can be considered as an extension of the (single-item) second-price auction to multi-item auctions. As detailed in Section 1.1, it is the most used auction type by the biggest publishers in sponsored search ads, such as Google (see [4] and [5]) and Alibaba (see [6]). The definition of GSP auctions from [46] is shown as follows.

Definition 2.1.23 *Generalized second-price (GSP) auction*

Considering N bidders, $K \leq N$ items sold, and the set of (decreasing) bids $\mathbf{B} = \{b_1, \dots, b_i, \dots, b_N\}$, in a GSP auction, the bidder i :

- receives item i at price b_{i+1} if $i \leq K$;
- receives and pays nothing if $i > K$.

To formalize the definition of GSP auctions as repeated Bayesian games, some elements of the tuple described in Definition 2.1.9 are specified:

- The set $\Theta = \{Beginning, winner_k(i), Finished\}$ of states of the world, where *Beginning* means that the game just started and there is still no winner, *winner_k(i)* means i is the current winner of item k , and *Finished* means that the game has already finished;
- The set of actions X , which represents the set of all bids for all players;
- The set of all signals that the players can receive $S = (S_i)_{i \in N}$, where $S_i = \{Beginning, winner_1, \dots, winner_k, \dots, winner_K, not_winner, Finished\}$;
- The private information of bidder i : $T_i = \{t_i^{(1)}, \dots, t_i^{(K)}\}$, which represents her true value for each item;
- The utility function for agent i : $u_i = \sum_{k=1}^K \max(0, t_i^{(k)} - \max_{j \in U_{k+1}} x_j)$, where U_{k+1} is defined recursively:
 - $B_1 = \mathbf{B}$;
 - $U_i = \{x \in B_i : x \geq b \quad \forall b \in B_i\}$, $B_{i+1} = B_i \setminus U_i$.

The other elements of the tuple are straightforwardly defined. Regarding the property of being DSIC or not, the following corollary proven in [44] and [51] states GSP auctions are not DSIC.

Corollary 2.1.2

Truth-telling is not a dominant strategy under GSP. Therefore, a GSP auction is not DSIC.

As bidding the true value is not a dominant strategy under the GSP mechanism, bidders tend to deviate from their true value. In theory, advertisers could gradually learn the values from other advertisers using automated robots, consequently having more private information about the auction. However, in practice, they may not be able to, as their bidding software must first be authorized by the publishers, which are unlikely to permit strategies which would substantially reduce their revenues. [51]

Although truth-telling is not an equilibrium of GSP, it has an important property. Reminding the concept of best response (Definition 2.1.13) and ex-post Nash equilibrium (Definition 2.1.14), the following corollary was proven in [51].

Corollary 2.1.3

A GSP auction in continuous strategies has a unique ex-post Nash equilibrium. In this equilibrium, the strategy of each advertiser is a best response to other advertisers' strategies regardless of their realized values.

In other words, even if a particular player learned the values of other players before the game, she would not want to change her strategy, implying that the equilibrium is robust. Moreover, it has been proven in [51] that, in this equilibrium, the expected revenue of the seller is at least as high as in the dominant-strategy equilibrium of the VCG auction. This explains why GSP auctions are the most common type of auction in the sponsored search ad market.

2.1.3 Auctions with reserve price

After discussing the main types of auctions that will be used in this thesis, it is important to highlight an additional variable that became common in sponsored search auctions: the reserve price, a limit on the price of the auctioned item, which is chosen

by the publisher. As mentioned in Section 1.1, the choice of the reserve price is the main mechanism through which a seller can influence the auction revenue. [14] As this thesis focuses on the point of view of the publishers, the definition of reserve price on the supply side from [28] is shown as follows.

Definition 2.1.24 *Reserve price*

On the supply side, reserve (or reservation) price is the minimum bid that the seller would accept from bidders in an auction.

A reserve price is easy to set and adjust, and has become a common tool for search engines to manage their revenues. [52] It can be shared in advance with the advertisers, then being an additional private information for all the bidders (see [52]), or be kept as a secret until the end of the auction (see [53]), simply excluding any bid below the reserve price. As shown in [53], a seller can increase her revenues by conducting an auction in which her reserve price is kept as secret compared to conducting an auction while announcing the reserve price.

When considering (single-item or generalized) second-price auctions, the role of the reserve price is more than only filtering bids above a certain value. Considering the case of a single-item second-price auction, the advertiser whose bid was the highest among all the bids will pay the maximum between the reserve price and the second-highest bid, if her bid was higher than the reserve price. This way, choosing the reserve price wisely is important for increasing the publishers' revenues. If the reserve price is set too high, too many bids will be excluded and it will adversely affect the revenue stream because unoccupied ad positions are wasted. On the other hand, if the reserve price is too low, search engines may lose the opportunity of increasing their revenues when advertisers are willing to pay more. [52]

The two types of auctions studied in this thesis (second-price and GSP auctions) will be detailed when in the presence of reserve prices.

Second-price auction with reserve price

The reserve price not only sets which bids are admitted in second-price auctions but also affects the payment made by the winner. So, it changes both the allocation and the payment rules. Based on Definition 2.1.18, a second-price auction with reserve price is defined as follows.

Definition 2.1.25 *Second-price auction with reserve price*

In a second-price (or *Vickrey*) auction with reserve price, there are N bidders who simultaneously submit a real-valued positive bid, while the seller sets a reserve price r . The bidder with the highest bid wins the good if her bid is above the reserve price and pays the maximum between r and b_2 , the second-highest bid. The other bidders get and pay nothing (in case of a tie, the winner is determined by a flip of a coin). If there is no bid above the reserve price, the bidder with the highest bid also gets and pays nothing.

As done before, a second-price auction with reserve price can be formalized as a repeated Bayesian game when specifying the following elements:

- The set $\Theta = \{Beginning, winner(1, r), winner(1, b_2), \dots, winner(i, r), winner(i, b_2), \dots, winner(N, r), winner(N, b_2), Finished\}$ of states of the world, where *Beginning* means that the game just started and there is still no winner, *winner(i, r)* means i is the current winner paying the reserve price, *winner(i, b_2)* means i is the current winner paying the second-highest bid, and *Finished* means that the game has already finished;
- The set of actions X , which represents the set of all bids for all players;
- The set of all signals that the players can receive $S = (S_i)_{i \in N}$, where $S_i = \{Beginning, winner^r, winner^{b_2}, not_winner, Finished\}$;
- The utility function for agent i :

$$u_i = \begin{cases} x_i - \max(\max_{j \neq i} x_j, r) & \text{if } x_i \geq \max_{j \neq i} x_j, \quad x_i \geq r \\ 0 & \text{otherwise} \end{cases}.$$

GSP auction with reserve price

As in a single-item second-price auction, the reserve price changes the allocation and the payment rules of GSP auctions. There are two possible ways to add reserve prices to a GSP auction. First, it is possible to have a specific reserve price for each item. The second option is to have a single reserve price for all the items. Starting with the first possibility and based on Definition 2.1.23, a GSP auction with specific reserve prices is defined as follows.

Definition 2.1.26 *GSP auction with specific reserve prices*

Considering N bidders, $K \leq N$ items sold, the set of (decreasing) bids $\mathbf{B} = \{b_1, \dots, b_i, \dots, b_N\}$, and the set of reserve prices $\mathbf{R} = \{r_1, \dots, r_k, \dots, r_K\}$ in a GSP auction, the bidder i :

- receives item i at price $\max(r_i, b_{i+1})$ if $i \leq K$ and $b_i \geq r_i$;
- receives and pays nothing otherwise.

Also based on Definition 2.1.23, a GSP auction with single reserve price is defined as follows.

Definition 2.1.27 *GSP auction with single reserve price*

Considering N bidders, $K \leq N$ items sold, the set of (decreasing) bids $\mathbf{B} = \{b_1, \dots, b_i, \dots, b_N\}$, and a reserve price r in a GSP auction, the bidder i :

- receives item i at price $\max(r, b_{i+1})$ if $i \leq K$ and $b_i \geq r$;
- receives and pays nothing otherwise.

As mentioned in [31], the case of GSP auctions with single reserve price is the one that is used the most in reality. So, this thesis will focus on this case when analyzing GSP auctions with reserve price, and it can be detailed as a repeated Bayesian game by defining the following elements:

- The set $\Theta = \{Beginning, winner_k(1, r), winner_k(i, b_{i+1}), Finished\}$ of states of the world, where *Beginning* means that the game just started and there is still no winner, $winner_k(i, r)$ means i is the current winner of item k paying the reserve price r , $winner_k(i, b_{i+1})$ means i is the current winner of item k paying the next highest bid b_{i+1} , and *Finished* means that the game has already finished;
- The set of actions X , which represents the set of all bids for all players;
- The set of all signals that the players can receive $S = (S_i)_{i \in N}$, where $S_i = \{Beginning, winner_1^r, winner_1^{b_{i+1}}, \dots, winner_k^r, winner_k^{b_{i+1}}, \dots, winner_K^r, winner_K^{b_{i+1}}, not_winner, Finished\}$;
- The private information of bidder i : $T_i = \{t_i^{(1)}, \dots, t_i^{(K)}\}$, which represents her true value for each item;

- The utility function for agent i : $u_i = \sum_{k=1}^K \max\{0, t_i^{(k)} - \max(r, \max_{j \in U_{k+1}} x_j)\}$, where U_{k+1} is defined recursively:
 - $B_1 = \mathbf{B}$;
 - $U_i = \{x \in B_i : x \geq b \quad \forall b \in B_i\}$, $B_{i+1} = B_i \setminus U_i$.

Although it was possible to see the existence of different equilibria and dominant strategies for second-price and GSP auctions from a game-theoretic point of view, reserve prices and other practical constraints that occur in sponsored search auctions (such as the advertisers' budgets) are usually not covered by theoretical dominant strategies. It confirms the importance of using optimization models when adding complexity to these auctions.

2.2 Stochastic programming

As we have just seen, optimization models are important when adding more complex practical elements to sponsored search ads. In this thesis, a multi-stage stochastic MILP model is derived, and this subsection presents the most relevant concepts of Stochastic Programming for understanding this model.

2.2.1 Two-stage stochastic programming

Before moving to multi-stage models, it is important to understand two-stage cases before. This subsection addresses the modelling of uncertainty in linear programs as two-stage stochastic programs. The evaluation of the performance of these programs' solutions is also described. The main references for the definitions, propositions, and proofs shown in the course of this subsection are the lecture notes from Papavasiliou in [54].

Two-stage stochastic linear programs

The analysis of two-stage models in this thesis will be restricted to the linear case. *Stochastic linear programs* (SLP) are linear programs with uncertain data. The main aspect of stochastic linear programs that make their optimization challenging to solve is the *recourse*: the ability to react after uncertainty has been revealed to the decision maker. The sequence of events in two-stage stochastic linear programs is illustrated in Figure 2.1.




Figure 2.1: Sequence of events in two-stage stochastic linear programs. [54]

Typically, a two-stage stochastic program involves two sets of decisions:

- *First-stage decisions*: set of decisions that need to be fixed prior to the revelation of uncertainty;
- *Second-stage decisions*: set of decisions determined based on the realization of uncertainty once it is revealed.

Given these concepts and considering a probability space $(\Omega, \mathcal{B}(\Omega), \mathbb{P})$, two-stage stochastic programs can be formulated as follows:

Model 2.2.1 *Two-stage stochastic linear program*

$$\begin{aligned}
 \min \quad & c^T x + \mathbb{E}[\min q(\omega)^T y(\omega)] \\
 \text{s.t.} \quad & Ax = b, \\
 & T(\omega)x + W(\omega)y(\omega) = h(\omega), \\
 & x \geq 0, \quad y(\omega) \geq 0, \quad \omega \in \Omega.
 \end{aligned}$$

In the model, $x \in \mathbb{R}^{n_1}$ represents the first-stage decisions, while the second-stage decisions for a given realization ω are denoted as $y(\omega) \in \mathbb{R}^{n_2}$. The parameters of the first stage are deterministic and denoted as $c \in \mathbb{R}^{n_1}$, $b \in \mathbb{R}^{m_1}$, and $A \in \mathbb{R}^{m_1 \times n_1}$, while the parameters of the second stage are $q(\omega) \in \mathbb{R}^{n_2}$, $h(\omega) \in \mathbb{R}^{m_2}$, $T(\omega) \in \mathbb{R}^{m_2 \times n_1}$, and $W(\omega) \in \mathbb{R}^{m_2 \times n_2}$.

An important specific case is the one of stochastic linear programs with *fixed recourse*. A stochastic linear program is said to obey fixed recourse when W (see Model 2.2.1) does not depend on ω .

Performance of two-stage stochastic programming solutions

When dealing with stochastic programs, it is useful to know what are the real benefits of adopting a sophisticated model of uncertainty (such as Stochastic Programming). One can straightforwardly think about two alternatives that could be checked:

- What performance could be attained if the future were known in advance;
- What performance could be attained by a simpler policy.

One can denote all random parameters of a two-stage stochastic linear program as a vector

$$\xi^T = (q^T, T_1^T, \dots, T_{n_1}^T, W_1^T, \dots, W_{n_2}^T, h^T),$$

where T_i and W_i denote the i -th column of T and W , respectively. One can also define the set

$$K_1 = \{x | Ax = b, x \geq 0\},$$

representing the first-stage decisions that respect first-stage constraints, which in turn is useful to define the function $z(x, \xi)$:

$$z(x, \xi) = c^T x + \mathcal{Q}(x, \xi) + \delta(x | K_1),$$

where

$$\mathcal{Q}(x, \xi) = \min_y \{q(\omega)^T y | W(\omega)y = h(\omega) - T(\omega)x\}$$

and $\delta(x | K_1)$ is the *indicator function*, defined as $\delta(x | K_1) = 0$ if $x \in K_1$ and $\delta(x | K_1) = +\infty$ otherwise. The function $z(x, \xi)$ is the cost incurred if one were to optimize only for scenario ω , where this cost is parametrized by x .

With these concepts described, it is possible to define an important concept for evaluating the performance of two-stage stochastic programming solutions, the *wait-and-see value*, whose definition originates from [54].

Definition 2.2.1 *Wait-and-see value (of a two-stage stochastic program)*

The wait-and-see value of a two-stage stochastic program is the expected value of reacting with perfect foresight $x^*(\xi)$ to any outcome ω :

$$\begin{aligned} WS &= \mathbb{E}[\min_x z(x, \xi)] \\ &= \mathbb{E}[z(x^*(\xi), \xi)]. \end{aligned}$$

Moreover, the definition of the *here-and-now value* of a two-stage stochastic program from [54] is shown as follows.

Definition 2.2.2 *Here-and-now value (of a two-stage stochastic program)*

The here-and-now value (or stochastic programming solution) is the expected value of the two-stage stochastic program:

$$SP = \min_x \mathbb{E}[z(x, \xi)].$$

One can notice that the only difference between WS and SP in terms of formulation is that the minimization and expectation operators have been swapped. However, in terms of computation, SP requires the solution of a stochastic program, often requiring non-trivial algorithms, while WS requires the solution of $|\Omega|$ linear programs, where $|\Omega|$ is the number of second-stage outcomes, and therefore is generally easier to compute than SP .

After defining these two values, it is possible to present the concept of *expected value of perfect information*.

Definition 2.2.3 *Expected value of perfect information*

The expected value of perfect information is the difference between the wait-and-see value and the stochastic programming solution:

$$EVPI = SP - WS.$$

In other words, the *EVPI* is the additional benefit that could be obtained if a perfect forecast of the future were possible, comparatively to the here-and-now value.

After comparing the stochastic programming solution to the case where the future is known in advance, as we have just seen, another alternative is comparing it to a simpler policy. The simpler policy usually considered as benchmark is the *expected value problem*, described as follows as done in [54].

Definition 2.2.4 *Expected value problem (of a two-stage stochastic program)*

Noting $\bar{\xi} = \mathbb{E}[\xi]$, the expected (or mean) value of the random parameters, the expected value problem of a two-stage stochastic program is written as

$$EV = \min_x z(x, \bar{\xi}).$$

Finding the optimal solution of the expected value problem consists in solving a linear program only with the average scenario. However, even if this problem is smaller and easier to compute when compared to the stochastic program, it is important to observe the performance of this optimal first-stage decision, denoted as $x^*(\bar{\xi})$, when applied to the real problem. For that, the *expected value of using the EV solution* is presented as follows, using the definition from [54].

Definition 2.2.5 *Expected value of using the EV solution (for a two-stage stochastic program)*

For a two-stage stochastic program, the expected value of using the EV solution measures the performance of $x^*(\bar{\xi})$ by reacting optimally in the second stage given the first-stage decision $x^*(\bar{\xi})$:

$$EEV = \mathbb{E}[z(x^*(\bar{\xi}), \xi)].$$

It is possible to see that *EEV* is less challenging to compute than *WS*. To compute *EEV*, it is necessary to solve, for every $\omega \in \Omega$, a linear program only over the second-stage variables, while the first-stage variables also need to be taken into account when computing *WS*. The concept of *EEV* allows the definition of the *value of the*

stochastic solution.

Definition 2.2.6 *Value of the stochastic solution*

The value of the stochastic solution is defined as the benefit of the stochastic programming solution relative to the simpler decision rule $x^*(\bar{\xi})$:

$$VSS = EEV - SP.$$

Two general and intuitive inequalities can be extracted when comparing SP , EEV , and WS . The first one relates SP and WS and is shown in the following proposition, originated from [54]:

Proposition 2.2.1

Given a minimization stochastic linear program, it holds that:

$$WS \leq SP.$$

Proof For every ω , $z(x^*(\xi), \xi) \leq z(x^*, \xi)$, where x^* is the optimal first-stage stochastic programming solution. One can take expectations on both sides to obtain $WS \leq SP$.

□

The second inequality compares EEV and SP . It also originates from [54] and is shown as follows.

Proposition 2.2.2

Given a minimization stochastic linear program, it holds that:

$$SP \leq EEV.$$

Proof Reminding that x^* is the optimal solution of

$$\min_x \mathbb{E}[z(x, \xi)],$$

one can easily see that the vector $x^*(\bar{\xi})$ is a solution for the same minimization problem. Therefore, $\min_x \mathbb{E}[z(x, \xi)] = SP \leq EEV = \min_x \mathbb{E}[z(x(\bar{\xi}), \xi)]$.

□

These inequalities are summarized in Figure 2.2.




Figure 2.2: Relations between WS , SP , and EEV .

2.2.2 Multi-stage stochastic programming

After introducing some of the main concepts of two-stage stochastic programming, it is possible to move to multi-stage models. The main references to the concepts presented in this subsection are lecture notes from Papavasiliou in [54] and the work of Escudero et al. in [55].

Scenario trees and lattices

Before presenting multi-stage stochastic programs, it is important to know how uncertainty is modelled in these models. Considering that the system can be modelled as a Markov decision model, one can define a set S_t at each stage $t = 1, \dots, H$ of the problem and the sample space as $\Omega = S_1 \times \dots \times S_H$. The probability space $(\Omega, \mathcal{A}_H, \mathbb{P})$ can then be defined, where $\mathcal{A}_H = 2^\Omega$ is the total set of events at stage H and $\mathbb{P}$ is a probability measure on these events.

Multi-stage stochastic programs can be also defined in more concise representations than their probability spaces. One of these representations is the *scenario tree*. Before defining a scenario tree, it is important to present the concepts of *tree* and *rooted tree*, described as follows using the definitions from [54].

Definition 2.2.7 *Tree*

A tree (N, E) is a connected acyclic graph, where N denotes its nodes and E denotes its edges.

Definition 2.2.8 *Rooted tree*

A rooted tree is a tree in which there exists a unique node that is singled out. This node is designated as the *root*.

Once rooted trees are defined, it is possible to define a scenario tree, whose definition originates from [54].

Definition 2.2.9 *Scenario tree*

A scenario tree is a rooted tree which is used in multi-stage programs for describing uncertainty in a model. The nodes of a scenario tree represent histories of realizations of the random input $\xi_{[t]} = (\xi_1, \dots, \xi_t)$, with each node characterized by a unique time index. The edges represent transition probabilities from a history $\xi_{[t]} \in \Xi_{[t]}$ in stage t to a history $\xi_{[t+1]} \in \Xi_{[t+1]}$, where $\Xi_{[t]} = \Xi_1 \times \dots \times \Xi_t$. The root of a scenario tree corresponds to the first stage $t = 1$.

Denoting $\omega_{[t]} = (\omega_1, \dots, \omega_t)$ as an index in the set $\Xi_{[t]} = \Xi_1 \times \dots \times \Xi_t$, where ω_t is an index in the set Ξ_t , the definitions of *ancestor* and *descendants* of a node can now be presented and are shown as follows, based on the definitions presented in [54].

Definition 2.2.10 *Ancestor (of a node)*

The ancestor of a node $\omega_{[t]}$, denoted as $a(\omega_{[t]})$, is the unique adjacent node which precedes $\omega_{[t]}$, i.e.

$$a(\omega_{[t]}) = \{\omega_{[t-1]} : (\omega_{[t-1]}, \omega_{[t]}) \in E\}.$$

Definition 2.2.11 *Descendants (of a node)*

The descendants (or *children*) of a node $\omega_{[t]}$, denoted as $c(\omega_{[t]})$, correspond to the set of nodes that are adjacent to $\omega_{[t]}$ and are indexed by a subsequent time stage, i.e.

$$c(\omega_{[t]}) = \{\omega_{[t+1]} : (\omega_{[t]}, \omega_{[t+1]}) \in E\}.$$

An alternative way to represent a Markov process is a *lattice*, described by the following definition from [54].

Definition 2.2.12 *Lattice*

A lattice consists of

- (i) a set of nodes, each of which is characterized by a time stage, and each of which corresponds to a realization of ξ_t ;
- (ii) a set of edges connecting all nodes in stage t to all nodes in stage $t + 1$, corresponding to transition probabilities.

While the nodes of a scenario tree represent a history $\xi_{[t]}$, the nodes of a lattice represent the realization ξ_t of a process. Figure 2.3 illustrates these two representations.




Figure 2.3: Graphical representations of a scenario tree and a lattice. [54] One can see that a lattice is a compact version for displaying the realizations of uncertainty while a scenario tree contains an exponential number of nodes.

Multi-stage stochastic linear programs

After presenting how uncertainty is modelled in multi-stage stochastic programming, it is possible to describe a multi-stage stochastic linear program (MS-SLP). Considering a finite set of outcomes Ω and a probability space $(\Omega, 2^\Omega, \mathbb{P})$, the following *extensive formulation* of an MS-SLP is presented.

Model 2.2.2 *Multi-stage stochastic linear program*

$$\begin{aligned} \min \quad & c^T x_1 + \mathbb{E}[\min c_2(\omega)^T x_2(\omega) + \cdots + c_H(\omega)^T x_H(\omega)] \\ \text{s.t.} \quad & W_1 x_1 = h_1, \\ & T_1(\omega)x_1 + W_2(\omega)x_2(\omega) = h_2(\omega), & \omega \in \Omega, \\ & \vdots \\ & T_{t-1}(\omega)x_{t-1} + W_t(\omega)x_t(\omega) = h_t(\omega), & \omega \in \Omega, \\ & \vdots \\ & T_{H-1}(\omega)x_{H-1} + W_H(\omega)x_H(\omega) = h_H(\omega), & \omega \in \Omega, \\ & x \geq 0, \quad x_t(\omega) \geq 0, & t = 2, \dots, H, \omega \in \Omega. \end{aligned}$$

In this model, $\xi_t^T(\omega) = [c_t^T(\omega), h_t^T(\omega), \text{vec}(T_{t-1}^T(\omega)), \text{vec}(W_t^T(\omega))]$ is the random input and x_t is the decision vector at stage t .

Multi-stage stochastic MILPs

Once MS-SLPs are defined, it is possible to move to multi-stage stochastic mixed-integer linear programs (MS-SMILPs). Considering that each decision vector x_t can be denoted as

$$x_t = [x_t^{(1)}, \dots, x_t^{(n_1)}],$$

and that the set of stages is denoted as T , it is possible to write the following extensive formulation for MS-SMILPs.

Model 2.2.3 *Multi-stage stochastic MILP*

$$\begin{aligned}
\min \quad & c^T x_1 + \mathbb{E}[\min c_2(\omega)^T x_2(\omega) + \dots + c_H(\omega)^T x_H(\omega)] \\
\text{s.t.} \quad & W_1 x_1 = h_1, \\
& T_1(\omega)x_1 + W_2(\omega)x_2(\omega) = h_2(\omega), & \omega \in \Omega, \\
& \vdots \\
& T_{t-1}(\omega)x_{t-1} + W_t(\omega)x_t(\omega) = h_t(\omega), & \omega \in \Omega, \\
& \vdots \\
& T_{H-1}(\omega)x_{H-1} + W_H(\omega)x_H(\omega) = h_H(\omega), & \omega \in \Omega, \\
& x_1 \geq 0, \quad x_t(\omega) \geq 0, & t \in T, \omega \in \Omega, \\
& x_1^{(j_1)} \in \{0, 1\}, & \text{for some } j_1 \in \{1, \dots, n_1\}, \\
& x_t^{(j_t)}(\omega) \in \{0, 1\}, & \text{for some } j_t \in \{1, \dots, n_t\}, t \in T.
\end{aligned}$$

Performance of multi-stage stochastic programming solutions

Extending the parameters *VSS* and *EVPI* to multiple periods is not straightforward. In particular, it is not clear which variables must be fixed in the *WS* model and at which stages the decision variables should be fixed in the *EEV* model. However, there are studies to understand the best way to extend these concepts to the case of multi-stage stochastic programs. The approach used in this thesis has as main reference the work from Escudero et al. in [55].

Starting with the wait-and-see value, a trivial solution would be to fix only the first-stage decisions. However, as shown in [55], this procedure can lead to a paradox in some cases. *WS* models should, then, involve the models of all possible scenarios and consider that they are completely independent. This way, the wait-and-see value considered in this thesis is the solution of the *perfect foresight model*. Denoting S as the set of all possible scenarios and $|S|$ as the total number of scenarios, the perfect foresight model for an MS-SMILP is shown as follows.

Model 2.2.4 *Perfect foresight model of an MS-SMILP*

$$\begin{aligned}
\min \quad & \frac{1}{|S|} \left[c^T x_{1,s} + \cdots + c_{H,s}^T x_{H,s} \right] \\
\text{s.t.} \quad & W_1 x_{1,s} = h_1, & s \in S, \\
& T_{1,s} x_{1,s} + W_{2,s} x_{2,s} = h_{2,s}, & s \in S, \\
& \vdots \\
& T_{t-1,s} x_{t-1,s} + W_{t,s} x_{t,s} = h_{t,s}, & s \in S, \\
& \vdots \\
& T_{H-1,s} x_{H-1,s} + W_{H,s} x_{H,s} = h_{H,s}, & s \in S, \\
& x_{t,s} \geq 0, & t \in T, s \in S, \\
& x_{t,s}^{(j_t)}(\omega) \in \{0, 1\} & \text{for some } j_t \in \{1, \dots, n_t\}, t \in T, s \in S.
\end{aligned}$$

Considering WS as the solution of the perfect foresight model, the inequality $WS \leq SP$ is still valid for the same reasoning as the one presented in Proposition 2.2.1. The $EVPI$ for multiple periods is defined as before (see Definition 2.2.3). Moving to EV , the following definition is similar to the one for two-stage stochastic programs and is based on [55].

Definition 2.2.13 *Expected value problem (of a multi-stage stochastic program)*

Denoting $\bar{\xi}_{[H]} = (\bar{\xi}_1, \dots, \bar{\xi}_H)$ as the expected (or mean) value of the random parameters for all the stages, the expected value problem of a multi-stage stochastic program is written as

$$EV = \min_x z(x_1, \dots, x_H, \bar{\xi}_{[H]}).$$

The solution of the expected value problem is denoted as $x^* = (x_1^*, \dots, x_H^*)$. In [55], EEV is defined for each stage t and therefore is slightly different from the two-stage

case. This definition is shown as follows.

Definition 2.2.14 *Expected value of using the EV solution (for a multi-stage stochastic program)*

The expected value of using the *EV* solution at t for $t = 2, \dots, T$ is the optimal value of the stochastic model (2.2.3), where the decision variables until stage $t - 1$, $(x_1, \dots, x_{t-1})$, are fixed at the optimal values obtained as solution of the average scenario model.

$$EEV_t = \begin{cases} \text{Model 2.2.3} \\ s.t. & \begin{aligned} x_1 &= x_1^*, \\ x_2(\xi) &= x_2^*, \quad \forall \xi \in \Xi, \\ &\vdots \\ x_{t-1}(\xi) &= x_{t-1}^*, \quad \forall \xi \in \Xi. \end{aligned} \end{cases}$$

In this thesis, the focus will be on EEV_H , the expected value of using the *EV* solution at H , the last stage. For this case, the inequality shown in Proposition 2.2.2 is still valid (for the same reasoning) and the value of the stochastic solution at H , denoted as VSS_H , comes from Definition 2.2.6:

$$VSS_H = EEV_H - SP.$$

3 | A new multi-stage stochastic MILP approach

This chapter is devoted to presenting a new multi-stage stochastic MILP approach to optimize the reserve price in (single-item and generalized) second-price auctions.

First, the rules and details of the auctions considered in this thesis are explained. The main assumptions are highlighted and the stages are illustrated by scenario trees. The simplified versions of the auctions resulting from these assumptions are formalized.

Next, the focus is set to the formulation of the models. Starting with the single-item case, the objective function and the constraints are explained in details. It is also shown how the model is modified to become a MILP. This model is then extended to rank-by-bid and rank-by-revenue GSP auctions.

Finally, the single-item case is further analyzed. Considering four other approaches to choose the reserve price in second-price auctions (including the optimal reserve price if the budget constraints were ignored), it is shown that the optimal value from the proposed model always leads to greater or equal revenues for the publishers, considering the presented assumptions.

3.1 Auction rules and details

This thesis focuses on optimizing the reserve price in second-price and GSP sponsored search auctions in order to maximize the revenue of a publisher. To do that, these auctions are modelled as MS-SMILPs. It means that they still follow the rules seen in their game-theoretic models (see Subsections 2.1.2 and 2.1.3), but they are not modelled

as infinitely repeated Bayesian games. It is then assumed that these auctions can be described as a Markov decision model with stages $t = 1, \dots, H$. Modelling the auctions as multi-stage stochastic programs allows the addition of new constraints that may affect the auctions. In this thesis, the advertisers' budget constraints are added to the model.

The case considered here is the one in which the reserve price is kept as a secret. As explained in Subsection 2.1.3, by not announcing the reserve price to the advertisers, a seller can increase her revenues. The uncertainty comes from the advertisers' bids at stages $t = 2, \dots, H$, in which the auctions happen. It is considered that each advertiser has a bid distribution for each of these stages, a commonly seen assumption (see [8]). It is also assumed that these distributions can be learned by the publisher (see [31]) and that they can be approximated by discrete distributions. Moreover, the advertisers' budget constraints are considered to be known by the publisher, which indeed happens in real sponsored search auctions (see [56] and [57]).

Second-price auction

Starting with the single-item second-price auction, Figure 3.1 illustrates its scenario tree. At stage $t = 1$, the publisher decides the reserve price of the auction that will run at $t = 2$. At $t = 2, \dots, H - 1$, the winner of the auction (if any) is chosen according to the rules of the auction and the reserve price of the following auction is also decided. Finally, at $t = H$, the only decision is the winner of the last auction.




Figure 3.1: Scenario tree of a (single-item) second-price auction. Note that the reserve prices are chosen at $t = 1, \dots, H - 1$ while the winners are decided at $t = 2, \dots, H$.

There are some simplifications done to model the single-item second-price auction as an MS-SMILP. An advertiser may bid more than her budget left. However, if an advertiser has the highest bid but does not have enough budget left to pay it, there is no winner. Moreover, it is considered that the number of advertisers during the stages is constant.

The version of the second-price auction considered in this thesis is summarized as follows.

Definition 3.1.1 *Simplified version of the second-price auction*

The second-price auction considered in this thesis can be described as a Markov decision process happening at stages $t = 1, \dots, H$. This auction is detailed as follows :

- (i) The auctions run at stages $t = 2, \dots, H$;
- (ii) There is a single item being auctioned;
- (iii) The advertisers $i \in I$ participate in the auction during all the stages, and no new advertiser join them at any of these stages;
- (iv) Each advertiser i has a (fixed) maximum budget $\tilde{B}_i$ to be spent during these stages and this information is given to the publisher;
- (v) At $t = 1, \dots, H - 1$, the publisher chooses the reserve price r_t for the auction that happens at stage $t + 1$;
- (vi) At $t = 2, \dots, H$, each advertiser i makes a random bid $b_{i,t}^{\omega_{[t]}}$ drawn from a known discrete distribution;
- (vii) For advertiser i^* to win the auction at $t = 2, \dots, H$, her bid at t needs to be the highest among all the bids and above the reserve price r_t , and she must have enough budget left to pay for it. If there is no advertiser satisfying these conditions, the auction at t has no winner;
- (viii) If i^* wins an auction, she pays the maximum between the second-highest bid and the reserve price.

GSP auction

The stages for the GSP auction are similar to the ones presented for the second-price auction. At stage $t = 1$, a single reserve price is decided. The use of a single reserve price is more common than the use of specific reserve prices for each item in GSP sponsored search auctions (see [31]). At $t = 2, \dots, H - 1$, the winners of items $k = 1, \dots, K$ (if any) are chosen according to the rules of the auction and the reserve price of the following

auction is also decided. At the last stage, $t = H$, only the winners of the last auction are chosen. A scenario tree representing these stages is shown in Figure 3.2.




Figure 3.2: Scenario tree of a GSP auction. Note that the reserve prices are chosen at $t = 1, \dots, H - 1$ while the $|K|$ winners are decided at $t = 2, \dots, H$.

It is important to highlight that some publishers use different clearing rules when auctioning their ad slots as GSP auctions. The payment and allocation rules already described in Subsection 2.1.2 represent the so-called *rank-by-bid* GSP auction. However, there is a version of the GSP auctions that is used by some publishers called *rank-by-revenue* GSP auction, defined as follows as done in [58].

Definition 3.1.2 *Rank-by-revenue GSP auction with reserve price*

Considering that the click-through rate of an ad can be decomposed into a position-discount term θ_k and an advertiser-specific term q_i , i.e. $CTR_{k,i} = \theta_k q_i$, and that bidders are numbered in the descending order of $f_i = q_i b_i$, where b_i is the bid of advertiser i , the advertisers in a rank-by-revenue GSP:

- Are also ranked by descending order of $f_i = q_i b_i$;
- If $i \leq K$, pay $\max\{\frac{f_{i+1}}{q_i}, r\} = \max\{\frac{q_{i+1} b_{i+1}}{q_i}, r\}$.

Focusing on the clearing rules that had already been presented before (rank-by-bid GSP auctions), the assumptions done to second-price auctions are straightforwardly

extended as follows.

Definition 3.1.3 *Simplified version of the rank-by-bid GSP auction*

The GSP auction considered in this thesis can be described as a Markov decision process happening at stages $t = 1, \dots, H$, where :

- (i) The auctions run at stages $t = 2, \dots, H$;
- (ii) Items $k = 1, \dots, K$ are being auctioned;
- (iii) The advertisers $i \in I$, where $I \geq K$, participate in the auction during all the stages, and no new advertiser join them at any of these stages;
- (iv) Each advertiser i has a (fixed) maximum budget $\tilde{B}_i$ to be spent during these stages and this information is given to the publisher;
- (v) At $t = 1, \dots, H - 1$, the publisher chooses the reserve price r_t for the auction that happens at stage $t + 1$;
- (vi) At $t = 2, \dots, H$, each advertiser i makes a random bid $b_{i,t}^{\omega_{[t]}}$ drawn from a known discrete distribution;
- (vii) For advertiser i^* to win item k in the auction happening at $t = 2, \dots, H$, her bid at t needs to be the k -th highest among all the bids and above the reserve price r_t , and she must have enough budget left to pay for it. If there is no advertiser satisfying these conditions, item k has no winner at t ;
- (viii) If i^* wins item k , she pays the maximum between the following highest bid and the reserve price.

The same assumptions and simplifications done to the rank-by-bid GSP auction and formalized in Definition 3.1.3 can be also done to the rank-by-revenue GSP auction. As this extension is straightforward and for the sake of avoiding repetition, the simplified version of the rank-by-revenue GSP auction will not be formalized.

3.2 Model formulation

Once the main assumptions and details of the auctions are formalized, it is possible to proceed to the formulation of their models.

3.2.1 Second-price auction

Starting with the single-item second-price auction, the objective of the optimization model maximizes the expected total revenue of the publisher among all the

auctions:

$$\max \sum_{t=2}^T \mathbb{E}_{\omega_{[t]}} \left[\sum_{i \in I} v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + \sum_{i \in I} r_{t-1}^{a(\omega_{[t]})} y_{i,t}^{\omega_{[t]}} \right],$$

where $x_{i,t}^{\omega_{[t]}}$ is a binary variable which is equal to 1 if advertiser i wins auction t and pays the second-highest bid (otherwise it is equal to zero), $y_{i,t}^{\omega_{[t]}}$ is a binary variable which is equal to 1 if advertiser i wins auction t and pays the reserve price (otherwise it is equal to zero), $r_{t-1}^{a(\omega_{[t]})}$ is the nonnegative reserve price chosen at time $t-1$ (for the auctions that occur at t), and $v_{2,t}^{\omega_{[t]}}$ is the uncertain second-highest bid at time t .

As the second-price auction has only one item, there can only be one winner per auction, which is represented by the following constraint:

$$\sum_i x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}} \leq 1, \quad t \geq 2, \omega_{[t]} \in \Omega_{[t]}.$$

To ensure that the auctions follow the rules described in Definition 3.1.1, the following constraints allow only the advertiser with the highest bid to win and just let her win if her bid is higher or equal to the reserve price:

$$\begin{aligned} b_{i,t}^{\omega_{[t]}} &\geq (x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}}) v_{1,t}^{\omega_{[t]}}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ b_{i,t}^{\omega_{[t]}} &\geq (x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}}) r_{t-1}^{a(\omega_{[t]})}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \end{aligned}$$

where $b_{i,t}^{\omega_{[t]}}$ is the uncertain bid of advertiser i at time t and $v_{1,t}^{\omega_{[t]}}$ is the uncertain highest bid at time t . To make sure that the winner will pay the maximum between the reserve price and the second-highest bid, the following constraints are added:

$$\begin{aligned} x_{i,t}^{\omega_{[t]}} r_{t-1}^{a(\omega_{[t]})} &\leq x_{i,t}^{\omega_{[t]}} v_{2,t}^{\omega_{[t]}}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ y_{i,t}^{\omega_{[t]}} v_{2,t}^{\omega_{[t]}} &\leq y_{i,t}^{\omega_{[t]}} r_{t-1}^{a(\omega_{[t]})}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}. \end{aligned}$$

Finally, it is important to ensure that the advertisers do not spend more than their budget, which is done by the following sets of constraints:

$$\begin{aligned} B_{i,1}^{\omega_{[1]}} &= \tilde{B}_i, & i \in I, \omega_{[1]} \in \Omega_{[1]}, \\ B_{i,t}^{\omega_{[t]}} &= B_{i,t-1}^{a(\omega_{[t]})} - \left(v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + r_{t-1}^{a(\omega_{[t]})} y_{i,t}^{\omega_{[t]}} \right), & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ B_{i,t}^{\omega_{[t]}} &\geq 0, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \end{aligned}$$

where $\tilde{B}_i$ is the total budget of advertiser i and $B_{i,t}^{\omega_{[t]}}$ is the budget left for advertiser i at the end of period t . With that being explained, the model is shown as follows.

Model 3.2.1 *Multi-stage stochastic program for second-price auction*

$$\max_{x,y,r,B} \sum_{t=2}^T \mathbb{E}_{\omega_{[t]}} \left[\sum_{i \in I} v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + \sum_{i \in I} r_{t-1}^{a(\omega_{[t]})} y_{i,t}^{\omega_{[t]}} \right] \quad (3.1)$$

$$\text{s.t. } \sum_i x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}} \leq 1, \quad t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.2)$$

$$x_{i,t}^{\omega_{[t]}} r_{t-1}^{a(\omega_{[t]})} \leq x_{i,t}^{\omega_{[t]}} v_{2,t}^{\omega_{[t]}}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.3)$$

$$y_{i,t}^{\omega_{[t]}} v_{2,t}^{\omega_{[t]}} \leq y_{i,t}^{\omega_{[t]}} r_{t-1}^{a(\omega_{[t]})}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.4)$$

$$B_{i,1}^{\omega_{[1]}} = \tilde{B}_i, \quad i \in I, \omega_{[1]} \in \Omega_{[1]}, \quad (3.5)$$

$$B_{i,t}^{\omega_{[t]}} = B_{i,t-1}^{a(\omega_{[t]})} - \left(v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + r_{t-1}^{a(\omega_{[t]})} y_{i,t}^{\omega_{[t]}} \right), \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.6)$$

$$B_{i,t}^{\omega_{[t]}} \geq 0, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.7)$$

$$b_{i,t}^{\omega_{[t]}} \geq (x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}}) v_{1,t}^{\omega_{[t]}}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.8)$$

$$b_{i,t}^{\omega_{[t]}} \geq (x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}}) r_{t-1}^{a(\omega_{[t]})}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.9)$$

$$x_{i,t}^{\omega_{[t]}} \in \{0, 1\}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.10)$$

$$y_{i,t}^{\omega_{[t]}} \in \{0, 1\}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.11)$$

$$r_t^{\omega_{[t]}} \geq 0, \quad i \in I, t \leq T-1, \omega_{[t]} \in \Omega_{[t]}. \quad (3.12)$$

Model 3.2.1, however, has nonlinearities and therefore is not an SMILP. These nonlinearities are seen:

- (i) in the objective (3.1);
- (ii) in constraint (3.3);
- (iii) in constraint (3.4);
- (iv) in constraint (3.6);
- (v) and in constraint (3.9).

These nonlinearities are due to the terms $r_{t-1}^{a(\omega_{[t]})} y_{i,t}^{\omega_{[t]}}$ and $r_{t-1}^{a(\omega_{[t]})} x_{i,t}^{\omega_{[t]}}$, products of two variables of the model. In order to obtain an SMILP problem, which allows the use of specific algorithms to solve it, these terms can be linearized. It can be done by introducing new variables and constraints. Denoting these new variables as $z_{i,t}^{\omega_{[t]}}$ and

$u_{i,t}^{\omega_{[t]}}$, they are defined by:

$$\begin{aligned} z_{i,t}^{\omega_{[t]}} &= r_{t-1}^{a(\omega_{[t]})} y_{i,t}^{\omega_{[t]}}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ u_{i,t}^{\omega_{[t]}} &= r_{t-1}^{a(\omega_{[t]})} x_{i,t}^{\omega_{[t]}}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}. \end{aligned}$$

It is also necessary to add some new constraints to obtain an SMILP that corresponds to Model 3.2.1. For the variable $z_{i,t}^{\omega_{[t]}}$, the following constraints are added:

$$\begin{aligned} z_{i,t}^{\omega_{[t]}} &\leq M y_{i,t}^{\omega_{[t]}}, & t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ z_{i,t}^{\omega_{[t]}} &\geq 0, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ z_{i,t}^{\omega_{[t]}} &\leq r_{t-1}^{a(\omega_{[t]})}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ z_{i,t}^{\omega_{[t]}} &\geq r_{t-1}^{a(\omega_{[t]})} - (1 - y_{i,t}^{\omega_{[t]}})M, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \end{aligned}$$

where the *Big M* method is used when defining M as a large positive penalty constant. For this model, M is the highest possible bid among all scenarios. The same technique is used for $u_{i,t}^{\omega_{[t]}}$, leading to the following constraints:

$$\begin{aligned} u_{i,t}^{\omega_{[t]}} &\leq M x_{i,t}^{\omega_{[t]}}, \\ u_{i,t}^{\omega_{[t]}} &\geq 0, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ u_{i,t}^{\omega_{[t]}} &\leq r_{t-1}^{a(\omega_{[t]})}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\ u_{i,t}^{\omega_{[t]}} &\geq r_{t-1}^{a(\omega_{[t]})} - (1 - x_{i,t}^{\omega_{[t]}})M, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}. \end{aligned}$$

Finally, considering that the bid distributions for each advertiser are finite, the objective (3.1) can be rewritten as:

$$\max \sum_{t \geq 2} \sum_{\omega_{[t]} \in \Omega_{[t]}} \mathbb{P}(\omega_{[t]}) \left[\sum_{i \in I} v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + \sum_{i \in I} z_{i,t}^{\omega_{[t]}} \right].$$

It is possible now to formulate the problem as an MS-SMILP. A list of the entire nomenclature of this model can be found in Appendix A.

Model 3.2.2 *MS-SMILP for second-price auction*

$$\max_{x,y,u,z,r,B} \sum_{t \geq 2} \sum_{\omega_{[t]} \in \Omega_{[t]}} \mathbb{P}(\omega_{[t]}) \left[\sum_{i \in I} v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + \sum_{i \in I} z_{i,t}^{\omega_{[t]}} \right] \quad (3.13)$$

$$\text{s.t.} \quad \sum_{i \in I} x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}} \leq 1, \quad t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.14)$$

$$u_{i,t}^{\omega_{[t]}} \leq x_{i,t}^{\omega_{[t]}} v_{2,t}^{\omega_{[t]}}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.15)$$

$$y_{i,t}^{\omega_{[t]}} v_{2,t}^{\omega_{[t]}} \leq z_{i,t}^{\omega_{[t]}}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.16)$$

$$B_{i,1}^{\omega_{[1]}} = \tilde{B}_i, \quad i \in I, \omega_{[1]} \in \Omega_{[1]}, \quad (3.17)$$

$$B_{i,t}^{\omega_{[t]}} = B_{i,t-1}^{a(\omega_{[t]})} - \left(v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + z_{i,t}^{\omega_{[t]}} \right), \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.18)$$

$$b_{i,t}^{\omega_{[t]}} \geq (x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}}) v_{1,t}^{\omega_{[t]}}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.19)$$

$$b_{i,t}^{\omega_{[t]}} \geq u_{i,t}^{\omega_{[t]}} + z_{i,t}^{\omega_{[t]}}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.20)$$

$$z_{i,t}^{\omega_{[t]}} \leq M y_{i,t}^{\omega_{[t]}}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.21)$$

$$z_{i,t}^{\omega_{[t]}} \leq r_{t-1}^{a(\omega_{[t]})}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.22)$$

$$z_{i,t}^{\omega_{[t]}} \geq r_{t-1}^{a(\omega_{[t]})} - (1 - y_{i,t}^{\omega_{[t]}}) M, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.23)$$

$$u_{i,t}^{\omega_{[t]}} \leq M x_{i,t}^{\omega_{[t]}}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.24)$$

$$u_{i,t}^{\omega_{[t]}} \leq r_{t-1}^{a(\omega_{[t]})}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.25)$$

$$u_{i,t}^{\omega_{[t]}} \geq r_{t-1}^{a(\omega_{[t]})} - (1 - x_{i,t}^{\omega_{[t]}}) M, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.26)$$

$$x_{i,t}^{\omega_{[t]}} \in \{0, 1\}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.27)$$

$$y_{i,t}^{\omega_{[t]}} \in \{0, 1\}, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.28)$$

$$r_t^{\omega_{[t]}} \geq 0, \quad t \leq T - 1, \omega_{[t]} \in \Omega_{[t]}, \quad (3.29)$$

$$u_{i,t}^{\omega_{[t]}} \geq 0, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.30)$$

$$z_{i,t}^{\omega_{[t]}} \geq 0, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \quad (3.31)$$

$$B_{i,t}^{\omega_{[t]}} \geq 0, \quad i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}. \quad (3.32)$$

Model 3.2.2 does not have fixed recourse: see, for example, that $v_{2,t}^{\omega_{[t]}}$ multiplies $x_{i,t}^{\omega_{[t]}}$ in constraint (3.18). However, it can be seen as a mixed-integer linear program (MILP), with a linear objective function, linear constraints, and some variables restricted to be integer (in this case, binary). Therefore, specific algorithms for this type of problem can be used.

3.2.2 GSP auction

The extension of Model 3.2.2 to a rank-by-bid GSP auction with $|K| \leq |I|$ items (denoted as $k \in K$) is straightforward and is shown as follows.

Model 3.2.3 MS-SMILP for rank-by-bid GSP auction

$$\begin{aligned}
& \max_{x,y,u,z,r,B} \sum_{t \geq 2} \sum_{\omega_{[t]} \in \Omega_{[t]}} \mathbb{P}(\omega_{[t]}) \left[\sum_{i \in I} \sum_{k \in K} v_{k+1,t}^{\omega_{[t]}} x_{i,k,t}^{\omega_{[t]}} + \sum_{i \in I} \sum_{k \in K} z_{i,k,t}^{\omega_{[t]}} \right] \\
& \text{s.t.} \quad \sum_{i \in I} x_{i,k,t}^{\omega_{[t]}} + y_{i,k,t}^{\omega_{[t]}} \leq 1, & k \in K, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,k,t}^{\omega_{[t]}} \leq x_{i,k,t}^{\omega_{[t]}} v_{k+1,t}^{\omega_{[t]}}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad y_{i,k,t}^{\omega_{[t]}} v_{k+1,t}^{\omega_{[t]}} \leq z_{i,k,t}^{\omega_{[t]}}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad \sum_{k \in K} x_{i,k,t}^{\omega_{[t]}} + y_{i,k,t}^{\omega_{[t]}} \leq 1, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad B_{i,1}^{\omega_{[1]}} = \tilde{B}_i, & i \in I, \omega_{[1]} \in \Omega_{[1]}, \\
& \quad B_{i,t}^{\omega_{[t]}} = B_{i,t-1}^{a(\omega_{[t]})} - \left(\sum_{k \in K} v_{k+1,t}^{\omega_{[t]}} x_{i,k,t}^{\omega_{[t]}} + \sum_{k \in K} z_{i,k,t}^{\omega_{[t]}} \right), & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad b_{i,t}^{\omega_{[t]}} \geq (x_{i,k,t}^{\omega_{[t]}} + y_{i,k,t}^{\omega_{[t]}}) v_{k,t}^{\omega_{[t]}}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad b_{i,t}^{\omega_{[t]}} \geq u_{i,k,t}^{\omega_{[t]}} + z_{i,k,t}^{\omega_{[t]}}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad z_{i,k,t}^{\omega_{[t]}} \leq M y_{i,k,t}^{\omega_{[t]}}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad z_{i,k,t}^{\omega_{[t]}} \leq r_{t-1}^{a(\omega_{[t]})}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad z_{i,k,t}^{\omega_{[t]}} \geq r_{t-1}^{a(\omega_{[t]})} - (1 - y_{i,k,t}^{\omega_{[t]}}) M, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,k,t}^{\omega_{[t]}} \leq M x_{i,k,t}^{\omega_{[t]}}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,k,t}^{\omega_{[t]}} \leq r_{t-1}^{a(\omega_{[t]})}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,k,t}^{\omega_{[t]}} \geq r_{t-1}^{a(\omega_{[t]})} - (1 - x_{i,k,t}^{\omega_{[t]}}) M, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad x_{i,k,t}^{\omega_{[t]}} \in \{0, 1\}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad y_{i,k,t}^{\omega_{[t]}} \in \{0, 1\}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad r_t^{\omega_{[t]}} \geq 0, & t \leq T-1, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,k,t}^{\omega_{[t]}} \geq 0, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad z_{i,k,t}^{\omega_{[t]}} \geq 0, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad B_{i,t}^{\omega_{[t]}} \geq 0, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}.
\end{aligned}$$

The constraint $\sum_{k \in K} x_{i,k,t}^{\omega[t]} + y_{i,k,t}^{\omega[t]} \leq 1$ guarantees that advertiser i wins at most one item at t . The entire nomenclature of the model is shown in Appendix A. The model can be also extended to a rank-by-revenue GSP auction, which is shown as follows.

Model 3.2.4 *MS-SMILP for rank-by-bid GSP auction*

$$\begin{aligned}
& \max_{x,y,u,z,r,B} \sum_{t \geq 2} \sum_{\omega[t] \in \Omega[t]} \mathbb{P}(\omega[t]) \left[\sum_{i \in I} \sum_{k \in K} \frac{f_{k+1,t}^{\omega[t]}}{q_{k,t}^{\omega[t]}} x_{i,k,t}^{\omega[t]} + \sum_{i \in I} \sum_{k \in K} z_{i,k,t}^{\omega[t]} \right] \\
& \text{s.t.} \quad \sum_{i \in I} x_{i,k,t}^{\omega[t]} + y_{i,k,t}^{\omega[t]} \leq 1, & k \in K, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad u_{i,k,t}^{\omega[t]} \leq x_{i,k,t}^{\omega[t]} f_{k+1,t}^{\omega[t]}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad y_{i,k,t}^{\omega[t]} v_{k+1,t}^{\omega[t]} \leq z_{i,k,t}^{\omega[t]}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad \sum_{k \in K} x_{i,k,t}^{\omega[t]} + y_{i,k,t}^{\omega[t]} \leq 1, & i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad B_{i,1}^{\omega[1]} = \tilde{B}_i, & i \in I, \omega[1] \in \Omega[1], \\
& \quad B_{i,t}^{\omega[t]} = B_{i,t-1}^{a(\omega[t])} - \left(\sum_{k \in K} \frac{f_{k+1,t}^{\omega[t]}}{q_{k,t}^{\omega[t]}} x_{i,k,t}^{\omega[t]} + \sum_{k \in K} z_{i,k,t}^{\omega[t]} \right), & i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad q_i b_{i,t}^{\omega[t]} \geq (x_{i,k,t}^{\omega[t]} + y_{i,k,t}^{\omega[t]}) f_{k,t}^{\omega[t]}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad b_{i,t}^{\omega[t]} \geq u_{i,k,t}^{\omega[t]} + z_{i,k,t}^{\omega[t]}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad z_{i,k,t}^{\omega[t]} \leq M y_{i,k,t}^{\omega[t]}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad z_{i,k,t}^{\omega[t]} \leq r_{t-1}^{a(\omega[t])}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad z_{i,k,t}^{\omega[t]} \geq r_{t-1}^{a(\omega[t])} - (1 - y_{i,k,t}^{\omega[t]}) M, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad u_{i,k,t}^{\omega[t]} \leq M x_{i,k,t}^{\omega[t]}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad u_{i,k,t}^{\omega[t]} \leq r_{t-1}^{a(\omega[t])}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad u_{i,k,t}^{\omega[t]} \geq r_{t-1}^{a(\omega[t])} - (1 - x_{i,k,t}^{\omega[t]}) M, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad x_{i,k,t}^{\omega[t]} \in \{0, 1\}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad y_{i,k,t}^{\omega[t]} \in \{0, 1\}, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad r_t^{\omega[t]} \geq 0, & t \leq T - 1, \omega[t] \in \Omega[t], \\
& \quad u_{i,k,t}^{\omega[t]} \geq 0, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad z_{i,k,t}^{\omega[t]} \geq 0, & k \in K, i \in I, t \geq 2, \omega[t] \in \Omega[t], \\
& \quad B_{i,t}^{\omega[t]} \geq 0, & i \in I, t \geq 2, \omega[t] \in \Omega[t].
\end{aligned}$$

The main changes in this model are the allocation and the payment rules (see Definition 3.1.2). The parameter $f_{k,t}^{\omega_{[t]}}$ is the k -th highest $f_i = q_i b_i$ for all advertisers and represents the allocation parameter. Therefore, the winner of item k pays the maximum between $\frac{f_{k+1,t}^{\omega_{[t]}}}{q_{k,t}}$ and the reserve price at stage t .

3.3 Analysis of the model

Once the models are presented, it is important to analyze their validity. As Models 3.2.3 and 3.2.4, respectively for rank-by-bid and rank-by-revenue GSP auctions, are extensions of Model 3.2.2, the analysis will be restricted to the last and simplest model.

This analysis will be done by comparing the proposed model to four different ways of choosing the reserve price:

- No reserve price;
- Reserve price equal to the last highest bid;
- Reserve price equal to the smallest possible highest bid;
- Not considering the advertisers' bids.

Starting with the approach of not setting a reserve price, the following is proposed.

Proposition 3.3.1

Denote the optimal value of Model 3.2.2 as $\bar{w}^*$ and the average revenue when not setting a reserve price as $w_{(1)}^*$. When considering the auction described in Definition 3.1.1, the following inequality holds:

$$\bar{w}^* \geq w_{(1)}^*.$$

Proof The approach of not setting a reserve price can be seen as adding the constraint $r_t^{\omega_{[t]}} = 0 \quad \forall t \leq T - 1, \omega_{[t]} \in \Omega_{[t]}$ to Model 3.2.2. Denote $\bar{\mathcal{F}}$ as the feasible set of Model 3.2.2. Moreover, denote $\mathcal{F}^{(1)}$ as the feasible set when adding the mentioned constraint, which has $\pi_{(1)}^*$ as optimal solution. Thus, it can be seen that

$$\mathcal{F}^{(1)} \subseteq \bar{\mathcal{F}} \implies \pi_{(1)}^* \in \bar{\mathcal{F}} \implies \bar{w}^* \geq w_{(1)}^*.$$

□

The second approach used as benchmark is setting the reserve price as the highest bid of the previous auction. It represents the situation where a publisher considers that the highest bid at stage $t + 1$ could be the same as the one at stage t , and therefore fixing the reserve price at this value would make the winner pay exactly her valuation for the item. At stage $t = 1$, as there is no previous auction, it is considered that the reserve price is set as zero. The following proposition is presented for this situation.

Proposition 3.3.2

Denote $\bar{w}^*$ as the optimal value of Model 3.2.2 and $w_{(2)}^*$ as the average revenue when setting a reserve price of 0 at stage $t = 1$ and a reserve price that equals the highest bid at $t - 1$ for stages $t = 2, \dots, T - 1$. When considering the auction described in Definition 3.1.1, the following inequality holds:

$$\bar{w}^* \geq w_{(2)}^*.$$

Proof Setting a reserve price of 0 at stage $t = 1$ means adding the constraint $r_1^{\omega_{[1]}} = 0 \quad \forall \omega_{[1]} \in \Omega_{[1]}$. Moreover, choosing the reserve price as the previous highest bid can be seen as adding the constraint $r_t^{\omega_{[t]}} = v_1^{a(\omega_{[t]})} \quad \forall t = 2, \dots, T - 1, \omega_{[t]} \in \Omega_{[t]}$. Denoting $\mathcal{F}^{(2)}$ as the feasible set when adding the mentioned constraints, which has $\pi_{(2)}^*$ as optimal solution, one can see that

$$\mathcal{F}^{(2)} \subseteq \bar{\mathcal{F}} \implies \pi_{(2)}^* \in \bar{\mathcal{F}} \implies \bar{w}^* \geq w_{(2)}^*.$$

□

The third approach considered is setting the reserve prices at $t = 1, \dots, T - 1$ as the smallest possible highest bid among all scenarios at stage $t + 1$. This approach is based on the following reasoning:

- The publisher wants to make sure that there are no potential winners lost by setting the reserve price too high. Then, the reserve price needs to be smaller or equal to the highest bid in all possible scenarios;
- Still, the publisher wants to minimize her losses due to setting a reserve price that is too low. Therefore, she sets the highest possible value which avoids losing a potential winner: the smallest possible highest bid at that stage.

For this approach, the following is proposed:

Proposition 3.3.3

Denote $\bar{w}^*$ as the optimal value of Model 3.2.2 and $w_{(3)}^*$ as the average revenue when setting a reserve price at stages $t = 1, \dots, T - 1$ that equals the smallest possible highest bid at $t + 1$. When considering the auction described in Definition 3.1.1, the following inequality holds:

$$\bar{w}^* \geq w_{(3)}^*.$$

Proof Setting a reserve price that equals the smallest possible highest bid can be seen as adding the constraint $r_t^{\omega_{[t]}} = \min_{c(\omega_{[t]}) \in \Omega_{[t+1]}} v_{1,t+1}^{c(\omega_{[t]})} \quad \forall t = 1, \dots, T - 1, \omega_{[t]} \in \Omega_{[t]}$. Denoting $\mathcal{F}^{(3)}$ as the feasible set when adding the mentioned constraint, which has $\pi_{(3)}^*$ as optimal solution, one can see that

$$\mathcal{F}^{(3)} \subseteq \bar{\mathcal{F}} \implies \pi_{(3)}^* \in \bar{\mathcal{F}} \implies \bar{w}^* \geq w_{(3)}^*.$$

□

The last benchmark is the case where the budget constraints are not taken into account by the publishers when optimizing the reserve prices. Keeping the rules described in Definition 3.1.1, when an advertiser's budget left is not enough to pay the item's price, the publisher may lose a winner. A proposition for this case is presented as follows.

Proposition 3.3.4

Denote $\bar{w}^*$ as the optimal value of Model 3.2.2 and $w_{(4)}^*$ as the average revenue when optimizing the reserve price while ignoring the advertisers' budgets. When considering the auction described in Definition 3.1.1, the following inequality holds:

$$\bar{w}^* \geq w_{(4)}^*.$$

Proof Even if the reserve prices are set as the optimal ones when ignoring the advertisers' budget constraints, they still need to be respected. The other constraints of Model 3.2.2 also need to be respected, as they represent the rules of the auction. This way, this is a feasible solution of Model 2.2.2, but not necessarily the optimal one. Denoting $\pi_{(4)}^*$ as this solution, one can see that

$$\pi_{(4)}^* \in \bar{\mathcal{F}} \implies \bar{w}^* \geq w_{(4)}^*.$$

□

Once it is shown that the optimal value of Model 3.2.2 is at least as good as the revenues using the proposed approaches, it is important to prove that it may lead to strictly better outcomes. This will be done in the next chapter by presenting a numerical example and results of simulations.

4 | Numerical example and simulations

This chapter focuses on the application of the proposed model for single-item second-price auctions. It is divided in two sections: the first one presents a numerical example and the second one shows the results of simulations.

The numerical example allows the reader to understand all the computations and analyses of the model. After presenting its characteristics, all the steps to compute its performance are shown and it is compared with the four other benchmarks considered before. It illustrates a case where the optimal reserve prices obtained from the proposed model lead to a better revenue for the publisher than the other approaches.

In the second section, results from simulations are presented, also focusing on evaluating its performance and comparing it with other approaches. Moreover, it is shown that the number of simplex iterations to solve the problem grows exponentially with the number of scenarios.

4.1 Numerical example

After formulating and analyzing the model, a numerical example for the single-item second-price auction will be shown (as a reminder, the models for GSP auctions are extensions of this model). A simple example was selected to clearly illustrate the results and to have a reasonable computational time. Two main analyses will be done for this example. First, the performance of the solutions of the model (VSS and $EVPI$) will be analyzed. Then, a comparison of the results of the model with the approaches presented in Section 3.3 will be made (showing a case where it indeed has a better performance than these approaches, as a complement to Propositions 3.3.1 to 3.3.4).

4.1.1 Characteristics

The situation considered involves two advertisers, named Advertiser A and Advertiser B, bidding in three auctions (at stages $t = 2, 3, 4$). Their budgets and possible bids are presented in Table 4.1.

	Budget	$t = 2$		$t = 3$		$t = 4$	
		Bids	Prob	Bids	Prob	Bids	Prob
Advertiser A	\$11.70	\$2.50	100%	\$3.00	50%	\$3.50	50%
				\$4.00	50%	\$6.00	50%
Advertiser B	\$11.00	\$2.00	50%	\$3.50	100%	\$5.80	100%
		\$3.00	50%				

Table 4.1: Budgets and possible bids for advertisers A and B.

From Table 4.1, it is possible to obtain the parameters $\tilde{B}_1 = 11.7$ and $\tilde{B}_2 = 11$. The other parameters are presented in the scenario tree shown in Figure 4.1.




Figure 4.1: Scenario tree with the chosen parameters (numerical example).

From Figure 4.1, it is possible to see that there are eight possible scenarios at $t = 4$ (nodes 4.1 to 4.8) happening with probability 0.125. The reserve prices are chosen at

stages $t = 1, 2, 3$ and the winners of each auction are decided at $t = 2, 3, 4$.

4.1.2 Performance

The performance of the model will be evaluated according to two main metrics: the expected value of perfect information, $EVPI$, and the value of the stochastic solution at stage 4, VSS_4 . As explained in Subsection 2.2.2, it requires the computation of the following solutions: the stochastic programming solution (SP); the wait-and-see solution (WS); and the expected value of using the EV solution at stage 4 (EEV_4).

These solutions were obtained by solving different mixed-integer linear programming (MILP) problems. The problems were implemented using the programming language Julia and solved using the Gurobi solver, which generally solves MILP problems using a linear-programming based branch-and-bound algorithm.

SP solution

The SP solution was obtained by solving Model 3.2.2. The SP solution and the number of simplex iterations taken to solve the model are shown in Table 4.2.

Optimal value (SP)	Simplex iterations
\$11.65	10622

Table 4.2: SP solution and number of simplex iterations (numerical example).

Moreover, the optimal reserve prices chosen at $t = 1, 2, 3$ are presented in Table 4.3, where the labels of the nodes are the same as the ones shown in Figure 4.1.

$t = 1$		$t = 2$		$t = 3$	
Node	Reserve price	Node	Reserve price	Node	Reserve price
1.1	\$2.40	2.1	\$3.50	3.1	\$5.80
				3.2	\$5.80
		2.2	\$3.50	3.3	\$5.00
				3.4	\$5.80

Table 4.3: Optimal reserve prices of the SP solution (numerical example).

WS solution

As mentioned in Subsection 2.2.2, the wait-and-see value considered in this thesis is the solution of the perfect foresight model. Calling S the set of possible scenarios of the problem ($S = \{4.1, 4.2, 4.3, 4.4, 4.5, 4.6, 4.7, 4.8\}$), the perfect foresight model for Model 2.2.2 is shown as follows, where the superscript s represents each of the scenarios.

Model 4.1.1 *Perfect foresight model for second-price auction*

$$\begin{aligned}
& \max_{x,y,u,z,r,B} \sum_{s \in S} \sum_{t \geq 2} \mathbb{P}(s) \left[\sum_{i \in I} v_{2,t}^s x_{i,t}^s + \sum_{i \in I} z_{i,t}^s \right] \\
& \text{s.t.} \quad \sum_{i \in I} x_{i,t}^s + y_{i,t}^s \leq 1, & t \geq 2, s \in S, \\
& \quad u_{i,t}^s \leq x_{i,t}^s v_{2,t}^s, & i \in I, t \geq 2, s \in S, \\
& \quad y_{i,t}^s v_{2,t}^s \leq z_{i,t}^s, & i \in I, t \geq 2, s \in S, \\
& \quad B_{i,1}^s = \tilde{B}_i, & i \in I, s \in S, \\
& \quad B_{i,t}^s = B_{i,t-1}^s - \left(v_{2,t}^s x_{i,t}^s + z_{i,t}^s \right), & i \in I, t \geq 2, s \in S, \\
& \quad b_{i,t}^s \geq (x_{i,t}^s + y_{i,t}^s) v_{1,t}^s, & i \in I, t \geq 2, s \in S, \\
& \quad b_{i,t}^s \geq u_{i,t}^s + z_{i,t}^s, & i \in I, t \geq 2, s \in S, \\
& \quad z_{i,t}^s \leq M y_{i,t}^s, & i \in I, t \geq 2, s \in S, \\
& \quad z_{i,t}^s \leq r_{t-1}^s, & i \in I, t \geq 2, s \in S, \\
& \quad z_{i,t}^s \geq r_{t-1}^s - (1 - y_{i,t}^s) M, & i \in I, t \geq 2, s \in S, \\
& \quad u_{i,t}^s \leq M x_{i,t}^s, & i \in I, t \geq 2, s \in S, \\
& \quad u_{i,t}^s \leq r_{t-1}^s, & i \in I, t \geq 2, s \in S, \\
& \quad u_{i,t}^s \geq r_{t-1}^s - (1 - x_{i,t}^s) M, & i \in I, t \geq 2, s \in S, \\
& \quad x_{i,t}^s \in \{0, 1\}, & i \in I, t \geq 2, s \in S, \\
& \quad y_{i,t}^s \in \{0, 1\}, & i \in I, t \geq 2, s \in S, \\
& \quad r_t^s \geq 0, & t \leq T - 1, s \in S, \\
& \quad u_{i,t}^s \geq 0, & i \in I, t \geq 2, s \in S, \\
& \quad z_{i,t}^s \geq 0, & i \in I, t \geq 2, s \in S, \\
& \quad B_{i,t}^s \geq 0, & i \in I, t \geq 2, s \in S.
\end{aligned}$$

Table 4.4 shows the WS solution and the number of iterations taken to compute it.

Optimal value (WS)	Simplex iterations
\$12.1375	422

Table 4.4: WS solution and number of simplex iterations (numerical example).

It can be seen that the WS solution outperforms the SP solution, as expected (see Proposition 2.2.1). Moreover, solving this model requires solving 8 MILPs that are much simpler than the stochastic problem, explaining why less iterations are required to compute WS than to compute SP .

EEV_4 solution

As detailed in Subsection 2.2.2, computing EEV_4 demands solving the expected value (EV) problem. The EV problem for the second-price auction is shown as follows.

Model 4.1.2 *Expected value problem for second-price auction*

$$\begin{aligned}
& \max_{x,y,u,z,r,B} \sum_{t \geq 2} \left[\sum_{i \in I} \bar{v}_{2,t} x_{i,t} + \sum_{i \in I} z_{i,t} \right] \\
& \text{s.t.} \quad \sum_{i \in I} x_{i,t} + y_{i,t} \leq 1, & t \geq 2, \\
& \quad u_{i,t} \leq x_{i,t} \bar{v}_{2,t}, & i \in I, t \geq 2, \\
& \quad y_{i,t} \bar{v}_{2,t} \leq z_{i,t}, & i \in I, t \geq 2, \\
& \quad B_{i,1} = \tilde{B}_i, & i \in I, \\
& \quad B_{i,t} = B_{i,t-1} - \left(\bar{v}_{2,t} x_{i,t} + z_{i,t} \right), & i \in I, t \geq 2, \\
& \quad \bar{b}_{i,t} \geq (x_{i,t} + y_{i,t} \bar{v}_{1,t}), & i \in I, t \geq 2, \\
& \quad \bar{b}_{i,t} \geq u_{i,t} + z_{i,t}, & i \in I, t \geq 2, \\
& \quad z_{i,t} \leq M y_{i,t}, & i \in I, t \geq 2, \\
& \quad z_{i,t} \leq r_{t-1}, & i \in I, t \geq 2, \\
& \quad z_{i,t} \geq r_{t-1} - (1 - y_{i,t})M, & i \in I, t \geq 2, \\
& \quad u_{i,t} \leq M x_{i,t}, & i \in I, t \geq 2, \\
& \quad u_{i,t} \leq r_{t-1}, & i \in I, t \geq 2, \\
& \quad u_{i,t} \geq r_{t-1} - (1 - x_{i,t})M, & i \in I, t \geq 2, \\
& \quad x_{i,t} \in \{0, 1\}, & i \in I, t \geq 2, \\
& \quad y_{i,t} \in \{0, 1\}, & i \in I, t \geq 2, \\
& \quad r_t \geq 0, & t \leq T - 1, \\
& \quad u_{i,t} \geq 0, & i \in I, t \geq 2, \\
& \quad z_{i,t} \geq 0, & i \in I, t \geq 2, \\
& \quad B_{i,t} \geq 0, & i \in I, t \geq 2.
\end{aligned}$$

In Model 4.1.2, $\bar{b}_t = \mathbb{E}_{\omega_{[t]}}[b_t^{\omega_{[t]}}]$, $\bar{v}_{1,t} = \mathbb{E}_{\omega_{[t]}}[v_{1,t}^{\omega_{[t]}}]$, and $\bar{v}_{2,t} = \mathbb{E}_{\omega_{[t]}}[v_{2,t}^{\omega_{[t]}}]$ for all $t \geq 2$. Denoting r_t^* as the optimal value of r_t for the expected value problem, EEV_4 can be found by solving the following model.

Model 4.1.3 *Model for computing EEV_4 (second-price auction)*

$$\begin{aligned}
& \max_{x,y,u,z,r,B} \sum_{t \geq 2} \sum_{\omega_{[t]} \in \Omega_{[t]}} \mathbb{P}(\omega_{[t]}) \left[\sum_{i \in I} v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + \sum_{i \in I} z_{i,t}^{\omega_{[t]}} \right] \\
& \text{s.t.} \quad \sum_{i \in I} x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}} \leq 1, & t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,k,t}^{\omega_{[t]}} \leq x_{i,k,t}^{\omega_{[t]}} v_{k+1,t}^{\omega_{[t]}}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad y_{i,k,t}^{\omega_{[t]}} v_{k+1,t}^{\omega_{[t]}} \leq z_{i,k,t}^{\omega_{[t]}}, & k \in K, i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad B_{i,1}^{\omega_{[1]}} = \tilde{B}_i, & i \in I, \omega_{[1]} \in \Omega_{[1]}, \\
& \quad B_{i,t}^{\omega_{[t]}} = B_{i,t-1}^{a(\omega_{[t]})} - \left(v_{2,t}^{\omega_{[t]}} x_{i,t}^{\omega_{[t]}} + z_{i,t}^{\omega_{[t]}} \right), & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad b_{i,t}^{\omega_{[t]}} \geq (x_{i,t}^{\omega_{[t]}} + y_{i,t}^{\omega_{[t]}}) v_{1,t}^{\omega_{[t]}}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad b_{i,t}^{\omega_{[t]}} \geq u_{i,t}^{\omega_{[t]}} + z_{i,t}^{\omega_{[t]}}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad z_{i,t}^{\omega_{[t]}} \leq M y_{i,t}^{\omega_{[t]}}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad z_{i,t}^{\omega_{[t]}} \leq r_{t-1}^{a(\omega_{[t]})}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad z_{i,t}^{\omega_{[t]}} \geq r_{t-1}^{a(\omega_{[t]})} - (1 - y_{i,t}^{\omega_{[t]}}) M, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,t}^{\omega_{[t]}} \leq M x_{i,t}^{\omega_{[t]}}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,t}^{\omega_{[t]}} \leq r_{t-1}^{a(\omega_{[t]})}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,t}^{\omega_{[t]}} \geq r_{t-1}^{a(\omega_{[t]})} - (1 - x_{i,t}^{\omega_{[t]}}) M, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad x_{i,t}^{\omega_{[t]}} \in \{0, 1\}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad y_{i,t}^{\omega_{[t]}} \in \{0, 1\}, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad r_t^{\omega_{[t]}} = r_t^*, & t \leq T - 1, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad u_{i,t}^{\omega_{[t]}} \geq 0, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad z_{i,t}^{\omega_{[t]}} \geq 0, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}, \\
& \quad B_{i,t}^{\omega_{[t]}} \geq 0, & i \in I, t \geq 2, \omega_{[t]} \in \Omega_{[t]}.
\end{aligned}$$

In Model 4.1.3, the constraint $r_t^{\omega_{[t]}} = r_t^*$ ensures that the reserve prices of the EV solution are fixed for $t = 1, 2, 3$ when solving the model. First, the EV solution and the

number of simplex iterations required to compute it are shown in Table 4.5.

Optimal value (EV)	Simplex iterations
\$11.80	19

Table 4.5: EV solution and number of simplex iterations (numerical example).

The small number of iterations demanded to find the EV solution can be easily explained: while 8 MILPs were solved to find WS , only one MILP needs to be solved to find EV . Although its optimal value is greater than the optimal value of the stochastic model, it solves a different objective. Therefore, to truly measure its performance, it is necessary to compute the EEV solution (in this case, the EEV_4 solution), which is shown in Table 4.6, as well as the number of iterations required to find it.

Optimal value (EEV_4)	Simplex iterations
\$11.2625	125

Table 4.6: EEV_4 solution and number of simplex iterations (numerical example).

The number of iterations to obtain the EEV_4 solution is also smaller when compared to the WS solution. In order to compute EEV_4 , it is necessary to solve, for every $s \in S$, a linear program over x, y, u, z , and B , while to compute WS it is necessary to solve $|S|$ linear programs over these variables and also r . As expected (see Proposition 2.2.2), the EEV_4 solution is smaller than SP .

The optimal values and the number of iterations to compute each solution are shown in Figures 4.2 and 4.3, respectively.




Figure 4.2: Optimal value of each solution (numerical example). As stated in Propositions 2.2.1 and 2.2.2, $EEV_4 \leq SP \leq WS$.




Figure 4.3: Number of simplex iterations for each solution (numerical example). Computing SP is significantly more computationally demanding than EEV_4 and WS .

EVPI and VSS_4

Once SP , WS , and EEV_4 are computed, it is possible to obtain $EVPI$ and VSS_4 . Reminding that Model 3.2.2 is a maximization problem, the expected value of perfect information is given by

$$EVPI = WS - SP = 0.4875,$$

meaning that the average revenue could be increased in \$0.4875 (or approximately 4%) if the information about the advertisers' bids were certain. The value of the stochastic solution at $t = 4$ is given by

$$VSS_4 = SP - EEV_4 = 0.3875,$$

which means that the stochastic solution leads to an increase of \$0.3875 in the average revenue when compared to the expected value solution. These values are summarized in Figure 4.4.




Figure 4.4: Performance of the numerical example.

4.1.3 Comparisons

Once the performance of the model is evaluated for the given example, it will now be compared with the four approaches detailed in Section 3.3. First, when considering the case where no reserve price is set, the expected revenue at $t = 2$ is \$2.25, at $t = 3$ is \$3.25, and at $t = 4$ is \$4.65, leading to an expected total revenue of \$10.15.

The second approach considered is setting the reserve price as the previous winning bid. The reserve prices chosen in this approach are shown in Table 4.7.

$t = 1$		$t = 2$		$t = 3$	
Node	Reserve price	Node	Reserve price	Node	Reserve price
1.1	\$0.00	2.1	\$2.50	3.1	\$3.50
				3.2	\$4.00
		2.2	\$3.00	3.3	\$3.50
				3.4	\$4.00

Table 4.7: Reserve prices for the second approach (numerical example).

In the second approach, the expected revenue at $t = 2$ is \$2.25, at $t = 3$ is \$3.25, and at $t = 4$ is \$4.775, with an expected total revenue of \$10.275.

Moving to the third approach, in which the reserve price is set as the smallest possible highest bid, the reserve prices of the example are presented in Table 4.8.

$t = 1$	$t = 2$	$t = 3$
Reserve price	Reserve price	Reserve price
\$2.50	\$3.50	\$5.80

Table 4.8: Reserve prices for the third approach (numerical example).

The third approach leads to an expected revenue at $t = 2$ of \$2.50, at $t = 3$ of \$3.50, and at $t = 4$ of \$4.35, with an expected total revenue of \$10.35.

Finally, the fourth approach (optimizing the reserve prices while ignoring the budget constraints) leads to the solution shown in Table 4.9.

$t = 1$		$t = 2$		$t = 3$	
Node	Reserve price	Node	Reserve price	Node	Reserve price
1.1	\$2.50	2.1	\$3.50	3.1	\$5.80
				3.2	\$5.80
		2.2	\$3.50	3.3	\$5.80
				3.4	\$5.80

Table 4.9: Reserve prices for the fourth approach (numerical example).

It is possible to see that, in this case, the solution is the same as the one seen for the third approach. Therefore, it leads to the same expected revenues at all time steps and an expected total revenue of \$10.35.

After computing the results of the four approaches used as benchmarks, it is possible to see that the solution of Model 2.2.2 leads to the highest revenue (\$11.65). So, besides

always leading to greater or equal revenues when compared to these approaches (see Propositions 3.3.1 to 3.3.4), a case where the model actually leads to a better revenue is shown. The results of the model and of the benchmarks are shown in Figure 4.5. Detailed plots for each scenario are shown in Appendix B.




Figure 4.5: Average revenue for each approach (numerical example). Model 2.2.2 led to the highest average revenue over the three auctions (\$11.65).

4.2 Simulations

After presenting the numerical example in details, some simulations were done for different numbers of auctions and sets of randomly-generated parameters (distributions of the bids and budgets). The way these parameters were generated can be seen in Appendix C. In this section, the goal is to analyze the average performance of the model, to observe how the number of simplex iterations varies with the number of scenarios, and to compare the model with other approaches to choose the reserve price.

As in the numerical example, the implementation of the models was done using the programming language Julia and they were solved using the Gurobi solver. The computer used for the simulations has an Intel(R) Core(TM) i7-6500U CPU with four threads at 2.50GHz and 8GB of RAM. As explained in Subsection 2.2.1, solving the stochastic program is usually more computationally demanding than computing the other solutions (which is the reason why they are considered as benchmarks). The stochastic program could be solved in a few seconds when there were up to 8 scenarios. When 16 scenarios were considered, the computation always exceeded the maximum time of 5 minutes chosen for the problem, usually reaching more than a million simplex iterations. This way, 45 simulations were done, all of them with two advertisers per auction: 15 for the case with 2 scenarios (and one auction), 15 with 4 scenarios (and two auctions), and 15 with 8 scenarios (and three auctions). For these numbers of

scenarios, the average number of simplex iterations required to compute the solution of Model 2.2.2 is shown in Figure 4.6.




Figure 4.6: Average number of simplex iterations (simulations). One can see an exponential increase with the number of scenarios.

One can observe an exponential increase of the number of simplex iterations with the number of scenarios, explaining why it was not possible to solve it for 16 scenarios within 5 minutes. This is due to the fact that the size of optimization problems derived from multi-stage stochastic programs grows exponentially as the number of scenarios increases (with more stages or more uncertainties within a stage).

4.2.1 Performance

The average solutions for 2, 4 and 8 scenarios are shown in Figure 4.7.




Figure 4.7: Average solutions (simulations). For all the numbers of scenarios, the inequality $\overline{EEV}_H \leq \overline{SP} \leq \overline{WS}$ is satisfied.

As stated in Propositions 2.2.1 and 2.2.2, for all the numbers of scenarios, $EEV_H \leq SP \leq WS$, where H denotes the final stage. Therefore, their averages also follow the inequalities.

The absolute difference between the averages of SP and EEV_H , respectively denoted by $\overline{SP}$ and $\overline{EEV}_H$, was the greatest when the number of scenarios was 4. This difference represents the average value of stochastic solution at stage H , $\overline{VSS}_H$. For the simulations with 2 scenarios (one auction and two stages), $\overline{VSS}_2 = \$0.32$; for 4 scenarios (two auctions and three stages), $\overline{VSS}_3 = \$0.92$; and for 8 scenarios (three auctions and four stages), $\overline{VSS}_4 = \$0.18$. The greatest relative difference between $\overline{SP}$ and $\overline{EEV}_H$ was also obtained with 4 scenarios: the stochastic solution led to an average increase of approximately 20% in revenues when compared to $\overline{EEV}_3$.

When comparing $\overline{WS}$ (the average of WS among the 15 simulations) and $\overline{SP}$, the greater absolute difference between these solutions (representing the average expected value of perfect information, $\overline{EVPI}$) occurred with 8 scenarios. For the simulations with 2 scenarios, $\overline{EVPI} = \$0.13$; for 4 scenarios, $\overline{EVPI} = \$0.44$; and for 8 scenarios, $\overline{EVPI} = \$0.49$. The greatest relative difference between $\overline{WS}$ and $\overline{SP}$, however, occurred with 4 scenarios: having perfect information about the future would lead to an average increase of approximately 8% in the publisher's revenue.

For each number of scenarios, $\overline{VSS}_H$ and $\overline{EVPI}$ are shown in Figure 4.8.




Figure 4.8: Average performance (simulations). $\overline{EVPI}$ is the greatest with 4 scenarios, while $\overline{VSS}_H$ is the greatest with 8 scenarios.

4.2.2 Comparisons

For each number of scenarios, the average revenue when using the optimal solution of Model 2.2.2 was compared to the average revenue when using the four approaches presented in Section 3.3. This comparison is represented in Figure 4.9.




Figure 4.9: Average revenue for each approach (simulations). For all the numbers of scenarios, the highest revenue was obtained when solving Model 2.2.2.

For all the numbers of scenarios, it was possible to see that the revenues when using approaches 1 and 2 were similar, and approaches 3 and 4 led to the same revenue. Using the solution of Model 2.2.2 led to the greatest average revenue for all the numbers of scenarios. For 2 scenarios, it led to an average revenue of \$2.09, which is \$0.22 (or 11.7%) greater than the second-highest revenue, obtained using approaches 1 and 2. For 4 scenarios, an average revenue of \$5.51 was obtained when solving the model, representing an increase of \$0.57 (or 11.5%) from the second-highest revenue, using approach 1. Finally, for 8 scenarios, the average revenue when solving the model was \$11.72, which is \$0.76 (or 6.9%) greater than the second-highest revenue, obtained using approaches 3 and 4. From that, it is possible to see that taking the advertisers' budgets into account when choosing the reserve prices in second-price auctions may indeed lead to higher revenues for the publisher.

5 | Conclusion

This chapter concludes the work done in this thesis.

First, it summarizes what was done in the previous chapters, briefly explaining their contents and highlighting what was added to the literature.

After that, some ideas for future works on the framework of this thesis are presented, both to improve the models themselves and to add value going in other directions.

5.1 Summary

This thesis presented a study of the optimization of the reserve price in sponsored search auctions. Differently from what has been shown so far in the literature, a multi-stage stochastic mixed-integer linear program was proposed, allowing taking the advertisers' budgets into account when choosing the reserve price. The context and the state of the art on the topic are presented in Chapter 1, as well as the main contributions of this work.

In Chapter 2, the mathematical background necessary to understand the development of this work was presented. First, concepts related to Game Theory were explained. Basic game-theoretic concepts relevant for the comprehension of auctions were presented, followed by the usual approach of modelling auctions as repeated Bayesian games, focusing on two types of auctions: second-price and GSP auctions. The introduction of reserve prices in the mechanism of these auctions was also detailed. After that, as the model presented in this thesis is an MS-SMILP, concepts of Stochastic Programming were also explained. Starting with the simpler case of two-stage programs and then moving to multi-stage cases, their meanings and formulations were detailed and it was shown how their performances can be evaluated.

Chapter 3 was devoted to proposing and detailing the MS-SMILPs to optimize the reserve price in second-price and GSP auctions. First, the main assumptions were explained and the rules of the auctions considered in this thesis were formalized. Then, the formulation of the model was presented for second-price auctions, and its extensions to rank-by-bid and rank-by-revenue GSP auctions were also shown. In this part, all the variables, parameters, constraints, and objective functions were explained in details, as well as the manipulations done to obtain the models as mixed-integer linear programs. Chapter 3 ended by proposing and proving that the solution of the proposed model for second-price auctions always leads to revenues for the publisher that are greater or equal to the revenues when using four other approaches to choose the reserve price.

Finally, a numerical example and results of simulations were presented in Chapter 4. First, the example was shown in details, allowing the reader to understand all the computations and analyses. Its characteristics were presented and schematized in a scenario tree. Then, its performance was evaluated and the way that each metric was computed was explained. The results of the example were then compared to the results of the four approaches previously considered. This way, not only it was proven that the revenues using the proposed model are always greater or equal to the revenues using these four approaches, but also a case where it leads to a strictly greater revenue was shown. After that, the performance of some simulations was also analyzed, and it was possible to see gains when using the proposed the model in comparison to the other four approaches. Moreover, the exponential increase of the number of iterations with the number of scenarios was highlighted and explained.

It is possible to conclude, then, that the main objective of this thesis was attained. The formulation of second-price and GSP auctions with reserve price as multi-stage stochastic mixed-integer linear programs allowed including the advertisers' budgets as constraints in the model while maintaining the main characteristics and rules of these auctions. Under the proposed hypotheses, the solution of the model for second-price auctions was proven to always lead to revenues for the publisher that are greater or equal to the revenues when using other approaches, including optimizing the reserve price without considering the budget constraints. Moreover, by showing a numerical example and analyzing results of simulations, it was possible to see an increase in the publisher's revenues when using the solution of the proposed MS-SMILP.

5.2 Future directions

While this thesis has some noticeable contributions to the current literature on sponsored search auctions, there are many improvements that can be done to enhance its results. Some of these improvements are listed here.

The first improvement is related to the assumptions made to build the model. In this

thesis, it is considered that if an advertiser has the highest bid and her bid is greater than her budget left, the auction has no winner. To make the model closer to reality, one could assume that, in this case, this advertiser would not participate in the auction or that she would bid her budget left. This however requires further work to include these instances in a multi-stage stochastic program.

Moreover, while this thesis presented simulations of second-price auctions with up to eight scenarios, given the exponential increase of the number of iterations with the number of scenarios and the limitations of the machine used for these simulations, it was not possible to go further than that. One could work on simulating a model to optimize the reserve price considering the budget constraints for larger scales. This would allow considering more time steps, more advertisers, more uncertainties, and GSP auctions (with more than one item).

Even if the budget constraint is one of the most common practical constraints in sponsored search auctions, there are other instances that could be considered. One example concerns adaptive pacing strategies, in which advertisers adjust the pace at which they spend their budget according to their expenditures. One could work on adding these other practical constraints in a model to optimize the reserve price.

Finally, it is important to highlight that this thesis proposes a model to optimize the reserve price once the publisher has access to discrete bid distributions for each advertiser. Further studies on how to estimate these distributions would provide a great complement to the proposed model.

Appendices

A | Nomenclature of the optimization model

A.1 Second-price auction

Parameters

- $\tilde{B}_i$: total budget of advertiser i ;
- $b_{i,t}^{\omega_{[t]}}$: bid of advertiser i at time t for realization $\omega_{[t]}$;
- $v_{1,t}^{\omega_{[t]}}$: highest bid at time t for realization $\omega_{[t]}$;
- $v_{2,t}^{\omega_{[t]}}$: second-highest bid at time t for realization $\omega_{[t]}$;
- M : highest possible bid among all scenarios.

Variables

- $x_{i,t}^{\omega_{[t]}}$: binary variable which is equal to 1 if advertiser i wins auction t for realization $\omega_{[t]}$ and pays the second-highest bid, and is equal to 0 otherwise;
- $y_{i,t}^{\omega_{[t]}}$: binary variable which is equal to 1 if advertiser i wins auction t for realization $\omega_{[t]}$ and pays the reserve price, and is equal to 0 otherwise;
- $r_t^{\omega_{[t]}}$: reserve price at time t for realization $\omega_{[t]}$;
- $B_{i,t}^{\omega_{[t]}}$: budget left for advertiser i at the end of period t for realization $\omega_{[t]}$;
- $u_{i,t}^{\omega_{[t]}}$ and $z_{i,t}^{\omega_{[t]}}$: additional variables to remove nonlinearities.

A.2 GSP auction

Parameters

- $\tilde{B}_i$: total budget of advertiser i ;
- $b_{i,t}^{\omega_{[t]}}$: bid of advertiser i at time t for realization $\omega_{[t]}$;
- $v_{k,t}^{\omega_{[t]}}$: k -th highest bid at time t for realization $\omega_{[t]}$;
- M : highest possible bid among all scenarios;
- $f_{k,t}^{\omega_{[t]}}$ (only for rank-by-revenue GSP auctions): k -th highest allocation parameter $f_i = q_i b_i$ at time t for realization $\omega_{[t]}$, where q_i is the advertiser-specific term of the click-through rate and b_i is the bid of advertiser i ;
- $q_{k,t}^{\omega_{[t]}}$ (only for rank-by-revenue GSP auctions): k -th highest advertiser-specific term of the click-through rate at time t for realization $\omega_{[t]}$.

Variables

- $x_{i,k,t}^{\omega_{[t]}}$: binary variable which is equal to 1 if advertiser i wins item k in auction t for realization $\omega_{[t]}$ and pays the following highest bid, and is equal to 0 otherwise;
- $y_{i,k,t}^{\omega_{[t]}}$: binary variable which is equal to 1 if advertiser i wins item k in auction t for realization $\omega_{[t]}$ and pays the reserve price, and is equal to 0 otherwise;
- $r_t^{\omega_{[t]}}$: reserve price at time t for realization $\omega_{[t]}$;
- $B_{i,t}^{\omega_{[t]}}$: budget left for advertiser i at the end of period t for realization $\omega_{[t]}$;
- $u_{i,k,t}^{\omega_{[t]}}$ and $z_{i,k,t}^{\omega_{[t]}}$: additional variables to remove nonlinearities.

B | Results of the numerical example for each scenario




Figure B.1: Revenue for each approach (scenario 4.1).




Figure B.2: Revenue for each approach (scenario 4.2).




Figure B.3: Revenue for each approach (scenario 4.3).




Figure B.4: Revenue for each approach (scenario 4.4).




Figure B.5: Revenue for each approach (scenario 4.5).




Figure B.6: Revenue for each approach (scenario 4.6).




Figure B.7: Revenue for each approach (scenario 4.7).




Figure B.8: Revenue for each approach (scenario 4.8).

C | Generation of parameters for simulations

Noting $g_{l,h}$ randomly-generated integers between -5 and 5 , for $l = 1, \dots, 24$ and simulations $h = 1, \dots, 15$ for each number of scenarios, the parameters were generated as shown in Tables C.1, C.2, and C.3.

	Budget	$t = 2$	
		Bids	Prob
Advertiser A	\$11.70 + $0.1g_{1,h}$	\$2.50 + $0.1g_{2,h}$	100%
Advertiser B	\$11.00 + $0.1g_{3,h}$	\$2.00 + $0.1g_{4,h}$	50%
		\$3.00 + $0.1g_{5,h}$	50%

Table C.1: Parameters generation (2 scenarios).

	Budget	$t = 2$		$t = 3$	
		Bids	Prob	Bids	Prob
Advertiser A	\$11.70 + $0.1g_{6,h}$	\$2.50 + $0.1g_{7,h}$	100%	\$3.00 + $0.1g_{8,h}$	50%
				\$4.00 + $0.1g_{9,h}$	50%
Advertiser B	\$11.00 + $0.1g_{10,h}$	\$2.00 + $0.1g_{11,h}$	50%	\$3.50 + $0.1g_{12,h}$	100%
		\$3.00 + $0.1g_{13,h}$	50%		

Table C.2: Parameters generation (4 scenarios).

	Budget	$t = 2$		$t = 3$		$t = 4$	
		Bids	Prob	Bids	Prob	Bids	Prob
Advertiser A	\$11.70 + 0.1 $g_{14,h}$	\$2.50 + 0.1 $g_{15,h}$	100%	\$3.00 + 0.1 $g_{16,h}$	50%	\$3.50 + 0.1 $g_{17,h}$	50%
				\$4.00 + 0.1 $g_{18,h}$	50%	\$6.00 + 0.1 $g_{19,h}$	50%
Advertiser B	\$11.00 + 0.1 $g_{20,h}$	\$2.00 + 0.1 $g_{21,h}$	50%	\$3.50 + 0.1 $g_{22,h}$	100%	\$5.80 + 0.1 $g_{23,h}$	100%
		\$3.00 + 0.1 $g_{24,h}$	50%				

Table C.3: Parameters generation (8 scenarios).

Bibliography

- [1] Interactive Advertising Bureau - PWC. *Internet Advertising Revenue Report - 2020 full year results*. https://www.iab.com/wp-content/uploads/2021/04/IAB_2020-Internet-Advertising-Revenue-Report-Webinar_4.7.21-PwC-1.pdf. Accessed: 26/02/2022.
- [2] Statista Research Department. *Digital advertising spending worldwide 2018-2023*. <https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/>. Accessed: 26/02/2022.
- [3] Ethan Cramer-Flood. *Worldwide Digital Ad Spending Year-End Update*. Ed. by eMarketer. <https://content-na1.emarketer.com/worldwide-digital-ad-spending-year-end-update>. Accessed: 27/02/2022.
- [4] Kevin McLaughlin and Daniel Friedman. “Online ad auctions: An experiment”. In: (2016).
- [5] Brendan Lucier, Renato Paes Leme, and Éva Tardos. “On revenue in the generalized second price auction”. In: *Proceedings of the 21st international conference on World Wide Web*. 2012, pp. 361–370.
- [6] Gang Bai, Zhihui Xie, and Liang Wang. “Practical constrained optimization of auction mechanisms in E-commerce sponsored search advertising”. In: *arXiv preprint arXiv:1807.11790* (2018).
- [7] Yi Wang et al. “Peacock: Learning long-tail topic features for industrial applications”. In: *ACM Transactions on Intelligent Systems and Technology (TIST)* 6.4 (2015), pp. 1–23.
- [8] Weiran Shen et al. “Reinforcement mechanism design: With applications to dynamic pricing in sponsored search auctions”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34. 02. 2020, pp. 2236–2243.
- [9] Matthew Cary et al. *On best-response bidding in GSP auctions*. Tech. rep. National Bureau of Economic Research, 2008.
- [10] Deepak Agarwal et al. “Budget pacing for targeted online advertisements at linkedin”. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 2014, pp. 1613–1619.

- [11] Cheng Li et al. “Click-through prediction for advertising in twitter timeline”. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2015, pp. 1959–1968.
- [12] Alvin E Roth and Axel Ockenfels. “Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet”. In: *American economic review* 92.4 (2002), pp. 1093–1103.
- [13] Hal R Varian and Christopher Harris. “The VCG auction in theory and practice”. In: *American Economic Review* 104.5 (2014), pp. 442–45.
- [14] Mehryar Mohri and Andres Munoz Medina. “Learning Theory and Algorithms for revenue optimization in second price auctions with reserve”. In: *Proceedings of the 31st International Conference on Machine Learning*. Ed. by Eric P. Xing and Tony Jebara. Vol. 32. Proceedings of Machine Learning Research 1. Beijing, China: PMLR, June 2014, pp. 262–270. URL: <https://proceedings.mlr.press/v32/mohri14.html>.
- [15] Wei Yang et al. “Optimal reserve price in static and dynamic sponsored search auctions”. In: *Journal of systems science and systems engineering* 22.4 (2013), pp. 440–456.
- [16] Ittai Abraham et al. “Peaches, lemons, and cookies: Designing auction markets with dispersed information”. In: (2013).
- [17] Elizabeth Aguirre et al. “Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness”. In: *Journal of retailing* 91.1 (2015), pp. 34–49.
- [18] Gabriel Y. Weintraub Santiago R. Balseiro Omar Besbes. “Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design.” In: *Management Science* 61.4 (2015), pp. 864–884. URL: <https://doi.org/10.1287/mnsc.2014.2022>.
- [19] Han Cai et al. “Real-time bidding by reinforcement learning in display advertising”. In: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. 2017, pp. 661–670.
- [20] Hana Choi and Carl F Mela. *Display advertising pricing in exchange markets*. Tech. rep. Working paper, Duke University, 2018.
- [21] Jian Xu et al. “Smart pacing for effective online ad campaign optimization”. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2015, pp. 2217–2226.
- [22] Alexandre De Corniere and Romain De Nijs. “Online advertising and privacy”. In: *The RAND Journal of Economics* 47.1 (2016), pp. 48–72.
- [23] Sujith Ravi et al. “Automatic generation of bid phrases for online advertising”. In: *Proceedings of the third ACM international conference on Web search and data mining*. 2010, pp. 341–350.
- [24] Yejin Choi et al. “Using landing pages for sponsored search ad selection”. In: *Proceedings of the 19th international conference on World wide web*. 2010, pp. 251–260.

- [25] Chenyan Xiong et al. “Relational click prediction for sponsored search”. In: *Proceedings of the fifth ACM international conference on Web search and data mining*. 2012, pp. 493–502.
- [26] Nicola Gatti, Alessandro Lazaric, and Francesco Trovò. “A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities”. In: *Proceedings of the 13th ACM Conference on Electronic Commerce*. 2012, pp. 605–622.
- [27] Benjamin Edelman and Michael Schwarz. “Optimal auction design and equilibrium selection in sponsored search auctions”. In: *American Economic Review* 100.2 (2010), pp. 597–602.
- [28] Shuai Yuan et al. “An empirical study of reserve price optimisation in real-time bidding”. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 2014, pp. 1897–1906.
- [29] Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. “Regret Minimization for Reserve Prices in Second-Price Auctions”. In: *IEEE Transactions on Information Theory* 61.1 (2015), pp. 549–564. DOI: 10.1109/TIT.2014.2365772.
- [30] Mehryar Mohri and Andres Muñoz Medina. “Non-parametric Revenue Optimization for Generalized Second Price Auctions”. In: *CoRR* abs/1506.02719 (2015). arXiv: 1506.02719. URL: <http://arxiv.org/abs/1506.02719>.
- [31] Wei Yang, Baichun Xiao, and Lifang Wu. “Learning and pricing models for repeated generalized second-price auction in search advertising”. In: *European Journal of Operational Research* 282.2 (2020), pp. 696–711. ISSN: 0377-2217. DOI: <https://doi.org/10.1016/j.ejor.2019.09.051>. URL: <https://www.sciencedirect.com/science/article/pii/S0377221719308082>.
- [32] Santiago R Balseiro and Yonatan Gur. “Learning in repeated auctions with budgets: Regret minimization and equilibrium”. In: *Management Science* 65.9 (2019), pp. 3952–3968.
- [33] Vincent Conitzer et al. “Multiplicative pacing equilibria in auction markets”. In: *Operations Research* (2021).
- [34] Tsan-Ming Choi, Xun Li, and Cheng Ma. “Search-based advertising auctions with choice-based budget constraint”. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45.8 (2015), pp. 1178–1186.
- [35] Santiago Balseiro et al. “Budget-Management Strategies in Repeated Auctions”. In: *Operations Research* (2021).
- [36] Hana Choi et al. “Online display advertising markets: A literature review and future directions”. In: *Information Systems Research* 31.2 (2020), pp. 556–575.
- [37] Jason Rhuggenaath et al. “Optimal display-ad allocation with guaranteed contracts and supply side platforms”. In: *Computers & Industrial Engineering* 137 (2019), p. 106071.

- [38] Farzaneh Abbaspourtorbati and Marek Zima. “The Swiss reserve market: Stochastic programming in practice”. In: *IEEE Transactions on Power Systems* 31.2 (2015), pp. 1188–1194.
- [39] Patrizia Beraldi, Domenico Conforti, and Antonio Violi. “A two-stage stochastic programming model for electric energy producers”. In: *Computers & Operations Research* 35.10 (2008), pp. 3360–3370.
- [40] Laura Puglia, Daniele Bernardini, and Alberto Bemporad. “A multi-stage stochastic optimization approach to optimal bidding on energy markets”. In: *2011 50th IEEE Conference on Decision and Control and European Control Conference*. IEEE. 2011, pp. 1509–1514.
- [41] Trine Krogh Boomsma, Nina Juul, and Stein-Erik Fleten. “Bidding in sequential electricity markets: The Nordic case”. In: *European Journal of Operational Research* 238.3 (2014), pp. 797–809.
- [42] Yelena Vardanyan, Mikael Amelin, and Mohammad Hesamzadeh. “Short-term hydropower planning with uncertain wind power production”. In: *2013 IEEE Power & Energy Society General Meeting*. IEEE. 2013, pp. 1–5.
- [43] Thomas Volling et al. “A two-stage bid-price control for make-to-order revenue management”. In: *Computers & operations research* 39.5 (2012), pp. 1021–1032.
- [44] Roger B. Myerson. *Game theory - Analysis of Conflict*. Harvard University Press, 1997, pp. I–XIII, 1–568. ISBN: 978-0-674-34116-6.
- [45] Raphaël Jungers et al. *Game Theory: Rational decisions among individuals (lecture notes)*. Université catholique de Louvain, Feb. 2022.
- [46] Noam Nisan et al. *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007.
- [47] Tim Roughgarden. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.
- [48] Francoise Forges and Antoine Salomon. “Bayesian repeated games and reputation”. In: (2014).
- [49] L’udovi’ t Nastišin. “Facebook and Instagram analysis of CPC and CTR analysis by region and industry”. In: (2016).
- [50] W Sabrina Lin, H Vicky Zhao, and KJ Ray Liu. “Incentive cooperation strategies for peer-to-peer live multimedia streaming social networks”. In: *IEEE transactions on multimedia* 11.3 (2009), pp. 396–412.
- [51] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. “Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords”. In: *American economic review* 97.1 (2007), pp. 242–259.
- [52] Baichun Xiao, Wei Yang, and Jun Li. “OPTIMAL RESERVE PRICE FOR THE GENERALIZED SECOND-PRICE AUCTION IN SPONSORED SEARCH ADVERTISING.” In: *Journal of Electronic Commerce Research* 10.3 (2009).
- [53] Daniel R Vincent. “Bidding off the wall: Why reserve prices may be kept secret”. In: *Journal of Economic theory* 65.2 (1995), pp. 575–584.

- [54] Anthony Papavasiliou. *Operational Research (lecture notes)*. Université catholique de Louvain, Apr. 2018.
- [55] Laureano F Escudero et al. “The value of the stochastic solution in multistage problems”. In: *Top* 15.1 (2007), pp. 48–64.
- [56] Zoë Abrams. “Revenue maximization when bidders have budgets”. In: *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. 2006, pp. 1074–1082.
- [57] Zoe Abrams, Ofer Mendelevitch, and John Tomlin. “Optimal delivery of sponsored search advertisements subject to budget constraints”. In: *Proceedings of the 8th ACM conference on Electronic commerce*. 2007, pp. 272–278.
- [58] Tao Qin, Wei Chen, and Tie-Yan Liu. “Sponsored search auctions: Recent advances and future directions”. In: *ACM Transactions on Intelligent Systems and Technology (TIST)* 5.4 (2015), pp. 1–34.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain

Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | www.uclouvain.be/epl