POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Public Signaling in Vickrey-Clarke-Groves Ad Auctions

Tesi di Laurea Magistrale in
Mathematical Engineering - Ingegneria Matematica

Author: **Francesco Bacchiocchi**

Student ID: 952857
Advisor: Prof. Nicola Gatti
Co-advisors: Matteo Castiglioni, Alberto Marchesi, Giulia Romano
Academic Year: 2021-2022

Abstract

We study signaling in Bayesian ad auctions, in which bidders' valuations depend on a random, unknown state of nature. The auction mechanism has complete knowledge of the actual state of nature, and it can send signals to bidders so as to disclose information about the state and increase revenue. For instance, a state may collectively encode some features of the user that are known to the mechanism only, since the latter has access to data sources unaccessible to the bidders. We study the problem of computing how the mechanism should send signals to bidders in order to maximize revenue. While this problem has already been addressed in the easier setting of second-price auctions, to the best of our knowledge, our work is the first to explore ad auctions with more than one slot. In this Thesis, we focus on public signaling and Vickrey-Clarke-Groves mechanisms, under which bidders truthfully report their valuations. We start with a negative result, showing that, in general, the problem does not admit a PTAS unless P = NP, even when bidders' valuations are known to the mechanism. The rest of the Thesis is devoted to settings in which such negative result can be circumvented. First, we prove that, with known valuations, the problem can indeed be solved in polynomial time when either the number of states or the number of slots is fixed. Moreover, in the same setting, we provide an FPTAS for the case in which bidders are single minded, but the number of states of nature and the number of slots can be arbitrary. Then, we switch to the random valuations setting, in which these are randomly drawn according to some probability distribution. In this case, we show that the problem admits an FPTAS, a PTAS, and a QPTAS, when, respectively, the number of states of nature is fixed, the number of slots is fixed, and bidders' valuations are bounded away from zero.

Keywords: Ad Auctions, Bayesian Persuasion, Public Signaling Scheme.

Abstract in lingua italiana

In questa tesi si affronta lo studio di aste pubblicitarie Bayesiane, in cui le valutazioni degli offerenti dipendono da un stato di natura aleatorio. Il meccanismo dell'asta ha una conoscenza completa dello stato di natura reale e può inviare segnali agli offerenti in modo da rivelare parzialmente l'informazione su di esso e aumentare il proprio guadagno. Ad esempio, uno stato di natura può codificare alcune caratteristiche dell'utente che sono note solo al meccanismo, poiché quest'ultimo ha accesso a fonti di dati inaccessibili agli offerenti. In questa tesi si studia come un meccanismo d'asta debba inviare segnali agli offerenti al fine di massimizzare le proprie entrate. Mentre questo problema è già stato affrontato nel caso più semplice di aste *second-price*, al meglio della nostra conoscenza, in questa tesi si affronta per la prima volta il caso di aste pubblicitarie con più di un slot. In particolare, viene studiato il caso di uno schema di segnali pubblico e di un meccanismo Vickrey-Clarke-Groves, in base ai quali gli offerenti riportano in modo veritiero le loro valutazioni. In primo luogo si dimostra che, in generale, il problema non ammette un PTAS a meno che P = NP, anche quando le valutazioni degli offerenti sono note al meccanismo. Per questo motivo, nel resto della tesi, vengono affrontati scenari in cui tale risultato negativo può essere aggirato. In primo luogo, si dimostra che, quando le valutazioni sono note al meccanismo, il problema può essere risolto in tempo polinomiale se il numero degli stati di natura o il numero di slots è fissato. Inoltre, nello stesso contesto, viene mostrata l'esistenza di un FPTAS per il caso in cui gli offerenti sono *single-minded*, ma il numero di stati di natura e slots può essere arbitrariamente grande. Si passa poi al caso di valutazioni aleatorie, in cui queste sono estratte casualmente secondo una certa distribuzione di probabilità. In questo caso, si mostra che il problema ammette un FPTAS, un PTAS e un QPTAS, quando, rispettivamente, il numero di stati è fissato, il numero di slots è fissato e le valutazioni degli offerenti sono tutte maggiori di una certa quantità.

Parole chiave: Aste Pubblicitarie, Persuasione Bayesiana, Schema di Segnali Pubblico.

Contents

List of Figures

List of Tables

1 | Introduction

1.1. General Overview

Nowadays, worldwide spending in digital advertising is skyrocketing, and this growth is primarily driven by ad auctions. These account for almost all market share, since they are at the core of popular advertising platforms, such as, e.g., those by Google, Amazon, and Facebook. According to a recent report by eMarketer [2021], digital ad spending will reach over $490 billion in 2021 and zoom past half a trillion in 2022 (see Figure 1.1).

We study signaling in ad auction settings by means of the Bayesian persuasion framework [18]. In a standard ad auction, the advertisers (also called bidders) compete for displaying their ads on a limited number of slots, and each bidder has their own private valuation representing how much they value a click on their ad. In this work, we study Bayesian ad auctions, which are characterized by the fact that bidders' valuations depend on a random, unknown state of nature. The auction mechanism has complete knowledge of the actual state of nature, and it can send signals to bidders so as to disclose information about the state and increase revenue. In particular, the auction mechanism commits to a signaling scheme, which is defined as a randomized mapping from states of nature to signals being sent to the bidders.

Our model fits many real-world applications that are not captured by classical ad auctions. For instance, a state of nature may collectively encode some features of the user visualizing the ads such as, e.g., age, gender, or geographical region that are known to the mechanism only, since the latter has access to data sources unaccessible to the bidders.

We study the problem of computing a revenue-maximizing signaling scheme for the mechanism. In particular, in this Thesis we focus on public signaling, in which the mechanism can only send a single signal that is observed by all the bidders. Moreover, we restrict our attention to VCG mechanisms, which are widely used in practice and have the appealing property of inducing bidders to truthfully report their valuations. While the signaling problem studied in this Thesis has already been addressed in the easier setting of second-price auctions Badanidiyuru et al. [2018], to the best of our knowledge, our work is the first to explore algorithmic signaling in general ad auctions with more than one slot.

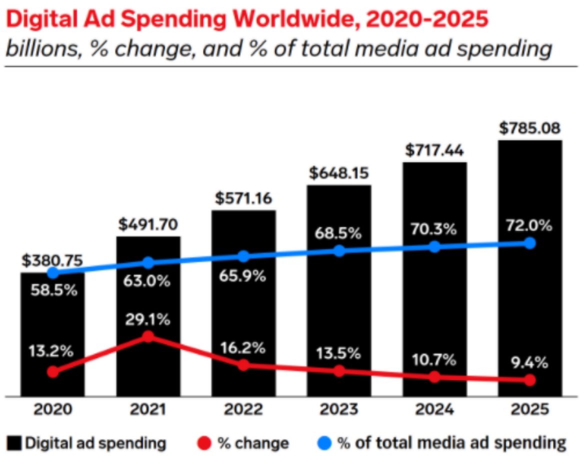

Figure 1.1: Digital ad spending worldwide.

1.2. Related Works

We study signaling in ad auction settings by means of the Bayesian persuasion framework by Kamenica and Gentzkow [2011]. Over the last years, this framework has received considerable attention from the scientific community, due to its applicability to many real-world scenarios, such as online advertising [5, 7, 17], voting [2, 9, 10, 15], traffic routing [6, 13, 21], recommendation systems [19], security [20, 23], and product marketing [3, 8].

To the best of our knowledge, the algorithmic study of signaling in auctions is limited to the second-price auction, which can be seen as a special ad auction with a single slot. A remarkable exception is a recent work by Castiglioni et al. [2022], which studies signaling in posted price auctions. Emek et al. [2014] study second-price auctions in the known valuations setting. They provide an LP to compute an optimal public signaling scheme. Moreover, they show that it is NP-Hard to compute an optimal signaling scheme in the random valuations setting. In our work, we generalize their positive result, in order to provide our polynomial-time algorithm working when the number of slots m is fixed. Cheng et al. [2015] complement the hardness result of Emek et al. [2014] by providing a PTAS for the random valuations setting. This result cannot be extended to ad auctions, as we show in our first negative result. However, we provide two generalizations of the result by Cheng et al. [2015]: we provide a PTAS for the random valuations setting with a fixed number of slots m, and a QPTAS when the bidder's valuations are bounded away from zero. Badanidiyuru et al. [2018] study algorithms whose running time does not depend on the number of states of nature. Moreover, they initiate the study of private signaling schemes, showing that, in second-price auctions, private signaling introduces non-trivial equilibrium selection problems. Finally, Castiglioni et al. [2022] study signaling in posted price auctions. They first prove that the problem of maximizing the seller's revenue does not admit an FPTAS unless P = NP, even for basic instances with a single buyer. Due to that, they focus on designing PTASs working both with public and private signaling schemes.

This work is also related to the line of research that employs q-uniform posteriors distributions in general Bayesian persuasion problems. Cheng et al. [2015] introduce q-uniform posteriors for the first time, showing that, under some assumptions, distributions over q-uniform posteriors approximate a sender-optimal public signaling scheme. They use this argument to provide a PTAS working for second-price auctions with Bayesian valuations. Xu [2020] focuses on the case of receivers' binary actions and employs q-uniform posteriors in order to show that the problem of computing an optimal public signaling scheme admits a bi-criteria PTAS for monotone submodular sender's utility functions. Castiglioni et al. [2020] consider the case of receivers' binary actions and general sender's utility functions, and they provide a tight bi-criteria QPTAS for the problem. All the works mentioned above focus on public signaling. The only exception is Castiglioni et al. [2020], which studies a specific case in between private and public signaling, though restricted to a voting scenario.

1.3. Original Contributions

We start our analysis with a negative result, showing that, in general, the revenue-maximizing problem with public signaling does not admit a PTAS unless P=NP, even when bidders' valuations are known to the mechanism. Thus, in the rest of the Thesis, we address settings in which we can prove that such a negative result can be circumvented. First, we show that, in the known valuations setting, the problem admits a polynomial-time algorithm when either the number of slots m or the number of states d is fixed. The proposed algorithms work by solving suitably-defined linear programs (LPs) of polynomial size, thanks to the crucial property that, when either m or d is fixed, there always exists an optimal signaling scheme using a polynomial number of different signals. Moreover, we also study special instances in which the bidders are single minded, but m and d can be arbitrary. In this case, each bidder positively values a click on their ad only when the actual state of nature is a specific single state, and all the bidders interested in the same state value a click on their ad for the same amount. By exploiting a particular combinatorial structure of the set of bidders' posterior distributions induced by signaling schemes, we are able to provide an FPTAS in such setting. The algorithm works by applying the ellipsoid method in a non-trivial way, with only access to an approximate polynomial-time separation oracle. The latter is implemented by a rather involved dynamic programming algorithm, which works thanks to the particular structure of the set of bidders' posteriors. Then, we switch the attention to the random valuations setting, where bidders' valuations are unknown to the mechanism but randomly drawn according to some probability distribution. In this case, we first provide some preliminary results that establish useful connections between the optimal value of the revenue-maximizing problem and that of optimal signaling schemes restricted to suitably-defined finite sets of posterior distributions. These sets are defined so that the expected revenue of the mechanism is stable, meaning that it does not decrease too much when restricting signaling schemes to use posteriors in such sets. As a preliminary step, we also show that it is possible to compute an approximately-optimal signaling scheme having only access to a finite number of samples from the distribution of bidders' valuations. In conclusion, all the preliminary results described so far allow us to prove that in the random valuations setting, the problem admits an FPTAS, a PTAS, and a QPTAS, when, respectively, d is fixed, m is fixed, and bidders' valuations are bounded away from zero. Finally, the work presented in this Thesis has been recently collected in a scientific paper by Bacchiocchi et al. [2022].

1.4. Thesis Structure

The Thesis is structured as follows.

- In *Chapter 2* we present the theoretical groundings on which this work is based. In particular, we discuss some key concepts about: Ad Auctions, Computational Complexity, and Information Design.
- In *Chapter 3* we introduce the problem of computing an optimal revenue-maximizing signaling scheme in ad auctions through the Bayesian persuasion framework.
- In *Chapter 4* we discuss a general inapproximability result related to the computation of a revenue-maximizing signaling scheme which also holds in the easier case of the known valuation setting.
- In *Chapter 5* we study the problem of computing an optimal revenue-maximizing signaling scheme in a setting in which the mechanism knows the matrix of bidders' valuations.

- In *Chapter 6* we study the problem of computing an optimal revenue-maximizing signaling scheme when the bidders' valuations are drawn from a probability distribution unknown to the mechanism.
- In *Chapter 7* we present some experimental results aimed at understanding the scenarios in which committing to a signaling scheme gives the mechanism higher expected revenue than not doing so.

2 | Preliminaries

2.1. Vickrey-Clarke-Groves Mechanism

This section presents some preliminary concepts helpful in understanding the Vickrey-Clarke-Groves (VCG) mechanism and its main properties. First, we introduce the following notation:

- $\mathcal{N} := [n]$ is the set of players.[1]
- $A := A_1 \times A_2 \times ... \times A_n$ is a set of possible joint actions profiles, where A_i is the action set of player $i \in \mathcal{N}$.
- $\Theta := \Theta_1 \times ... \times \Theta_n$ is a set of possible joint types profiles, where Θ_i is the set of types of player $i \in \mathcal{N}$.
- X is the set of outcomes.
- $g : A \to X$ is the outcome function.
- λ is a probability distribution over the set of types Θ.
- $u := (u_1, ..., u_n)$ is the tuple of utility functions, where $u_i : X \times \Theta \to [0, 1]$ is the utility of player $i \in \mathcal{N}$.

We also introduce the following definitions.

Definition 2.1 (Economic mechanism). *An Economic mechanism is a tuple* (A, X, g).

Definition 2.2 (Bayesian Game). *A Bayesian Game is a tuple* $(A, X, g, \Theta, \lambda, u)$.

Definition 2.3 (Social Choice Function). *A social choice function* $f : \Theta \to X$ *is a function that associates each possible tuple of players' types to an outcome.*

Basically, given the preferences of the receivers, a social choice function investigates ways in which they can be aggregated, while respecting some properties.

[1] In the following, given a natural number $n \in \mathbb{N}_+$, we denote with $[n]$ the set $\{1, ..., n\}$.

Definition 2.4 (Direct (revelation) economic mechanism). *Given a social choice function $f : \Theta \rightarrow X$, a direct (revelation) economic mechanism is a mechanism (Θ, X, f).*

We also introduce the definition of quasi linear environments in which the utility of a player linearly depends on her valuation.

Definition 2.5 (Quasi Linear Environment). *A quasi linear environment is given by:*

- *An outcome space given by: $X = \{(k, p_1, ..., p_n) : k \in K,\ p_i \in [0,1]\ \forall i \in \mathcal{N}\}$.*
- *Utility functions: $u_i(k, \theta_i) = v_i(k, \theta_i) - p_i$ with $v_i : K \times \Theta_i \rightarrow [0,1]$.*

Where K is the set of possible allocations.

Quasi linear environments play a crucial role since they describe the key features of many auctioning scenarios. Indeed, an allocation $k \in K$ indicates which items are assigned to each player, while payments represent the money each player has to provide to the auctioneer. We now introduce an important characterization of a social choice function in which players are incentivized to report their true type.

Definition 2.6 (Dominant-Strategy Incentive Compatibility). *A social choice function $f : \Theta \rightarrow X$ is dominant strategy incentive compatible (DSIC) if the Bayesian game induced by the direct revelation economic mechanism has a pure dominant strategy equilibrium $(s_1^*, ..., s_n^*)$ such that $s_i^*(\theta_i) = \theta_i$ for every player $i \in \mathcal{N}$ and type $\theta_i \in \Theta_i$.*

The definition means that truthfully responding is a dominant strategy for the players, no matter what other players do. We finally present a mechanism having the crucial property of incentivizing players to report their true type.

Definition 2.7 (Vickrey–Clarke–Groves (VCG) mechanism). *A direct-revelation economic mechanism (Θ, X, f) in which $f(\theta) = (k(\theta), p_1(\theta), ..., p_n(\theta))$ is Vickrey–Clarke–Groves (VCG) mechanism if :*

- *The allocation is given by: $k^* \in \arg\max_{k' \in K} \sum_{i \in \mathcal{N}} v_i(k', \theta_i)$.*
- *The payments are given by: $p_i(k^*, \theta_i) = \max_{k' \in K_{-i}} \sum_{i \neq j} v_j(k', \theta_j) - \sum_{i \neq j} v_j(k^*, \theta_j)\ \forall i \in \mathcal{N}$.*

Where K_{-i} is the set of possible allocations without player $i \in \mathcal{N}$.

Theorem 2.1. *Every VCG Mechanism is Dominant Strategy Incentive Compatible (DSIC).*

Proof. Let $k_1 \in K$ be the allocation that maximizes the players' valuations when player $i \in \mathcal{N}$ reports her true type $\theta_i \in \Theta_i$, while let $k_2 \in K$ be the allocation that maximizes the palyers' valuations when player $i \in \mathcal{N}$ misreports her true type $\tilde{\theta}_i \in \Theta_i$. Formally, we have that:

$$k_1 \in \arg\max_{k \in K} v_i(k, \theta_i) + \sum_{i \neq j} v_j(k, \theta_j) \qquad k_2 \in \arg\max_{k \in K} v_i(k, \tilde{\theta}_i) + \sum_{i \neq j} v_j(k, \theta_j).$$

The utility of player $i \in \mathcal{N}$ when she reports her true type is given by:[2]

$$\begin{aligned} u_i(k_1, \theta_i) = v_i(k_1, \theta_i) - p_i(k_1) &= v_i(k_1, \theta_i) + \sum_{i \neq j} v_j(k_1, \theta_j) + \max_{k' \in \mathrm{K}_{-i}} \sum_{i \neq j} v_j(k', \theta_j) \\ &\geq v_i(k_2, \tilde{\theta}_i) + \sum_{i \neq j} v_j(k_2, \theta_j) + \max_{k' \in \mathrm{K}_{-i}} \sum_{i \neq j} v_j(k', \theta_j) \\ &= u_i(k_2, \tilde{\theta}_i). \end{aligned}$$

The inequality holds because of the definition of allocations $k_1, k_2 \in K$ and shows that the utility of player $i \in \mathcal{N}$ is maximized when she reports her true type, concluding the proof. □

We now present an introductory example of an auction implementing a VCG mechanism.

Example 2.1. *We consider an auction with three bidders and two items. This example assumes that the bidders' valuations coincide with their types, as reported in the following.*

	v_1	v_2	v_3
$\emptyset$	0.0	0.0	0.0
$\{A\}$	0.2	0.0	0.0
$\{B\}$	0.0	0.5	0.0
$\{A, B\}$	0.2	0.5	0.6

Table 2.1: Bidders' valuations.

The VCG allocation prescribes item $\{A\}$ to player 1 while it prescribes item $\{B\}$ to player 2, since this allocation is the one that maximizes the reported valuations. Finally, the

[2] W.l.o.g., we assume that each player $j \in \mathcal{N}$ reports her true type.

prices prescribed by the VCG mechanism are reported below:

$$p_1 = \max_{k' \in K_{-1}} \sum_{i \neq 1} v_i(k') - \sum_{j \neq 1} v_j(k^*) = 0.6 - 0.2 = 0.4$$
$$p_2 = \max_{k' \in K_{-2}} \sum_{i \neq 2} v_i(k') - \sum_{j \neq 2} v_j(k^*) = 0.6 - 0.5 = 0.1$$
$$p_3 = \max_{k' \in K_{-3}} \sum_{i \neq 3} v_i(k') - \sum_{j \neq 3} v_j(k^*) = 0.7 - 0.7 = 0$$

where $k^* \in K$ is the allocation presented in Definition 2.7, which maximizes the reported valuations.

2.2. Ad Auctions

In this section, we introduce and characterize ad auctions (or sponsored search auctions). In particular, as customary in the literature, we assume there is a set $\mathcal{N} := [n]$ of advertisers (or bidders) who compete for displaying their ads on a set $\mathcal{M} := [m]$ of slots, with $m \leq n$. Each bidder $i \in \mathcal{N}$ is characterized by a private valuation $v_i \in [0, 1]$, which represents how much they value a click on their ad. Moreover, each slot $j \in \mathcal{M}$ is associated with a click through rate (CTR) parameter $\lambda_j \in [0, 1]$, which is the probability of a user clicking the slot and we indicate with $\Lambda := \{\lambda_1, ..., \lambda_m\}$ the set of these parameters. In this work, for the ease of presentation, we assume that the click through rate only depends on the slot and not on the ad being displayed. In general, each slot may have its own prominence value the probability with which a user observes it and each bidder may have their own ad quality the probability with which their ad is clicked once observed, so that the click through rate is defined as the product of these two quantities. All the results in this thesis can be easily extended to such general model. Moreover, for ease of notation, we assume that the slots are ordered so that $\lambda_1 \geq \ldots \geq \lambda_m$. The auction goes on as follows: first, each bidder $i \in \mathcal{N}$ separately reports a bid $b_i \in [0, 1]$ to the auction mechanism; then, based on the bids, the latter allocates an ad to each slot and defines how much each bidder has to pay the mechanism for a click on their ad. We focus on truthful mechanisms, and the VCG mechanism in particular. In truthful mechanisms, as shown in previous section, allocation and payments are defined so that it is a dominant strategy for each bidder to report their true valuation to the mechanism, namely $b_i = v_i$ for every $i \in \mathcal{N}$. In particular, the allocation implemented by the VCG mechanism orderly assigns the first m bidders in decreasing value of b_i to the first m slots (those with the highest click through rates). At the same time, assuming without loss of generality that bidder i

is assigned to slot i, the mechanism defines the payments according to Definition 2.7:

$$\begin{aligned} p_i &= \Big(\sum_{j=i+1}^{m+1} b_j \lambda_{j-1} + \sum_{j=1}^{i-1} b_j \lambda_j \Big) - \Big(\sum_{j=1}^{m+1} b_j \lambda_j - \lambda_i b_i \Big) \\ &= \sum_{j=i+1}^{m+1} b_j (\lambda_{j-1} - \lambda_j), \end{aligned}$$

for each bidder $i \in [m]$, where, for the ease of notation, we let $\lambda_{m+1} = 0$. The payment is zero for all the other bidders. In practice, each bidder $i \in [m]$ has to pay p_i/λ_i whenever a user clicks on their ad, so that their utility is $\lambda_i v_i - p_i$ in expectation over the clicks. The expected utility of all the other bidders is zero. In the following we provide an example of ad auction implementing a VCG allocation mechanism.

Example 2.2. *We consider an ad auction with $m = 3$ slots and $n = 6$ bidders interested in displaying their ad. In the following tables we summarize all the information concerning the bidders' valuations and the slots' features. In particular the CTRs are given by:*

λ_1	λ_2	λ_3
0.8	0.3	0.2

Table 2.2: Slots' CTRs.

while the bidders' valuations are given by:

v_1	v_2	v_3	v_4	v_5	v_6
0.5	0.6	0.9	0.4	0.4	0.8

Table 2.3: Bidders' valuation in an ad auction.

It is easy to check that, accordingly to what prescribed by the VCG allocation mechanism, the first slot is allocated to advertiser 3, the second to advertiser 6, while the last one to advertiser 2, as shown in Figure 2.1. In addition, the payments prescribed by the mechanism are reported below.

$$\begin{aligned} p_3 &= v_6(\lambda_1 - \lambda_2) + v_2(\lambda_2 - \lambda_3) + v_1 \lambda_3 = 0.56 \\ p_6 &= v_2(\lambda_2 - \lambda_3) + v_1 \lambda_3 = 0.16 \\ p_2 &= v_1 \lambda_3 = 0.1. \end{aligned}$$

We remark that the expected payments associated to the three remaining receivers are equal to zero since their ad is not displayed.

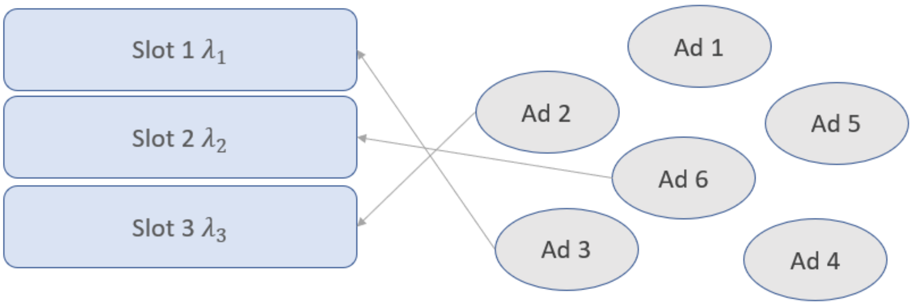

Figure 2.1: VCG allocation in ad auctions.

2.3. Computational Complexity

In this section, we present some definitions related to the Computational Complexity Theory, which is the field of Computer Science that aims at classifying problems in complexity classes according to their difficulty. In particular, we introduce some preliminary concepts and the most important complexity classes. In general, the complexity of an algorithm can be measured in two dimensions:

- the time complexity, which measures the maximum number of computational steps that the algorithm performs in order to output a solution;
- the space complexity, which measures the amount of memory used by the algorithm.

Usually, the complexity is expressed as a function of the size of the algorithm's input. However, since it is often quite involved to provide an analytical expression of such a function, it results easier just to provide an upper bound of it.

Definition 2.8. *The function $f(n)$ is asymptotically upper-bounded by $g(n)$, we write it as $f(n) = O(g(n))$, if there exist two positive constants c and n_0 such that:*

$$0 \leq f(n) \leq cg(n), \quad \forall n \geq n_0.$$

Thanks to the previous definitions we can now distinguish the two main classes of algorithms.

Definition 2.9. *An algorithm is polynomial if it has a worst-case time complexity $O\left(n^d\right)$, where d is a constant and n is the size of the input.*

Definition 2.10. *An algorithm is exponential if it has a worst-case time complexity $O\left(2^n\right)$ where n is the size of the input.*

It is possible to show that there exist problems that are particularly difficult and can be solved only in exponential time. Due to that, problems are classified in complexity classes, depending on their difficulty. We present now complexity classes with respect to two types of problems: decision problems and functional problems. Decision problems are problems that admit only YES or NO answers and can be split into three complexity classes:

- **P**. The class of decision problems that can be solved by a polynomial time algorithm.
- **NP**. The class of decision problems such that for every YES-instance there is a concise certificate which allows to verify, in polynomial time, that the instance really admits a YES answer.
- **Co-NP**. The class of decision problems whose complements are in NP; where the complement of a decision problem is obtained reversing the YES and NO answers.

Functional problems are problems where a single output, if it exists, is expected for every input, but the outcome is more complex than one of decision problems. Also these problems can be split into complexity classes, including:

- **FP**. The class of functional problems that can be solved by a polynomial time algorithm;
- **FNP**. The class of functional problems such that there exists an algorithm that, given a problem instance I and a solution y, can verify, in polynomial time, whether y is a solution of I.

In NP, there is a class of problems, called NP-Complete problems, which includes the most difficult problems in NP. Each problem in NP can be transformed into one of these problems in polynomial time. In the following, we introduce the formal definition of reduction, NP-Hardness, and NP-Completeness.

Definition 2.11. *A problem π can be reduced in polynomial time to another problem π', written as $\pi \leq_P \pi'$, if there exists a transformation that allows to build an instance I' of π' from an instance I of π in polynomial time, and such that a solution of I can be derived, in polynomial time, from one of I'.*

If a problem π is reducible in polynomial time to a problem π', then π' is at least as difficult as π. For example, if π' can be solved in polynomial time also π has a polynomial time algorithm. In fact, we can solve π by reducing π to π' and solving π'. Reductions are also useful to determine membership to another class of problems, called NP-Hard problems.

Definition 2.12. *A problem π is* NP*-Hard if every problem $\pi' \in \mathrm{NP}$ is such that $\pi' \leq_P \pi$.*

Definition 2.13. *A problem π is* NP*-Complete if it is in* NP *and it is* NP*-Hard.*

NP-Complete problems play a crucial role because, if there exists a polynomial algorithm for a NP-Complete, then P = NP, while proving that does not exist a polynomial algorithm for a NP-Complete implies that P $\subset$ NP and thus some problem in NP cannot be solved in polynomial time. The fact that no polynomial time algorithm for an NP-Complete problem has ever been found motivates the conjecture that P $\neq$ NP. Figure 2.2 shows the two different scenarios. Since it is possible to reduce every problem in NP to an NP-Complete problem, to show that a problem is NP-Hard it is enough to prove that an NP-Complete problem can be reduced to it. We now restrict our attention to optimization problems since they play a crucial role in the development of this thesis. Sometimes, optimization problems cannot be solved efficiently by an exact algorithm, but there exist algorithms that can approximate an optimal solution in polynomial time.

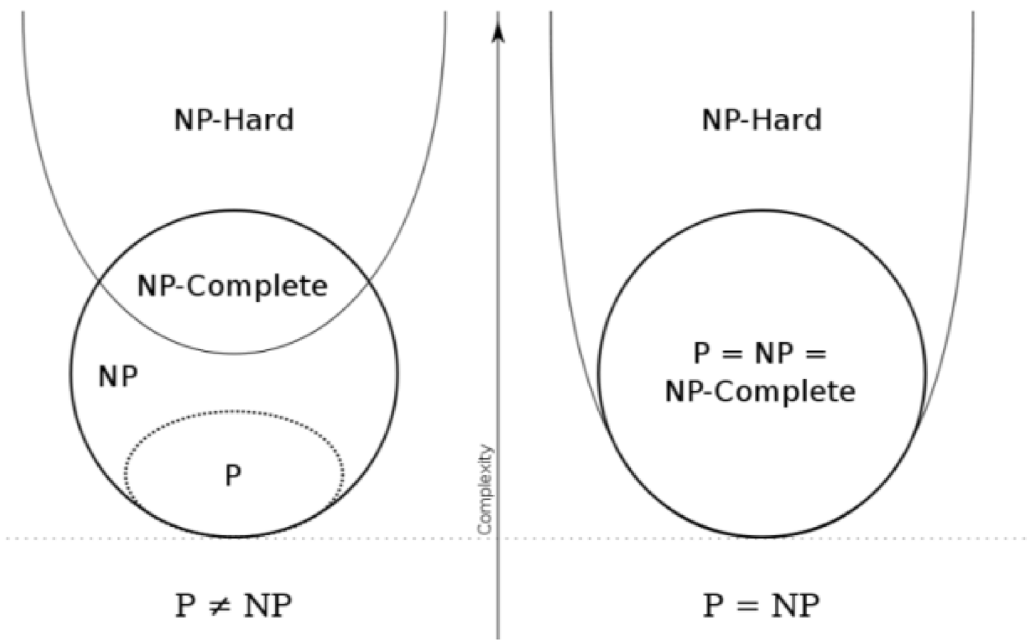

Figure 2.2: The two possible relations between P and NP.

Since an approximate solution could be acceptable for a given application, approximation algorithms can be usefully employed, as they are usually much faster than exact algorithms and guarantee a bound on the error. Since we will mainly rely on maximization problems in this work, we present the following definitions in order to maximize a given objective function. Nevertheless, it is possible to extend these definitions for other kinds of problems. In the following we let OPT be the value of the optimal solution while we let APX be the value returned by the approximation algorithm.

Definition 2.14 (Multiplicative approximation algorithm)**.** *Given a maximization problem $\mathcal{P}$ and $\varepsilon > 0$. An algorithm provides an ε-multiplicative approximation for the problem $\mathcal{P}$, if for all instances I of $\mathcal{P}$ it delivers a feasible solution such that:*

$$\frac{\mathrm{APX}}{\mathrm{OPT}} \geq (1-\varepsilon)$$

Definition 2.15 (Additive approximation algorithm)**.** *Given a maximization problem $\mathcal{P}$ and $\varepsilon > 0$. An algorithm provides an ε-additive approximation for the problem $\mathcal{P}$, if for all instances I of $\mathcal{P}$ it delivers a feasible solution such that:*

$$\mathrm{APX} \geq \mathrm{OPT} - \varepsilon$$

We also introduce the definition of polynomial-time approximation scheme (PTAS) where the algorithm gives arbitrarily good approximations for the problem in polynomial time.

Definition 2.16 (Polynomial time approximation scheme)**.** *Let $\mathcal{P}$ be a maximization problem. A polynomial time approximation scheme ($PTAS$) is an algorithm that takes an instance I of $\mathcal{P}$ of size n and a parameter $\varepsilon > 0$, and in time poly(n), outputs a solution that is a (multiplicative or additive) ε-approximation of the optimal solution.*

Similarly, we introduce the strongest notion of fully polynomial time approximation scheme (FPTAS).

Definition 2.17 (Fully polynomial-time approximation scheme)**.** *Let $\mathcal{P}$ be a maximization problem. A fully polynomial time approximation scheme (FPTAS) is an algorithm that takes an instance I of $\mathcal{P}$ of size n and a parameter $\varepsilon > 0$, and in time poly $(n, 1/\epsilon)$, outputs a solution that is a (multiplicative or additive) ε-approximation of the optimal solution.*

2.4. Information Design

Information design is the field in economics and game theory that studies how the strategic provision of information can affect the interaction between rational players. In this scenario, the interaction between agents can be described as a game with incomplete information. The main important feature of these games is that the payoff function depends on a finite set of parameters $\Theta \coloneqq \{\theta_1, \dots, \theta_d\}$, denoted as states of nature, in order to model the environment randomicity. The primary goal of Information Design is to study and characterize the information structure of such games, which can be modeled as a map between states of nature and a countable set of signals $\mathcal{S}$. Signals are designed to spread information to the agents and, by doing so, influence the decisions they will consequently take. Information design addresses the study of information structures by using two approaches based on a descriptive or a prescriptive question. The descriptive question aims at understanding the consequences of the provision of information on the set of potential equilibria of a game with incomplete information. While the prescriptive question, also called persuasion, deals with computing the information structure which optimizes a given objective function. Bayesian persuasion tries to understand all those situations in which an informed principal (the sender) tries to persuade the behavior of self-interested agent(s) (the receiver(s)) via the provision of payoff-relevant information.

2.4.1. Persuading a Single Agent

This model describes the interaction of two agents, an informed sender and a receiver. The receiver may pick one of several actions $a \in \mathcal{A}$ and her utility is specified by the function $u : \mathcal{A} \times \Theta \to \mathbb{R}$, highlighting the fact that she is influenced by the selected action and by a random state of nature $\theta \in \Theta$. In turn, the choice of the receiver effects together with the state of nature $\theta \in \Theta$ the sender's objective function $f : \mathcal{A} \times \Theta \to \mathbb{R}$. The state of nature $\theta \in \Theta$ is drawn according to a prior distribution $\mu \in \Delta_\Theta$, which is known by the sender and the receiver.[3] However, since the sender can observe the actual realization of the state of nature $\theta \in \Theta$, she can persuade the receiver to take a more profitable action. In addition, before the state of nature is realized, the sender publicly commits to a signaling scheme (also called policy), which represents a randomized map from the set of states of nature Θ to the set of signals $\mathcal{S}$. A signal can be considered as an abstract element without intrinsic meaning. Nevertheless, since the receiver is rational and Bayesian, she interprets it depending on how the signaling scheme uses it. Finally,

[3] Given a finite set X, we denote with Δ_X the $(|X|-1)$-dimensional simplex defined over the elements of X.

the interaction between the sender and the receiver is reported below:

1. The sender commits to signaling scheme $\phi : \Theta \to \Delta_{\mathcal{S}}$.
2. The receiver observes the signaling scheme chosen by the sender.
3. The true state $\theta \sim \mu$ is realized and the sender observes it.
4. A signal $s \in \mathcal{S}$ is drawn from ϕ_θ.
5. The receiver observes the signal s and rationally updates his belief over Θ, according to the Bayes rule.
6. The receiver selects an action a maximizing her expected utility.
7. The sender and the receiver receive utility $f(\theta, a)$ and $u(\theta, a)$, respectively.

A signaling scheme $\phi : \Theta \to \Delta_{\mathcal{S}}$ is a map from the set of states of nature Θ to distributions over the set of signals $\mathcal{S}$. Given a public signaling scheme, the probability that the sender selects a signal $s \in \mathcal{S}$ after observing $\theta \in \Theta$ is given by $\phi_\theta(s)$. On the other hand, the probability to receive the signal $s \in \mathcal{S}$ is equal to $\Pr(s) = \sum_{\theta \in \Theta} \mu_\theta \phi_\theta(s)$. After receiving the signal $s \in \mathcal{S}$, the receiver infers a probability distribution $\xi(\theta) = \mu_\theta \phi_\theta(s) / \Pr(s)$ over the set of states of nature by means of Bayes theorem. Then, the receiver takes an action $a^*(\xi) \in \arg\max_{a \in \mathcal{A}} \sum_{\theta \in \Theta} \xi(\theta) u(\theta, a)$, maximizing her expected utility. In case of indifference, it is usually assumed that the tie is broken in favor of the sender. Finally, the sender gets utility $f(\theta, a^*(\xi))$. Algorithmic Bayesian persuasion studies the algorithmic problem of computing an optimal signaling scheme that maximizes the expected sender's utility. In particular, since each signal corresponds to an action, it is possible to label each one of them with the action it generates. From this observation, Kamenica and Gentzkow [2011] show by a revelation-principle style argument that it is possible to assume that the sender commits to signaling schemes with two different properties. Direct signaling schemes, meaning that the signals $\mathcal{S}$ are action recommendations, and persuasive signaling schemes, which means that the receiver has no incentive in deviating from the recommended action. As a result, $|\mathcal{S}| = |\mathcal{A}|$, with each signal $s \in \mathcal{S}$ interpretable as a recommendation to play action $a \in \mathcal{A}$. Moreover, the recommendation of the sender is followed by the receiver since no other action provides her a better utility. In addition, the presence of incentive compatibility (IC) constraints in the formulation of the problem ensures the persuasiveness requirement. Finally, restricting our attention to direct and persuasive signaling schemes, the following linear program (LP) describes the sender's optimization problem.

$$\begin{aligned}
\max_{\phi\geq 0} \sum_{\theta\in\Theta} \mu_\theta \sum_{a\in\mathcal{A}} \phi_\theta(a) f(\theta,a) \quad & \text{s.t.} \\
\sum_{\theta\in\Theta} \mu_\theta \phi_\theta(a) \left(u\left(\theta,a\right) - u\left(\theta,a'\right)\right) \geq 0 \quad & \forall a,a' \in \mathcal{A} \\
\sum_{a\in\mathcal{A}} \phi_\theta(a) = 1 \quad & \forall \theta \in \Theta \\
\phi_\theta(a) \geq 0 \quad & \forall a \in \mathcal{A}, \forall \theta \in \Theta.
\end{aligned}$$

Notice that in the previous LP there is a variable $\phi_\theta(a) \in [0,1]$ for each state of nature $\theta \in \Theta$ and action $a \in \mathcal{A}$. Solving this LP is often impractical, due to that, it is possible to re-formulate it in the space of the receiver's posterior probabilities. As we will see along this thesis, this approach will allow us to tackle many different problems and it will play a crucial role in the development of all our analysis. In particular, as previously observed, given a signaling scheme $\phi_\theta(s)$, each signal $s \in \mathcal{S}$ is played with probability $\Pr(s) = \sum_{\theta\in\Theta} \mu_\theta \phi_\theta(s)$ and once the receiver observes it, she infers a posterior distribution over the set of states of nature $\xi(\theta) = \mu(\theta)\phi_\theta(s)/\Pr(s)$. Therefore, a signaling scheme can be thought as a distribution $\gamma \in \Delta\left(\Delta_\Theta\right)$ over posterior distributions, one for each signal, whose expectation equals the prior $\mu \in \Delta_\Theta$. In other words, a signaling scheme is a way of writing $\mu \in \Delta_\Theta$ as a convex combination of the posterior distributions in the simplex.

$$\begin{aligned}
\max_{\gamma\geq 0} \sum_{\xi\in\mathrm{supp}(\gamma)} \gamma(\xi) \sum_{\theta\in\Theta} \xi(\theta) f\left(\theta, a^*(\xi)\right) \quad & \text{s.t.} \\
\sum_{\xi\in\mathrm{supp}(\gamma)} \gamma(\xi)\xi(\theta) = \mu_\theta \quad & \forall \theta \in \Theta.
\end{aligned}$$

In principle, since the space of possible posterior distributions over the states of natures is infinite-dimensional, also the latter LP results impractical to be solved. Nevertheless, in this work we will see how to solve the previous LP for particular instances of the problem we consider. The following example presents a possible scenario of application of the Bayesian persuasion model.

Example 2.3. *Consider an academic advisor (the sender) who is writing a recommendation letter (the signal) for his graduating student to send to a company (the receiver), which in turn must decide whether or not to hire the student. The advisor gets utility 1 if his student is hired, and 0 otherwise. The state of nature determines the quality of the student, and hence the company's utility from hiring the student. Suppose that the student is excellent with probability* 1/3*, and weak with probability* 2/3*. Moreover, suppose that* *the company gets utility 1 for hiring an excellent student, utility* -1 *for hiring a weak student, and utility 0 for not hiring. Consider the following signaling schemes:*

1. ***No information:*** *Given no additional information, the company maximizes its utility by not hiring. The advisor's expected utility is* 0.

2. ***Full information:*** *Knowing the quality of the student, the company hires him if and only if he is excellent. The advisor's expected utility is* $1/3$.

3. ***The optimal (partially informative):*** *The advisor recommends "hiring" when the student is excellent, and with probability* $1/2$ *when the student is weak. Otherwise, the advisor recommends "not hiring". The company maximizes its expected utility by following the recommendation, and the advisor's expected utility is* $2/3$.

2.4.2. Persuading Multiple Agents

On the ground of the foundational work by Kamenica and Gentzkow [2011], several extensions of the single-agent model have been proposed. One of the most important is the one that allows the presence of multiple receivers. Information structure design can get much more complex when a lot of receivers are involved. This is due to the following two key dimensions of challenges:

- **Externalities among receivers:** Whereas in the single-agent model the receiver selects the action exclusively on the ground of its virtue to maximise his own utility, by contrast the presence of multiple receivers introduces the additional problem of selecting and computing an equilibrium among them.
- **Coordinating receivers:** When the sender faces a multitude of receivers, she faces the additional problem of coordinating their actions.

Seeking to split these two different sources of complexity, it is common to introduce an additional hypothesis to the Bayesian persuasion model that is the so called no-externalities assumption, prescribing that each receiver's utility is independent from the actions chosen by other receivers. This assumption allows to circumvent the equilibrium selection and computational problem, which have been proven to make the problem intractable even in elementary settings, and focus on the critical problem of coordinating receivers' behaviors. However, even when accounting for the no-externalities assumption, the issue of computing an optimal signaling scheme remains largely intractable in a multi-receivers setting. To account for multiple receivers with no inter-agent externalities, the basic model presented before requires a series of changes. In particular, the single receiver is substituted with a finite set R of receivers, each one equipped with a finite set of available actions $\mathcal{A}^r$. Due to the no-externalities assumption, each receiver's payoff depends only on the action she selects and the state of nature $\theta \in \Theta$. The function $u_r : \Theta \times \mathcal{A}^r \to \mathbb{R}$ describes the utility of receiver $r \in R$, and the notation $u_r(\theta, a)$ denotes the utility she perceives by playing action $a \in \mathcal{A}^r$ being in state $\theta \in \Theta$. The sender's utility is a function of the state of nature $\theta \in \Theta$ and the actions selected by all the receivers. A tuple specifying an action for each receiver is usually indicated as action profile and it is denoted by $\mathbf{a} \in \mathcal{A}$, where $\mathcal{A} := \times_{r \in R} \mathcal{A}^r$. Hence, the sender's utility is described by $f : \Theta \times \mathcal{A} \to \mathbb{R}$ and $f(\theta, \mathbf{a})$ denotes the sender's payoff when receivers behave according to action profile $\mathbf{a} \in \mathcal{A}$ and the state of nature is $\theta \in \Theta$. The interaction between the sender and the receivers proceeds with the same steps that characterize the exchange of information in the single-agent persuasion model. The slight difference is that the sender must send signals to multiple receivers and, therefore, she sends a signal to each of them. How these signals are structured depends on the type of signaling scheme employed by the sender. In the literature of multi-receivers persuasion, two basic types of signaling scheme have been studied:

- **Private signaling scheme**: The sender can reveal different (possibly correlated) signals to different receivers through a private communication channel.
- **Public signaling scheme**: The sender is constrained to rely upon a public communication channel, thus she has to send a unique signal to every receiver, making the same information available to all of them.

Example 2.4. *Consider again an academic advisor (the sender) who is writing a recommendation letter (the signal) for his student. However, now the student has applied to two fellowship programs (the receivers), each of which must decide whether or not to award the student a fellowship funding as part of his graduate education. Suppose that the student can accept one or both fellowship awards. The advisor gets utility 1 if his student is awarded at least one fellowship, and o otherwise. A student is excellent with probability* $1/3$ *and weak with probability* $2/3$*, and a fellowship program gets utility 1 from awarding an excellent student,* -1 *from awarding a weak student, and 0 from not awarding the student. Obviously, a fellowship program makes an award if and only it believes its expected utility for doing so is non negative. Consider the following signaling schemes:*

- ***No Information:*** *Neither program makes the award, and the advisor's utility is 0.*
- ***Full information:*** *Both programs make the award if the student is excellent, and neither of them makes the award if the student is weak. The advisor's expected utility is* $1/3$.

- ***Optimal public scheme:*** *If the student is excellent, the advisor publicly signals "award". If the student is weak, the advisor publicly signals "award" or "don't award" with equal probability. Therefore, both programs are simultaneously persuaded to award the fellowship the student with probability* $2/3$*, and neither of them makes the award with probability 1/3. The advisor's expected utility is* $2/3$*.*

- ***Optimal private scheme:*** *If the student is excellent, the advisor recommends "award" to both fellowship programs. If the student is weak, the advisor recommends"award" to one fellowship program chosen uniformly at random, and recommends "don't award" to the other. The result is that both fellowship programs make the award when the student is excellent, and exactly one of the programs makes the award when the student is weak. This yields utility 1 for the advisor.*

2.4.3. Public Signaling Schemes

In this thesis we will focus only on public signaling schemes in which the sender reports the same identical signal $s \in \mathcal{S}$ to all the receivers and we indicate with $\mathcal{S}$ the set of such signals. Analogously to what we have presented before, a public signaling scheme $\phi : \Theta \to \Delta_{\mathcal{S}}$ is a map from states of nature Θ to distributions over public signals $\mathcal{S}$. Moreover, given a public signaling scheme, the probability with which the sender selects a signal $s \in \mathcal{S}$ after observing $\theta \in \Theta$ is denoted by $\phi_\theta(s)$. While, the probability with which the receiver receives signal s is given by $\Pr(s) = \sum_{\theta \in \Theta} \mu_\theta \phi_\theta(s)$. Since the prior is a common information among the receivers and also they observe the same signal $s \in \mathcal{S}$, they all perform the same Bayesian update and infer the same posterior belief $\xi(\theta) = \mu_\theta \phi_\theta(s) / \Pr(s)$. After inferring the posterior distribution $\xi \in \Delta_\Theta$, each receiver r solves a single-receiver decision problem to choose the action that maximizes his posterior expected utility, in a way that: $a^*(\xi) \in \arg\max_{a \in \mathcal{A}^r} \sum_{\theta \in \Theta} \xi(\theta) u_r(\theta, a)$. Analogously to the single-receiver case, ties are broken in favor of the sender. Again, applying the same argument employed in single-receiver persuasion, we can restrict attention to direct and persuasive public signaling schemes. Therefore, each signal $s \in \mathcal{S}$ can be equivalently expressed as an action profile $\mathbf{a} \in \mathcal{A}$, recommending an action to each receiver. As in the single-receiver persuasion case, the persuasiveness requirement is enforced by IC-constraints on the signaling scheme. Also in this case we focus our attention on direct and persuasive public signaling schemes and so the sender's optimization problem can be formulated with the following LP, where we denote with $\phi_\theta(\mathbf{a})$ the probability with which the sender draws the signal $s = \mathbf{a}$ when the realized state of nature is $\theta \in \Theta$. Notice that the following LP has an exponential number of variables and so it cannot be solved in polynomial time.

$$
\begin{aligned}
\max_{\phi \geq 0} \sum_{\theta \in \Theta} \mu_\theta \sum_{\mathbf{a} \in \mathcal{A}} \phi_\theta(\mathbf{a}) f(\theta, \mathbf{a}) \quad & \text{s.t.} \\
\sum_{\theta \in \Theta} \mu_\theta \phi_\theta(\mathbf{a}) \left(u_r\left(\theta, \mathbf{a}_r\right) - u_r\left(\theta, a'\right)\right) \geq 0 \quad & \forall r \in R, \forall \mathbf{a} \in \mathcal{A}, a' \in \mathcal{A}^r \\
\sum_{\mathbf{a} \in \mathcal{A}} \phi_\theta(\mathbf{a}) = 1 \quad & \forall \theta \in \Theta \\
\phi_\theta(\mathbf{a}) \geq 0 \quad & \forall \mathbf{a} \in \mathcal{A}, \forall \theta \in \Theta.
\end{aligned}
$$

Finally, as highlighted for the case of single-receiver persuasion, it is possible to re write the problem of computing an optimal public signaling scheme adopting a different approach. Indeed, adopting the same reasoning employed in previous section, a public signaling scheme can be thought as a distribution $\gamma \Delta \left(\Delta_\Theta\right)$ over posterior distributions such that $\sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi) \, \xi(\theta) = \mu_\theta$ for each $\theta \in \Theta$. Hence, the problem can be seen has the computation of a distribution $\gamma \in \Delta \left(\Delta_\Theta\right)$ that optimizes the sender's utility.

$$
\begin{aligned}
\max_{\gamma \geq 0} \sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi) \sum_{\theta \in \Theta} \xi(\theta) f\left(\theta, \mathbf{a}^*(\xi)\right) \quad & \text{s.t.} \\
\sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi) \, \xi(\theta) = \mu_\theta \quad & \forall \theta \in \Theta.
\end{aligned}
$$

We finally remark the fact that from an optimal solution of the previous LP we can always recover an optimal signaling scheme $\phi : \Theta \to \Delta_{\mathcal{S}}$ setting:

$$
\phi_\theta(s) = \frac{\gamma(\xi_s) \xi_s(\theta)}{\mu_\theta} \quad \forall \theta \in \Theta, \forall s \in \mathcal{S}.
$$

3 | Problem Formulation

3.1. Bayesian ad Auction

We present now how the Bayesian Persuasion framework applies to ad auctions. In particular we present a Bayesian ad auctions, which is characterized by a finite set $\Theta := \{\theta_1, \dots, \theta_d\}$ of d states of nature. Moreover each bidder[1] $i \in \mathcal{N}$ has a valuation vector $v_i \in [0,1]^d$, with $v_i(\theta)$ being bidder i's valuation in state $\theta \in \Theta$, and all such vectors are arranged in a matrix of bidders' valuations $V \in [0,1]^{n \times d}$, whose entries are defined as $V(i,\theta) := v_i(\theta)$ for all $i \in \mathcal{N}$ and $\theta \in \Theta$. We model signaling by means of the Bayesian persuasion framework (see Section 2.4). We consider the case in which the auction mechanism (the sender) knows the state of nature and acts as a sender by issuing signals to the bidders (the receivers), so as to partially disclose information about the state and increase revenue. As customary in the literature, we assume that the state is drawn from a common prior distribution $\mu \in \Delta_\Theta$, with μ_θ denoting the probability of state $\theta \in \Theta$. The mechanism publicly commits to a signaling scheme $\phi : \Theta \to \Delta_{\mathcal{S}}$, which is a randomized mapping from states of nature to signals for the bidders. We focus on the case of public signaling in which all the bidders receive the same signal from the auction mechanism. Formally, a signaling scheme is a function $\phi : \Theta \to \Delta_{\mathcal{S}}$, where $\mathcal{S}$ is a set of available signals. For the ease of notation, we let $\phi_\theta(s)$ be the probability of sending signal $s \in \mathcal{S}$ when the state is $\theta \in \Theta$. Finally, a Bayesian ad auction goes on as follows:

1. The auction mechanism chooses a signaling scheme ϕ and the bidders observe it.

2. The mechanism observes the state of nature $\theta \sim \mu$ and draws signal $s \sim \phi_\theta$.

3. The bidders observe the signal $s \in \mathcal{S}$ and rationally update their prior belief over states according to Bayes rule. In particular, given a signal $s \in \mathcal{S}$, they infer a posterior probability $\xi_s \in \Delta_\Theta$ over the set of States of Nature, prescribed by:

$$\xi_s(\theta) := \frac{\mu_\theta \phi_\theta(s)}{\sum_{\theta' \in \Theta} \mu_{\theta'} \phi_{\theta'}(s)}. \tag{3.1}$$

[1]In this thesis we will indifferently use the terms: receivers, bidders, advertisers.

4. Finally, each bidder $i \in \mathcal{N}$ truthfully reports to the mechanism its expected valuation given the posterior $\xi_s \in \Delta_\Theta$, namely $\xi_s^\top v_i = \sum_{\theta \in \Theta} v_i(\theta)\, \xi_s(\theta)$, and the mechanism allocates slots and defines payments as in a standard ad auction (see Section 2.2).

In figure 3.1 we report the time line of a Bayesian ad auction.

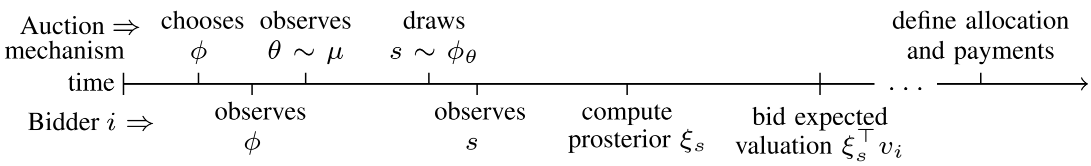

Figure 3.1: Time-line of a Bayesian ad auction.

3.2. The Revenue-maximization Problem

Our goal will be the one to design an optimal signaling scheme $\phi : \Theta \to \Delta_\mathcal{S}$ which maximizes the expected revenue of the mechanism. To do that it is oftentimes useful to represent signaling schemes as convex combinations of the posteriors they can induce. Formally, a signaling scheme $\phi : \Theta \to \Delta_\mathcal{S}$ induces a probability distribution γ over posteriors in Δ_Θ, with $\gamma(\xi)$ denoting the probability of posterior $\xi \in \Delta_\Theta$, defined as follows:

$$\gamma(\xi) \coloneqq \sum_{s \in \mathcal{S}: \xi_s = \xi} \sum_{\theta \in \Theta} \mu_\theta \phi_\theta(s).$$

Indeed, we can directly reason about distributions γ over Δ_Θ rather than about signaling schemes, provided that they are consistent with the prior. By letting $\text{supp}(\gamma) \coloneqq \{\xi \in \Delta_\Theta \mid \gamma(\xi) > 0\}$ be the support of γ, this requires that:

$$\sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi)\, \xi(\theta) = \mu_\theta \quad \forall \theta \in \Theta. \tag{3.2}$$

From now on we will use the term signaling scheme to refer to a consistent distribution γ over Δ_Θ. We focus on the problem of computing an optimal signaling scheme, *i.e.*, one maximizing the revenue of the mechanism. We study two settings:

- In Chapter 5 we discuss the Known Valuations setting in which the matrix of bidders' valuations V is known to the mechanism.
- In Chapter 6 we study the Random Valuations setting in which the matrix of bidders' valuations V is unknown, but randomly drawn according to a probability

distribution $\mathcal{V}$. Moreover, as it is customary in the literature, (see *e.g.*, Badanidiyuru et al. [2018]), in the Random Valuations setting we assume that algorithms have access to a black-box oracle returning *i.i.d.* samples drawn from $\mathcal{V}$ (rather than actually knowing such distribution).

Finally we denote by $\textsc{Rev}(V, \xi)$ the expected revenue of the mechanism when the bidders' valuations are given by V and the posterior induced by the mechanism is $\xi \in \Delta_\Theta$. Formally, given that bidders truthfully report their expected valuations and assuming w.l.o.g. that bidder $i \in \mathcal{N}$ is assigned by the mechanism to slot $i \in \mathcal{N}$, we can write $\textsc{Rev}(V, \xi) \coloneqq \sum_{j=1}^{m} j\, \xi^\top v_{j+1}(\lambda_j - \lambda_{j+1})$. Then, given a signaling scheme γ, the expected revenue of the mechanism is given by $\mathbb{E}_{\xi \sim \gamma}\left[\textsc{Rev}(V, \xi)\right]$. When the valuations are unknown, we let $\textsc{Rev}(\mathcal{V}, \xi) \coloneqq \mathbb{E}_{\xi \sim \gamma, V \sim \mathcal{V}}\left[\textsc{Rev}(V, \xi)\right]$ and define the expected revenue analogously.

4 | A General Inapproximability Result

We start our analysis with the following negative result.

Theorem 4.1. *The problem of computing an optimal signaling scheme does* not *admit a* PTAS *unless* $\mathsf{P} = \mathsf{NP}$*, even when it is restricted to the Known Valuations setting.*

Proof. We reduce from VERTEX COVER in cubic graphs. Formally, it is NP-Hard to approximate the minimum size vertex cover in cubic graph with an approximation $(1+\varepsilon)$, for a given constant $\varepsilon > 0$ (see Alimonti and Kann [2000]). Let $\eta = \varepsilon/7$ and $\delta = \eta/4$. We show that for $\delta > 0$, an $1 - \delta$ approximation to the signaling problem can be used to provide a $(1 + \varepsilon)$ approximation to vertex cover in polynomial time.

Given an instance of vertex cover (L, E) with nodes $\rho = |L|$ and edges E. For each $z \in [\rho]$, we build an instance as follows. There are $m_z = z + \rho|E| - 1$ slots and $\lambda_j = 1$ for each $j \in [m]$. The set of states is $\Theta = \{\theta_l\}_{l\in L}$ and the set of receivers is $\mathcal{N} = \{r_{e,i}\}_{e\in E, i\in[\rho]} \cup \{r_l\}_{l\in L} \cup \{r_i\}_{i\in[m+1]}$. The valuation of a receiver $r_{e,i}$, $e \in E$ and $i \in [\rho]$, is $v_{r_{e,i}}(\theta_v) = 1$ if $v \in e$, *i.e.*, e is an edge that includes v, and 0 otherwise. The valuation of a receiver r_l, $l \in L$, is $v_{r_v}(\theta_l) = 1$ and $v_{r_l}(\theta_{l'}) = 0$ for each $l' \neq l$. Moreover the valuation of a receiver $\{r_i\}$, $i \in [m + 1]$ is $v_{r_i}(\theta) = (1 - \eta)/z$ for each state $\theta \in \Theta$. Finally, the prior is uniform over all the states.

Let L^* be the minimum vertex cover and z^* be its size. We show how to build a vertex cover of size at most $z^*(1 + \varepsilon)$ from the solutions to the signaling problems instantiated with $z \in [\rho]$. For all $z \in [\rho]$, given a signaling scheme we recover a vertex cover $L(z)$ as follows. Take the posterior with larger sender's utility and add to the vertex cover $L(z)$ all the vertexes $l \in L$ such that the receiver r_l has valuation at least $1/z(1 - \frac{3}{4}\eta)$. Then, for each edge that is not covered, we add to $L(z)$ one arbitrary adjacent vertex. It is easy to see that the resulting solution $L(z)$ is a vertex cover. Finally, the algorithm returns the smallest among the vertex covers $L(z)$, $z \in [\rho]$.

We show that the vertex cover $L(z^*)$ has size at most $z^*(1+\varepsilon)$, concluding the proof. First, we show that the optimal solution of the signaling problem is at least $\frac{m_{z^*}}{z^*}(1-\frac{1}{2}\eta)$. Consider the signaling scheme with two signals s_1 and s_2 with $\phi_{\theta_l}(s_1)=1$ for each $l\in L^*$ and $\phi_{\theta_l}(s_2)=1$ for each $l\notin L^*$. In the posterior induced by s_1 the revenue is at least $\frac{m_{z^*}}{z^*}$ since all the receivers $r_{e,i}$, have expected at least $1/z^*$, while all the receivers $\{r_l\}_{l\in L^*}$ have utility at least $1/z^*$. Hence, there are at least $z^*+\rho|E|=m_{z^*}+1$ agents with valuation at least $1/z^*$. Moreover, in the posterior induced by s_2 the revenue is at least $(1-\eta)\frac{m_{z^*}}{z^*}$ since all the receivers $\{r_i\}_{i\in[m_{z^*+1}]}$ have expected valuation $(1-\eta)/z^*$. Since $z^*\geq|E|/3$ and $\rho=\frac{2}{3}|E|$ signal s_1 is sent with probability at least $z^*/\rho\geq\frac{1}{2}$ and the solution has value at least $\frac{1}{2}\frac{m_{z^*}}{z^*}+\frac{1}{2}(1-\eta)\frac{m_{z^*}}{z^*}=\frac{m_{z^*}}{z^*}(1-\frac{1}{2}\eta)$. Hence, a $1-\delta$ approximation algorithm for the signaling problem must return a signaling scheme with value at least $\frac{m_{z^*}}{z^*}(1-\frac{1}{2}\eta)(1-\delta)\geq\frac{m_{z^*}}{z^*}(1-\frac{3}{4}\eta)$. Since the expected revenue is of the signaling scheme is at least $\frac{m_{z^*}}{z^*}(1-\frac{3}{4}\eta)$, this signaling scheme sends a signal that induces a posterior $\xi\in\Delta_\Theta$ with revenue at least $\frac{m_{z^*}}{z^*}(1-\frac{3}{4}\eta)$. This implies that there are at least z^* receivers r_l with utility greater or equal to $\frac{1}{z^*}(1-\frac{3}{4}\eta)$ and that the utility of all the receivers $r_{e,i}$ is at least $\frac{1}{z^*}(1-\frac{3}{4}\eta)$. We show that our algorithm recovers a vertex cover with size at most $z^*(1+7\eta)=z^*(1+\varepsilon)$ from this posterior. Consider the set of vertexes L^1 with utility at least $\frac{1}{z^*}(1-\frac{3}{4}\eta)$. This set has size at most $z^*/(1-\frac{3}{4}\eta)$ since in each state only one receiver r_l has valuation 1 and all the other receivers r'_l, $l'\neq l$, have valuation 0. Consider the set $L^2=L\setminus L^1$ of vertexes not in this set. We have that $\sum_{l\in L^2}\xi(\theta_l)\leq\frac{3}{4}\eta$ since $\sum_{l\in L^1}\xi(\theta_l)=\sum_{l\in L^1}\xi^\top v_{r_l}\geq(1-\frac{3}{4}\eta)$. Let $\bar{E}$ be the set of edges not covered by L^1. Since each vertex has three edges, we have that $\sum_{e\in\bar{E}}\xi^\top v_{r_{e,i}}\leq 3\frac{3}{4}\eta$ for each $i\in[m_{z^*}+1]$. Moreover, since for each edge in $\bar{E}$, $\xi^\top v_{r_{e,i}}\geq(1-\frac{3}{4}\eta)/z^*$, we have that $|\bar{E}|\leq\frac{3\frac{3}{4}\eta}{(1-\frac{3}{4}\eta)/z^*}$. Then, the vertex cover build by the algorithm for $z=z^*$ includes at most

$$z^*(\frac{1}{1-\frac{3}{4}\eta}+\frac{3\frac{3}{4}\eta}{1-\frac{3}{4}\eta})\leq z^*(1+\frac{9}{4}\eta)(1+\frac{3}{2}\eta)\leq z^*(1+7\eta)=z^*(1+\varepsilon).$$

□

Theorem 4.1 is proved by a reduction from the VERTEX COVER problem in cubic graphs Alimonti and Kann [2000]. In the rest of this Thesis, we study several settings in which the negative result in Theorem 4.1 can be circumvented, by either fixing some parameters of the problem (see Section 5.2 and Section 6.2) or considering instances with a specific structure (see Sections 5.3 and 6.3).

5 | Known Valuations Setting

5.1. Introduction

In this section we focus on the problem of computing a revenue-maximizing signaling scheme when the bidder's valuation matrix $V \in [0,1]^{n \times d}$ is known by the mechanism. In particular we formulate the problem by means of the following LP[1].

$$\max_{\gamma \geq 0} \sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi) \, \textsc{Rev}(V, \xi) \quad \text{s.t.} \tag{5.1a}$$

$$\sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi) \, \xi(\theta) = \mu_\theta \qquad \forall \theta \in \Theta. \tag{5.1b}$$

Due the negative result presented in Chapter 4 we will focus on two particular setting. In the first one we will look for an optimal signaling scheme when either the number of states of nature d or the number of slots m is fixed (see Section 5.2), while in the second setting we will tackle the problem when the bidders are single minded according to the definition presented in Section 5.3.

5.2. Parametrized Complexity

We first study the parametrized complexity of the problem of computing an optimal signaling scheme, showing that it admits a polynomial-time algorithm when either the number of slots m or the number of states of nature d is fixed. In the following, we let $\Pi_l \subseteq 2^{\mathcal{N}}$ be the set of all the the possible permutations of $l \leq n$ bidders taken from $\mathcal{N}$, with $\pi = (i_1, ..., i_l) \in \Pi_l$ denoting a tuple made by bidders $i_1, \ldots, i_l \in \mathcal{N}$, in that order. We also let $\Xi_\pi \subseteq \Delta_\Theta$ be the (possibly empty) set of posteriors in which the expected valuations of bidders in $\pi \in \Pi_l$ are ordered (from the highest to the lowest) according to π. Formally we define Ξ_π as follows.[2]

[1] Since, in next sections, we will prove that there always exists an optimal distribution γ with finite support we can safely write the problem as follows.

[2] we denote with π_j the $j-th$ element of the tuple π.

Definition 5.1. *Given an integer* $l \leq n$ *and a tuple* $\pi = (i_1, ..., i_l) \in \Pi_l$ *we let* $\Xi_\pi \subseteq \Delta_\Theta$ *be such that:*

$$\Xi_\pi := \left\{\xi \in \Delta_\Theta \mid \xi^\top v_{\pi_1} \geq \xi^\top v_{\pi_2} \geq \ldots \geq \xi^\top v_{\pi_l}\right\}.$$

We first observe that for each possible integer $l \leq n$ and tuple $\pi \in \Pi_l$ the region $\Xi_\pi \subseteq \Delta_\Theta$ is a convex polytope. Moreover we observe that, given a permutation $\pi \in \Pi_l$ of $l \geq m+1$ bidders, the expected revenue of the mechanism in any posterior $\xi \in \Xi_\pi$ is $\textsc{Rev}(V, \xi) = \sum_{\theta \in \Theta} \xi(\theta) \sum_{j=1}^{m} j v_{\pi_{j+1}}(\theta)(\lambda_j - \lambda_{j+1})$, since the bidders truthfully report their expected valuations to the mechanism, and, thus, the latter allocates slots to bidders in π according to their order in the permutation. Thus, for any fixed $\pi \in \Pi_l$ with $l \geq m+1$, the term $\textsc{Rev}(V, \xi)$ is linear in $\xi \in \Xi_\pi$.

5.2.1. Fixing the Number of Slots

To find a revenue-maximizing signaling scheme in this setting we start proving the following lemma.

Lemma 5.1. *There always exists a revenue-maximizing signaling scheme* γ *such that* $|\Xi_\pi \cap \mathrm{supp}(\gamma)| \leq 1$ *for every* $\pi \in \Pi_{m+1}$.

Proof. To prove the result we show that, given a signaling scheme γ that induces two posteriors $\xi, \xi' \in \Xi_\pi$ for a $\pi \in \Pi_{m+1}$, we can recover a signaling scheme γ^* with at least the same revenue that replaces the two posteriors ξ and ξ' with a convex combination of them. Let $\xi, \xi' \in \Xi_\pi$ be two elements belonging to the support of an optimal signaling scheme γ. In order to show the result we introduce a posterior probability ξ^* as follows: $\xi^* = z\xi_1 + (1-z)\xi_2$ with $z = \gamma(\xi_1)/(\gamma(\xi_1) + \gamma(\xi_2))$. Since Ξ_π is a convex polytope each convex combination of a subset of its elements belongs to it, so that $\xi^* \in \Xi_\pi$. Moreover, we define a new signaling scheme γ^* as follows: $\gamma^*(\xi^*) = \gamma(\xi_1) + \gamma(\xi_2)$ and $\gamma^*(\xi_1) = \gamma^*(\xi_2) = 0$ while $\gamma^*(\xi_i) = \gamma(\xi_i)\ \forall i \neq 1, 2$. To conclude the proof, we observe that the two signaling schemes gain the same revenue. Indeed, by linearity we have: $\gamma(\xi^*)\textsc{Rev}(V, \xi^*) = (\gamma(\xi_1) + \gamma(\xi_2))\textsc{Rev}(V, z\xi_1 + (1-z)\xi_2) = \gamma(\xi_1)\textsc{Rev}(V, \xi_1) + \gamma(\xi_2)\textsc{Rev}(V, \xi_2)$. □

Intuitively, the lemma follows from the fact that, given any signaling scheme γ and two posteriors $\xi, \xi' \in \mathrm{supp}(\gamma)$ such that $\xi, \xi' \in \Xi_\pi$ for some $\pi \in \Pi_{m+1}$, it is always possible to define a new signaling scheme that replaces ξ and ξ' with a suitably-defined convex combination of them, without decreasing the expected revenue (since it is linear over Ξ_π). To formulate the problem as an LP we introduce a variable for each $\pi \in \Pi_{m+1}$ and $\theta \in \Theta$ in order to encode the products $\gamma(\xi_\pi)\xi_\pi(\theta)$. Formally given a tuple $\pi \in \Pi_{m+1}$ and $\xi \in \Xi_\pi$, we define $x_\pi(\theta) := \gamma(\xi)\xi(\theta)$ for each $\theta \in \Theta$ as the posterior probability multiplied by the probability of state $\theta \in \Theta$ in $\xi \in \Delta_\Theta$. Notice that by Lemma 5.1 there is at most one $\xi \in \Xi_\pi$ belonging to the support of an optimal γ for each possible tuple $\pi \in \Pi_{m+1}$. Thus, we can represent our optimization problem with the following LP.

$$\max_{x \in [0,1]^{|\Pi_{m+1}||\Theta|}} \sum_{\pi \in \Pi_{m+1}} \sum_{\theta \in \Theta} x_\pi(\theta) \sum_{j=1}^{m} j\, v_{\pi_{j+1}}(\theta)(\lambda_j - \lambda_{j+1}) \quad \text{s.t.} \tag{5.2a}$$

$$\sum_{\pi \in \Pi_{m+1}} x_\pi(\theta) = \mu_\theta \qquad \forall \theta \in \Theta \tag{5.2b}$$

$$\sum_{\theta \in \Theta} x_\pi(\theta)[v_{\pi_j}(\theta) - v_{\pi_{j+1}}(\theta)] \geq 0 \qquad \forall \pi \in \Pi_{m+1}, j \in [m] \tag{5.2c}$$

In Theorem 5.1 we finally prove that an optimal revenue-maximizing signaling scheme can be found in polynomial time if m is fixed. This holds because LP 5.2 has a polynomial number of variables and constraints and an optimal solution of LP 5.2 is always associated to a signaling scheme γ gaining the same revenue.

Theorem 5.1. *In the Known Valuations setting, if the number of slots m is fixed, then an optimal signaling scheme can be computed in polynomial time.*

Proof. As a first step we observe that the number of variables in LP 5.2 results equal to $O(n^{m+1}d)$ while the number of constraints is equal to $O(n^{m+1} + d)$. Due to that, as long as m is fixed, LP 5.2 is solvable in polynomial time. In addition, we show that from an optimal solution of LP 5.2 we can always find a signaling scheme that provides the same revenue. This holds setting $\gamma(\xi_\pi) = \sum_{\theta \in \Theta} x_\pi(\theta)$ for each $\pi \in \Pi_{m+1}$ and $\xi_\pi(\theta) = x_\pi(\theta)/\gamma(\xi_\pi)$ for each $\pi \in \Pi_{m+1}$ and $\theta \in \Theta$ if $\gamma(\xi_\pi) \neq 0$. Finally the two gains the same revenue indeed we have:

$$\begin{aligned}\sum_{\pi \in \Pi_{m+1}} \gamma(\xi_\pi) \sum_{\theta \in \Theta} \sum_{j=1}^{m} j\, \xi_\pi(\theta) v_{i_{j+1}}(\theta)(\lambda_j - \lambda_{j+1}) &= \sum_{\pi \in \Pi_{m+1}} \sum_{\theta \in \Theta} \gamma(\xi_\pi)\xi_\pi(\theta) \sum_{j=1}^{m} j\, v_{i_{j+1}}(\theta)(\lambda_j - \lambda_{j+1}) \\ &= \sum_{\pi \in \Pi_{m+1}} \sum_{\theta \in \Theta} x_\pi(\theta) \sum_{j=1}^{m} j\, v_{i_{j+1}}(\theta)(\lambda_j - \lambda_{j+1})\end{aligned}$$

□

We finally remark that setting $m = 1, \lambda_1 = 1$ and $\lambda_2 = 0$ in LP 5.2 we find what presented by Emek et al. [2014] in the easier case of a second price auction.

5.2.2. Fixing the Number of States of Nature

To design a polynomial-time algorithm in this scenario, we restrict the attention to those signaling schemes supported on a suitable defined finite set of posterior probabilities $\Xi^* \subset \Delta_\Theta$. Formally we define[3] $\Xi^* := \bigcup_{\pi \in \Pi_n} V(\Xi_\pi)$ as the union of the vertexes of the previous introduced regions Ξ_π over all the possible tuple $\pi \in \Pi_n$. In the following lemma we prove that there always exists an optimal signaling scheme supported in $\Xi^* \subset \Delta_\Theta$.

Lemma 5.2. *There always exists a revenue-maximizing signaling scheme γ such that supp$(\gamma) \subseteq \Xi^*$.*

Proof. First, we observe that for each $\xi \in \Delta_\Theta$ there exists a tuple $\pi \in \Pi_n$ such that $\xi \in \Xi_\pi$, this easily follow from the fact that $\bigcup_{\pi \in \Pi_n} \Xi_\pi = \Delta_\Theta$. As observed before in such regions the revenue is a linear function. Thus, it is possible to decompose each posterior $\xi \in \Xi_\pi$ by Caratheodory's theorem as a convex combination of the vertexes of Ξ_π without decreasing the revenue. Formally, for each $\pi \in \Pi_n$ and each posterior $\xi \in \Xi_\pi$, there exists a distribution $\gamma_\xi \in \Delta_{V(\Xi_\pi)}$ such that:

$$\xi(\theta) = \sum_{\tilde{\xi} \in \Xi^*} \gamma_\xi(\tilde{\xi}) \tilde{\xi}(\theta) \;\; \forall \theta \in \Theta.$$

We show that such a decomposition does not affect the final revenue. Indeed, by linearity we get the following:

$$\sum_{\tilde{\xi} \in V(\Xi_\pi)} \gamma_\xi(\tilde{\xi}) \textsc{Rev}(V, \tilde{\xi}) = \textsc{Rev}\Big(V, \sum_{\tilde{\xi} \in V(\Xi_\pi)} \gamma_\xi(\tilde{\xi})\tilde{\xi}\Big) = \textsc{Rev}(V, \xi).$$

To conclude the proof, we show that given the optimal distribution γ, we can recover a distribution $\gamma^* \in \Delta_{\Xi^*}$ with the same revenue. In particular, $\gamma^* \in \Delta_{\Xi^*}$ is such that:

$$\gamma^*(\tilde{\xi}) = \sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi)\gamma_\xi(\tilde{\xi}) \quad \forall \tilde{\xi} \in \Xi^*.$$

Since γ satisfies the consistency constraints (see Equation 3.2), it easy to see that also $\gamma^* \in \Delta_{\Xi^*}$ satisfies the constraints. Moreover, the two distributions provide the same

[3] Given a poplytope X we denote with $V(X)$ the set of its vertexes.

revenue. Indeed, we have

$$\begin{aligned}\sum_{\tilde{\xi}\in\Xi^*}\gamma^*(\tilde{\xi})\ \textsc{Rev}(V,\tilde{\xi}) &= \sum_{\xi\in\text{supp}(\gamma)}\gamma(\xi)\sum_{\tilde{\xi}\in\Xi^*}\gamma_\xi(\tilde{\xi})\textsc{Rev}(V,\tilde{\xi})\\ &= \sum_{\xi\in\text{supp}(\gamma)}\gamma(\xi)\ \textsc{Rev}(V,\xi).\end{aligned}$$

□

The lemma follows from the fact that, given any signaling scheme γ and posterior $\xi \in \text{supp}(\gamma)$ such that $\xi \in \Xi_\pi$ for some $\pi \in \Pi_n$, by Carathéodory's theorem it is always possible (since Ξ_π is a polytope) to decompose $\xi \in \Delta_\Theta$ into a convex combination of the vertices of Ξ_π, obtaining a new signaling scheme that provides the mechanism with an expected revenue at least as large as that of γ (since $\textsc{Rev}(V,\xi)$ is linear over Ξ_π). By observing that $|\Xi^*| = O((n^2+d)^{d-1})$, it is easy to show that an optimal signaling scheme can be computed by means of LP 5.1 taking $\text{supp}(\gamma) = \Xi^*$, which has a number of variables and constraints that is polynomial once d is fixed. This proves the following:

Theorem 5.2. *In the Known Valuations setting, if the number of states d is fixed, then an optimal signaling scheme can be computed in polynomial time.*

Proof. We first observe that the vertexes of each region Ξ_π are identify by the intersection of $d-1$ linear independent hyperplanes for each $\pi \in \Pi_n$. Moreover, we note that each one of these vertexes is identified by a subset of the $O(n^2)$ constraints $\xi^\top v_i \geq \xi^\top v_j$ for each $i \neq j \in \mathcal{N}$ and the d constraint ensuring that $\xi(\theta) \geq 0$ for each $\theta \in \Theta$. Hence, the total number of vertexes defining the previous discussed regions will be at most $\binom{n^2+d}{d-1} = O((n^2+d)^{d-1})$. Finally, we notice that, as long as d is a fixed parameter, it is possible to find an optimal signaling scheme in polynomial time solving LP 5.1 setting $\text{supp}(\gamma) = \Xi^*$. □

5.3. Single-Minded Bidders

In this section, we focus on particular Bayesian ad auctions where the bidders are single minded. Intuitively, in our setting, by single mindedness we mean that each bidder is interested in displaying their ad only when the realized state of nature is a specific (single) state, and that all the bidders interested in the same state value a click on their ad for the same amount. We introduce the following formal definition:

Definition 5.2 (Single-minded bidders)**.** *In a Bayesian ad auction, we say that bidders are* single minded *if there exist $\mathcal{N}_\theta \subseteq \mathcal{N}$ and $\delta_\theta \in [0,1]$ for all $\theta \in \Theta$ such that:*

(i) $\mathcal{N} = \bigcup_{\theta\in\Theta} \mathcal{N}_\theta$ *and* $\mathcal{N}_\theta \cap \mathcal{N}_{\theta'} = \emptyset$ *for all* $\theta \neq \theta' \in \Theta$*;*

(ii) for every $\theta \in \Theta$ *and* $i \in \mathcal{N}_\theta$*, it holds* $v_i(\theta) = \delta_\theta$ *and* $v_i(\theta') = 0$ *for all* $\theta' \in \Theta : \theta' \neq \theta$*.*

To better understand definition 5.2 we present an introductory example.

Example 5.1. *Let $\Theta = \{\theta_1, \theta_2, \theta_3\}$ be the set of states of Nature and let be $\mu \in \Delta_\Theta$ a uniform prior distribution. Let $|\mathcal{N}_{\theta_1}| = 10$, $|\mathcal{N}_{\theta_2}| = 3$ and $|\mathcal{N}_{\theta_3}| = 3$ and let $\delta_{\theta_i} = 1$ for $i = 1, 2, 3$. In addition, we consider an auction with $m = 5$ slots with CTRs given by $\Lambda = \{1, 0.9, 0.8, 0.7, 0.6\}$. In this case it is possible to design an optimal signaling (see Table 5.1) $\phi : \Theta \to \Delta_\mathcal{S}$ scheme just considering two different signals $\mathcal{S} = \{s_1, s_2\}$.*

	state θ_1	state θ_2	state θ_3
signal s_1	1	0	0
signal s_2	0	1	1

Table 5.1: Probability to send signal $s \in \mathcal{S}$ when $\theta \in \Theta$ is drawn.

Note that when bidders receive the signal s_1 the posterior probability they infer is given by $\xi_{s_1} = (1, 0, 0)$, on the contrary when they receive the signal s_2 the posterior probability they infer is equal to $\xi_{s_2} = (0, 0.5, 0.5)$. It is possible to show that the revenue prescribed by an optimal signaling scheme is equal to:

$$\sum_{s\in\mathcal{S}} \gamma(\xi_s)\mathrm{Rev}(\xi_s, V) = \gamma(\xi_{s_1})\mathrm{Rev}(\xi_{s_1}, V) + \gamma(\xi_{s_2})\mathrm{Rev}(\xi_{s_2}, V)$$
$$= \frac{4}{3} + \frac{4}{3} = \frac{8}{3}$$

Moreover, to satisfy constraints of Equation 3.2 we have that $\gamma(\xi_{s_1}) = 1/3$ and $\gamma(\xi_{s_2}) = 2/3$. On the other hand the expected revenue that does not commit to a signaling scheme is equal to[4]*:*

$$\sum_{\theta\in\Theta} \mu_\theta \mathrm{Rev}(\mathbb{I}_\theta, V) = \frac{1}{3}\left(4 + \frac{3}{5} + \frac{3}{5}\right) = \frac{26}{15}$$

Notice that, given a posterior $\xi \in \Delta_\Theta$ induced by the mechanism, all the bidders belonging to the same set $\mathcal{N}_\theta$ have the same expected valuation, namely $\xi^\top v_i = \delta_\theta \xi(\theta)$ for all $\theta \in \Theta$

[4]We denote with $\mathbb{I}_\theta \in \Delta_\Theta$ the vector having all components equal to zero except the one corresponding to $\theta \in \Theta$ which is equal to 1.

and $i \in \mathcal{N}_\theta$. As a result, given that bidders truthfully report their expected valuations, the mechanism will always receive at most d different bids, one per set $\mathcal{N}_\theta$. The last observation implies that, given $\xi \in \Delta_\Theta$, in order to unequivocally define an allocation of bidders to slots (and, thus, also define the expected payments) it is sufficient to know the relative ordering of the (at most) d different expected valuations associated to sets $\mathcal{N}_\theta$. This allows us to tackle the problem with an approach analogous to the one of Section 5.2, with the only difference that, in this case, we will reason about permutations of the groups of bidders $\mathcal{N}_\theta$, rather than about permutations of all the individual bidders. In the following, we let $\Pi \subseteq 2^\Theta$ be the set of all the permutations of the sates of nature $\Theta = \{\theta_1, \dots, \theta_d\}$, while we let $\pi = (\theta_{k_1}, \dots, \theta_{k_d}) \in \Pi$ be an ordered tuple made by states $\theta_{k_1}, \dots, \theta_{k_d} \in \Theta$, where $k_1, \dots, k_d \in [d]$. Moreover the polytope of posteriors in which the expected valuations associated to sets $\mathcal{N}_\theta$ are ordered according to π is given by:

$$\Xi_\pi := \left\{ \xi \in \Delta_\Theta \mid \delta_{\theta_{k_1}} \xi(\theta_{k_1}) \geq \dots \geq \delta_{\theta_{k_d}} \xi(\theta_{k_d}) \right\}$$

The first preliminary result that we need in order to derive our approximation algorithm is a characterization of the vertices of the sets Ξ_π for $\pi \in \Pi$, as follows.

Lemma 5.3. *Given $\pi \in \Pi$ and $\xi \in \Xi_\pi$, it holds that $\xi \in V(\Xi_\pi)$ if and only if there exists $\ell \in [d]$ such that (i) and (ii) are both satisfied:*

(i) $\delta_{\theta_{k_1}} \xi(\theta_{k_1}) = \dots = \delta_{\theta_{k_\ell}} \xi(\theta_{k_\ell}) > 0$

(ii) $\delta_{\theta_{k_{\ell+1}}} \xi(\theta_{k_{\ell+1}}) = \dots = \delta_{\theta_{k_d}} \xi(\theta_{k_d}) = 0$

Proof. First, we show that if a posterior $\xi \in \Xi_\pi$ satisfies (i) and (ii), then $\xi \in V(\Xi_\pi)$ and thus it is a vertex of Ξ_π. In particular, $\xi \in \Xi_\pi$ satisfies the linear independent equality $\delta_{\theta_{k_j}} \xi(\theta_{k_j}) = \delta_{\theta_{k_{j+1}}} \xi(\theta_{k_{j+1}})$ for each $j \in [\ell - 1]$. Moreover, it satisfies $\delta_{\theta_{k_j}} \xi(\theta_{k_j}) = 0$ for each $j \in \{\ell + 1, \dots, d\}$ and the simplex equality $\sum_{\theta \in \Theta} \xi(\theta) = 1$. Hence, $\xi \in \Xi_\pi$ is at the intersection of d linear independent hyperplanes defining Ξ_π and it is a vertex of Ξ_π. To conclude the proof, we show that each vertex $\xi \in \Xi_\pi$ satisfies (i) and (ii). In particular, we show that given a posterior $\xi \in \Xi_\pi$ such that $\delta_{\theta_{k_{j^*}}} \xi(\theta_{k_{j^*}}) > \delta_{\theta_{k_{j^*+1}}} \xi(\theta_{k_{j^*+1}}) > 0$ for a $j^* \in [d]$, and thus it does not satisfies (i) and (ii), the posterior is at the the intersection of at most $d - 1$ linear independent hyperplanes. Consider the hyperplanes $\delta_{\theta_{k_j}} \xi(\theta_{k_j}) = \delta_{\theta_{k_{j+1}}} \xi(\theta_{k_{j+1}})$ for $j \leq j^* - 1$. Notice that $\xi \in \Xi_\pi$ satisfies at most all the $j^* - 1$ equalities of this kind.[5] Moreover, consider all the $j > j^*$ with $\delta_{\theta_{k_{j^*}}} \xi(\theta_{k_{j^*}}) > 0$. Let j^{**} be the largest j that satisfies this condition. By a similar argument as above, we can show that there are

[5]Notice that the equality $\delta_{\theta_{k_j}} \xi(\theta_{k_j}) = \delta_{\theta_{k_{j'}}} \xi(\theta_{k_{j'}})$ with $j' > j$ and$|j - j'| \geq 2$ is linear dependent from the equalities $\delta_{\theta_{k_{\bar{j}}}} \xi(\theta_{k_{\bar{j}}}) = \delta_{\theta_{k_{\bar{j}+1}}} \xi(\theta_{k_{\bar{j}+1}})$ for each $\bar{j} \in \{j, \dots, j' - 2\}$.

at most $j^{**} - j^* - 1$ linear equalities $\delta_{\theta_{k_j}} \xi(\theta_{k_j}) = \delta_{\theta_{k_{j+1}}} \xi(\theta_{k_{j+1}})$ with $j^* + 1 \leq j \leq j^{**} - 1$. Finally, for all the $j > j^{**}$, the equalities $\delta_{\theta_{k_j}} \xi(\theta_{k_j}) = 0$ are satisfied. Hence, including the simplex constraint there are at most $j^* - 1 + j^{**} - j^* - 1 + d - j^{**} + 1 = d - 1$ linear independent equalities, concluding the proof. □

Intuitively, Lemma 5.3 states that the vertices of a set Ξ_π are all the posteriors $\xi \in \Delta_\Theta$ such that, for some $\ell \in [d]$, only the first ℓ states according to the ordering defined by π are assigned a positive probability, while all the remaining states have zero probability. Moreover, the positive probabilities of the posterior $\xi \in \Delta_\Theta$ are defined so that all the bidders belonging to the first ℓ sets $\mathcal{N}_\theta$, according to the ordering defined by π, are the same. Notice that, in the special case in which all the values δ_θ are equal to one, the vertices of all the sets Ξ_π are all the uniform probability distributions over subsets of ℓ states of nature, for any $\ell \in [d]$. By letting $\Xi^* = \bigcup_{\pi \in \Pi} V(\Xi_\pi)$, since the term $\textsc{Rev}(V, \xi)$ is linear in ξ over Ξ_π for every permutation $\pi \in \Pi$, we can conclude that $\text{OPT}_{\Xi^*} = \text{OPT}$ (the proof is analogous to that of Lemma 5.2). Thus, Lemma 5.3 allows us to find an optimal signaling scheme by solving LP 5.1 taking $\text{supp}(\gamma) = \Xi^*$ and the matrix of bidders' valuations V. However, notice that, since the size of Ξ^* is exponential in d, the resulting LP has exponentially-many variables. Nevertheless, since the LP has polynomially-many constraints, we can still solve it in polynomial time by applying the ellipsoid algorithm to its dual, provided that a polynomial-time separation oracle is available. In order to design a polynomial-time separation oracle, we apply the procedure described above to a relaxed version of LP 5.1, whose optimal value is sufficiently close to that of the original LP. Given $\beta \in \mathbb{R}_+$, the relaxed LP reads as follows:

$$\max_{\gamma \in \Delta_{\Xi^*}, z \leq 0} \sum_{\xi \in \Xi} \gamma(\xi) \, \textsc{Rev}(\mathcal{V}, \xi) + \beta z \quad \text{s.t.} \tag{5.4a}$$

$$\sum_{\xi \in \Xi^*} \gamma(\xi) \, \xi(\theta) - z \geq \mu_\theta \qquad \forall \theta \in \Theta. \tag{5.4b}$$

The dual problem of LP 5.4 reads as follows:

$$\min_{y \leq 0, t} \sum_{\theta \in \Theta} y_\theta \mu_\theta + t \quad \text{s.t.} \tag{5.5a}$$

$$\sum_{\theta \in \Theta} y_\theta \, \xi(\theta) + t \geq \textsc{Rev}(\mathcal{V}, \xi) \qquad \forall \xi \in \Xi^* \tag{5.5b}$$

$$\sum_{\theta \in \Theta} y_\theta \geq -\beta, \tag{5.5c}$$

where $y_\theta \in \mathbb{R}_-$ for $\theta \in \Theta$ are dual variables associated to Constraints (5.1b), while t is a dual variable for $\sum_{\xi \in \Xi^*} \gamma(\xi) = 1$. Notice that, by relaxing the LP, in the dual LP 5.5 we get the additional Constraint (5.5c) and that $y_\theta \leq 0$ for all $\theta \in \Theta$. This is crucial to design a polynomial-time separation oracle. The separation problem associated to Problem 5.5 reads as:

Definition 5.3 (Separation problem). *Given values for the dual variables $y_\theta \in [-\beta, 0]$ for all $\theta \in \Theta$, compute:*

$$\max_{\xi \in \Xi^*} \text{Rev}(V, \xi) - \sum_{\theta \in \Theta} y_\theta \, \xi(\theta). \tag{5.6}$$

The following Lemma 5.4 shows that Problem 5.6 can be solved optimally up to any given additive loss $\lambda > 0$, by means of a dynamic programming algorithm that runs in time polynomial in the size of the input, in $\frac{1}{\lambda}$, and in β. Formally:

Lemma 5.4. *Given $\lambda > 0$, there exists an algorithm that finds an additive λ-approximation to the Separation problem 5.6, in time polynomial in the size of the input, in $\frac{1}{\lambda}$, and in β.*

Proof. Let $f(v, j)$ be the revenue when $j \in [n]$ bidders have expected valuation $v \in [0, 1]$ and all the other bidders have expected valuation 0. Moreover, given a set $E \subseteq \mathbb{R}$ and an $x \in \mathbb{R}$, let $\lfloor x \rfloor_E$ be equal to the largest element $e \in E$ such that $e \leq x$. Similarly, we define $\lceil x \rceil_E$ be equal to the smallest element $e \in E$ such that $e \geq x$. We show that Algorithm 5.1 provides the desired guarantees. It is easy to see that the algorithm runs in polynomial time. Let $\xi^* \in \Xi^*$ be the optimal solution to $\max_{\xi \in \Xi^*} \text{Rev}(V, \xi) - \sum_{\theta \in \Theta} y_\theta \, \xi(\theta)$. Notice that by the definition of Ξ^* there exists a subset of states $\Theta^* \subseteq \Theta$ and a value v^* such that $\xi^*(\theta)\delta_\theta = v^*$ for each $\theta \in \Theta^*$ and $\xi^*(\theta) = 0$ for each $\theta \notin \Theta^*$. Our first step is to show that this solution $\xi^* \in \Xi^*$ corresponds to a feasible solution of the algorithm. Formally, we show that Θ^* is a feasible subset of states for $\Theta(v, d, w, \sum_{\theta \in \Theta^*} |\mathcal{N}_\theta|)$. In particular, it is sufficient to prove that $w = \sum_{\theta \in \Theta^*} \lfloor v/\delta_\theta \rfloor_E \leq \sum_{\theta \in \Theta^*} v/\delta_\theta \leq \sum_{\theta \in \Theta^*} v^*/\delta_\theta = 1$, with $v = \lfloor v^* \rfloor_G$. Since this solution is feasible, it provides a lower bound on the value $\max_{v \in G, j \in [n], w \in E} f(v, j) + M(v, d, w, j)$. Moreover, let $\theta^* \in \Theta^*$ be the state in Θ^* with smallest δ_θ, we provide a bound on v^*. In particular it holds that $v^* \leq \delta_{\theta^*}$, otherwise $\xi^*(\theta^*) = v^*/\delta_{\theta^*} > 1$ and $v^* \geq \delta_{\theta^*}/d$, otherwise $\sum_{\theta \in \Theta^*} v^*/\delta_\theta < 1$. This implies that for each $\theta \in \Theta^*$ we have $v/\delta_\theta = \lfloor v^* \rfloor_G/\delta_\theta \geq (v^* - \varepsilon\delta_{\theta^*})/\delta_\theta \geq v^*/\delta_\theta - \varepsilon$ and $v \geq v^* - \varepsilon$. Now, we can provide a lower bound on $\max_{v \in G, j \in [n], w \in E} f(v, j) + M(v, d, w, j)$. In particular, the solution that takes states Θ^*, $v = \lfloor v^* \rfloor_G$, and $w = \sum_{\theta\in\Theta^*} \lfloor v/\delta_\theta \rfloor_E$ has value at least:

$$\begin{aligned} f(v, \sum_{\theta\in\Theta^*} |\mathcal{N}_\theta|) - \sum_{\theta\in\Theta^*} \frac{y_\theta v}{\delta_\theta} &\geq f(v^*, \sum_{\theta\in\Theta^*} |\mathcal{N}_\theta|) - \varepsilon m - \sum_{\theta\in\Theta^*} \frac{y_\theta v}{\delta_\theta} \\ &\geq f(v^*, \sum_{\theta\in\Theta^*} |\mathcal{N}_\theta|) - \sum_{\theta\in\Theta^*} y_\theta \xi(\theta) - \varepsilon m - d\beta\varepsilon \\ &= \max_{\xi\in\Xi^*} [\, \textsc{Rev}(V, \xi) - \sum_{\theta\in\Theta} y_\theta\, \xi(\theta)] - \varepsilon m - d\beta\varepsilon, \end{aligned}$$

where the first inequality comes from Lipschitz continuity of $f(\cdot, x)$, *i.e.*, $f(v^*, x) - f(v, x) \leq m|v^* - v|$ for each $x \in \mathbb{N}_+$ and $v = \lfloor v^* \rfloor_G \geq v^* - \varepsilon$, while the second inequality comes from $-y_\theta v/\delta_\theta \geq -y_\theta(v^*/\delta_\theta - \varepsilon) \geq -y_\theta v^*/\delta_\theta - \beta\varepsilon$ for each $\theta \in \Theta^*$. To conclude the proof, we show that from a solution $(\hat{v}, \hat{j}, \hat{w})$ and $\hat{\Theta}$ the algorithm find a posterior $\hat{\xi} \in \Delta_\Theta$ with value at least $f(\hat{v}, \hat{j}) + M(\hat{v}, d, \hat{w}, \hat{j}) - \varepsilon dm - 2\beta d\varepsilon$. First, we bound the value of w_{real}. In particular, $w_{real} \leq 1 + \varepsilon d$ since for each state $\theta \in \Theta$, $\hat{v}/\delta_\theta - \lfloor \hat{v}/\delta_\theta \rfloor_E \leq \varepsilon$ and $\hat{w} \leq 1$. Hence, in the posterior $\hat{\xi} \in \Delta_\Theta$ all the bidders have valuation at least $\hat{v}/(1+\varepsilon d) \geq \hat{v} - \varepsilon d$ and $\textsc{Rev}(V, \hat{\xi}) = f(\hat{v} - \varepsilon d, \sum_{\theta\in\hat{\Theta}} |\mathcal{N}(\theta)|) \geq f(\hat{v}, \sum_{\theta\in\hat{\Theta}} |\mathcal{N}(\theta)|) - \varepsilon dm$ by the Lipschitz continuity of $f(\cdot, x)$. Now, we consider the component $M(\hat{v}, d, \hat{w}, \hat{j})$. In particular, we show that $\sum_{\theta\in\hat{\Theta}} -\hat{\xi}(\theta) y_\theta \geq M(\hat{v}, d, \hat{w}, \hat{j}) - 2\beta d\varepsilon$. Since $y_\theta \leq 0$ for each $\theta \in \Theta$, it holds:

$$\begin{aligned} \sum_{\theta\in\hat{\Theta}} -\hat{\xi}(\theta) y_\theta &= \sum_{\theta\in\hat{\Theta}} -\frac{\hat{v}}{\delta_\theta w_{real}} y_\theta \\ &= \frac{M(\hat{v}, d, \hat{w}, \hat{j})}{w_{real}} \\ &\geq \frac{M(\hat{v}, d, \hat{w}, \hat{j})}{(1 + d\varepsilon)} \\ &\geq M(\hat{v}, d, \hat{w}, \hat{j}) - \beta d\varepsilon - \beta d^2\varepsilon^2, \end{aligned}$$

where the last inequality comes from $M(\hat{v}, d, \hat{w}, \hat{j}) \leq \beta w_{real} \leq \beta(1 + d\varepsilon)$ and $1/(1 + d\varepsilon) \geq 1 - d\varepsilon$. To conclude, the value of the solution $\hat{\xi} \in \Delta_\Theta$ is an additive λ-approximation with $\lambda = \varepsilon m + \varepsilon dm + 2\beta d\varepsilon + \beta d^2\varepsilon^2$. □

Algorithm 5.1 Dynamic programming algorithm in the proof of Lemma 5.4

Require: $\varepsilon > 0$.

$c \leftarrow \lceil 1/\varepsilon \rceil$
$E \leftarrow \{i/c\}_{i=0}^{c}$
$G \leftarrow \cup_{\theta \in \Theta} \{\delta_\theta i/c\}_{i=0}^{c}$
initialize empty M and Θ with dimension $c \times d \times c \times n$
for $v \in G$ **do**
  for $w \in E, w \geq \lfloor v/\delta_{\theta_1} \rfloor_E$ **do**
    $M(v, 1, w, |\mathcal{N}_{\theta_1}|) \leftarrow -y_{\theta_1} v/\delta_{\theta_1}$
    $\Theta(v, 1, w, |\mathcal{N}_{\theta_1}|) \leftarrow \{\theta_1\}$
  end for
  for $i \in [d]$,$w \in E$,$j \in [n]$ **do**
    if $M(v, i-1, w, j) \geq M(v, i-1, w - \lceil v/\delta_{\theta_i} \rceil_E, j - |\mathcal{N}_{\theta_i}|) - y_{\theta_i} v/\delta_{\theta_i}$ **then**
      $M(v, i, w, j) \leftarrow M(v, i-1, w, j)$
      $\Theta(v, i, w, j) \leftarrow \Theta(v, i-1, w, j)$
    else
      $M(v, i, w, j) \leftarrow M(v, i-1, w - \lceil v/\delta_{\theta_i} \rceil_E, j) - y_{\theta_i} v/\delta_{\theta_i}$
      $\Theta(v, i, w, j) \leftarrow \Theta(v, i-1, w - \lceil v/\delta_{\theta_i} \rceil_E, j) \cup \{\theta_i\}$
    end if
  end for
end for
$(\hat{v}, \hat{j}, \hat{w}) \leftarrow \arg\max_{v \in E, j \in [n], w \in E} f(v, j) + M(v, d, w, j)$
$\hat{\Theta} \leftarrow \Theta(\hat{v}, d, \hat{w}, \hat{j})$
$w_{real} = \sum_{\theta \in \hat{\Theta}} \hat{v}/\delta_\theta,$
for $\theta \in \hat{\Theta}$ **do**
  $\hat{\xi}(\theta) = \frac{\hat{v}}{\delta_\theta w_{real}}$
end for
return $\hat{\xi} \in \Delta_\Theta$

The crucial observation that allows us to solve Problem 5.6 by means of dynamic programming is that, in any posterior $\xi \in \Xi^*$, bidders' expected valuations are either a positive, bidder-independent value or zero (see Lemma 5.3). This allows us to build a discretized range of possible bidders' valuation values, so that, for each discretized value, we can compute an optimal posterior $\xi \in \Xi^*$ inducing that value by adding states of nature incrementally in a dynamic programming fashion. Since the algorithm in Lemma 5.4 only returns an approximate solution to Problem 5.6, we need to carefully apply the ellipsoid algorithm to solve LP 5.5, so that it correctly works even with an approximated oracle.

Some non-trivial duality arguments allow us to prove that, indeed, this can be achieved by only incurring in a small additive loss on the quality of the returned solution, and without degrading the running time of the algorithm. Overall, this allows us to conclude that:

Theorem 5.3. *In the Known Valuations setting, if the bidders are single minded, then the problem of computing an optimal signaling scheme admits an additive FPTAS.*

Proof. Our FPTAS is described in Algorithm 5.2.

Algorithm 5.2 FPTAS in the proof of Theorem 5.3

Require: parameter of the relaxed LP $\beta > 0$, approximation factor of the approximation oracle $\lambda > 0$, error $\eta > 0$.

Initialization: $\rho_1 \leftarrow 0$, $\rho_2 \leftarrow m$, $H \leftarrow \emptyset, H^\star \leftarrow \emptyset$.
while $\rho_2 - \rho_1 > \eta$ **do**
  $\rho_3 \leftarrow (\rho_1 + \rho_2)/2$
  $H \leftarrow$ {posteriors relative to the violated constraints returned by the ellipsoid method on F with objective ρ_3 and approximation error η}
  if unfeasible **then**
    $\rho_1 \leftarrow \rho_3$
    $H^\star \leftarrow H$
  else
    $\rho_2 \leftarrow \rho_3$
  end if
  $(\gamma, z) \leftarrow$ solution to LP 5.14 with only posteriors in $H^\star$
  return the solution $\bar{\gamma}$ corresponding to the solution of the relaxed problem (γ, z)
end while

We start providing the following relaxation to LP 5.1 for a value $\beta \in \mathbb{R}_+$ defined in the following.

$$\max_{\gamma \in \Delta_{\Xi^*}, z \leq 0} \sum_{\xi \in \Xi^*} \gamma(\xi) \, \textsc{Rev}(\mathcal{V}, \xi) + \beta z \quad \text{s.t.} \tag{5.8a}$$

$$\sum_{\xi \in \Xi^*} \gamma(\xi) \, \xi(\theta) - z \geq \mu_\theta \qquad \forall \theta \in \Theta. \tag{5.8b}$$

Given a solution (γ, z) to LP 5.8, we can find an approximate solution to LP 5.1 as follows. First, notice that by the optimality of (γ, z), we have $\sum_{\xi \in \Xi^*} \gamma(\xi) \, \textsc{Rev}(\mathcal{V}, \xi) + \beta z \geq 0$ and

$z \geq -m/\beta$. Moreover we define $\bar{\mu}_\theta = \sum_{\xi \in \Xi^*} \gamma(\xi)\, \xi(\theta)$, so that we have $|\bar{\mu}_\theta - \mu_\theta| \leq m/\beta$ for each $\theta \in \Theta$. Consider a distribution $\bar{\gamma}$ such that: $\bar{\gamma}(\xi) = (1 - m/\beta)\gamma(\xi)$ for each $\xi \in \Delta_\Theta$ except for the vertexes of the simplex of Δ_Θ. Moreover, for each $\theta \in \Theta$ we define $\bar{\gamma}(\mathbb{I}_\theta) = \gamma(\mathbb{I}_\theta)(1 - m/\beta) + \mu_\theta - \sum_{\xi \in \Xi^*} \gamma(\xi)\xi(\theta)(1 - m/\beta)$. We observe that $\bar{\gamma}$ is a feasible solution to LP 5.1 since $\sum_{\xi \in \Xi^*} \bar{\gamma}(\xi)\xi(\theta) = \mu_\theta$ for each $\theta \in \Theta$ and it has value at least $\text{OPT}_{\Xi^*} - 2m^2/\beta$ since the distribution γ is scaled by a factor $(1 - m/\beta)$ and $\text{OPT}_{\Xi^*} \leq m$. Hence, to provide an approximation to LP 5.1, it is sufficient to provide an approximation to LP 5.8 for a sufficiently large $\beta > 0$. Since LP 5.8 has an exponential number of variables, the algorithm works by applying the ellipsoid method to the following dual problem.

$$\min_{y \leq 0, t} \sum_{\theta \in \Theta} y_\theta \mu_\theta + t \quad \text{s.t.} \tag{5.9a}$$

$$\sum_{\theta \in \Theta} y_\theta\, \xi(\theta) + t \geq \textsc{Rev}(V, \xi) \qquad \forall \xi \in \Xi^* \tag{5.9b}$$

$$\sum_{\theta \in \Theta} y_\theta \geq -\beta, \tag{5.9c}$$

where the dual variables are $y \in \mathbb{R}^d_-$ and $t \in \mathbb{R}$. Instead of an exact separation oracle, we use an approximate separation oracle that employs Algorithm 5.1 with a suitably-defined approximation $\lambda > 0$. We use a binary search scheme to find a value $\rho^\star \in [0, m]$ such that the dual problem with objective $\rho^\star$ is unfeasible, while the dual with objective $\rho^\star + \eta$ is approximately feasible, for some $\eta \geq 0$ defined in the following. The algorithm requires $\log(m/\eta)$ steps and, at each step, it works by determining, for a given value ρ_3, whether there exists a feasible solution for the following feasibility problem that we call F:

$$\sum_{\theta \in \Theta} y_\theta \mu_\theta + t \leq \rho_3 \tag{5.10a}$$

$$\sum_{\theta \in \Theta} y_\theta\, \xi(\theta) + t \geq \textsc{Rev}(V, \xi) \qquad \forall \xi \in \Xi^* \tag{5.10b}$$

$$\sum_{\theta \in \Theta} y_\theta \geq -\beta \tag{5.10c}$$

$$y_\theta \leq 0 \qquad \forall \theta \in \Theta. \tag{5.10d}$$

At each iteration of the bisection algorithm, the feasibility problem F is solved via the ellipsoid method. To do so, we need a separation oracle. We focus on an approximate separation oracle that returns a violated constraint. The oracle is implemented as follows. Given (y, t), first it check if all the $y_\theta \leq 0$ for each $\theta \in \Theta$, $\sum_{\theta \in \Theta} y_\theta \geq -\beta$, and $\sum_{\theta \in \Theta} y_\theta \mu_\theta + t \leq \rho_3$. If it is not the case, it returns a violated constraint. Otherwise, it uses Algorithm

5.1 to approximate:

$$\max_{\xi \in \Xi^*} \textsc{Rev}(V, \xi) - \sum_{\theta \in \Theta} y_\theta \, \xi(\theta). \tag{5.11}$$

Notice, that we guarantee that Algorithm 5.1 is called with $y_\theta \geq -\beta$ for each $\theta \in \Theta$. Let ξ be the returned posterior. If there returned posterior has value at least t, the separation oracle returns the constraint relative to posterior $\xi \in \Delta_\Theta$. Otherwise, it returns feasible. The bisection procedure terminates when it determines a value $\rho^\star$ such that on (F) the ellipsoid method returns unfeasible for $\rho^\star$, while returning feasible for $\rho^\star + \eta$. Then, the algorithm solves a modified primal LP 5.14 with only the subset of posteriors in $H^\star$, where $H^\star$ is the set of posteriors relative to the violated constraints returned by the ellipsoid method applied on the unfeasible problem with objective $\rho^\star$. Finally, it computes a solution $\bar{\gamma}$ from the solution γ of LP 5.14 as observed before. Now, we prove the approximation guarantees of the algorithm. The algorithm finds a $\rho^\star$ such that the problem is unfeasible, *i.e.*, the value of ρ_1 when the algorithm terminates, and a value smaller than or equal to $\rho^\star + \eta$ such that the ellipsoid method returns feasible, *i.e.*, the value of ρ_2 when the algorithm terminates. In particular, we show that OPT $\leq \rho^\star + \eta + \lambda$, where OPT is the value of LP 5.8. Since, the bisection algorithm returns that (F) is feasible with objective $\rho^\star + \eta$, it finds a solution (y, t) such that the approximate separation oracle did not find a violated constraint. We show that (y, t) is a solution to the following LP.

$$\sum_{\theta \in \Theta} y_\theta \mu_\theta + t \leq \rho^\star + \eta \tag{5.12a}$$

$$\sum_{\theta \in \Theta} y_\theta \, \xi(\theta) + t \geq \textsc{Rev}(V, \xi) - \lambda \qquad \forall \xi \in \Xi^* \tag{5.12b}$$

$$\sum_{\theta \in \Theta} y_\theta \geq -\beta \tag{5.12c}$$

$$y_\theta \leq 0 \qquad \forall \theta \in \Theta. \tag{5.12d}$$

This holds because we have shown that, when the separation oracle returns feasible, it holds $\max_{\xi \in \Xi^*} [\textsc{Rev}(V, \xi) - \sum_{\theta \in \Theta} y_\theta \, \xi(\theta)] \leq t + \lambda$ by the approximation guarantees of Algorithm 5.1, implying that all the Constraints (5.12b) are satisfied. Moreover, when the separation oracle returns feasible all the other constraints are satisfied. Then, by strong duality the value of the following LP is at most $\rho^\star + \eta$.

$$\max_{\gamma \in \Delta_{\Xi^*}, z \leq 0} \sum_{\xi \in \Xi^*} \gamma(\xi) \left(\textsc{Rev}(V, \xi) - \lambda \right) + \beta z \quad \text{s.t.} \tag{5.13a}$$

$$\sum_{\xi \in \Xi^*} \gamma(\xi) \, \xi(\theta) - z \geq \mu_\theta \qquad \forall \theta \in \Theta. \tag{5.13b}$$

Notice that any solution to LP 5.1 is also a feasible solution to the previous modified problem. Since in any feasible solution $\sum_{\xi \in \Xi^*} \gamma(\xi) = 1$ and LP 5.13 has value at most $\rho^\star + \eta$, then $\text{OPT} \le \rho^\star + \eta + \lambda$, where we recall that OPT is the value of LP 5.8. Let $H^\star$ be the set of posteriors relative to the constraints returned by the ellipsoid method run with objective $\rho^\star$. Since the ellipsoid method with the approximate separation oracle returns unfeasible, by strong duality LP 5.8 with only the variables $\gamma(\xi)$ relative to constraints in $H^\star$ has value at least $\rho^\star$. Moreover, since the ellipsoid method guarantees that $H^\star$ has polynomial size, the LP can be solved in polynomial time. Hence, solving the following LP, *i.e.*, the primal of LP 5.9 with only the variables $\gamma(\xi)$ in $H^\star$, we can find a solution with value at least $\rho^\star$.

$$\max_{\gamma \in \Delta_{H^\star}, z \le 0} \sum_{\xi \in H^\star} \gamma(\xi) \, \textsc{Rev}(V, \xi) + \beta z \quad \text{s.t.} \tag{5.14a}$$

$$\sum_{\xi \in H^\star} \gamma(\xi) \, \xi(\theta) - z \ge \mu_\theta \qquad \forall \theta \in \Theta. \tag{5.14b}$$

To conclude the proof, notice that the algorithm provides an $2m^2/\beta + \eta + \lambda$, where the term $2m^2/\beta$ is due to the relaxation of the primal and $\eta + \lambda$ to the use of an approximate separation oracle. Since the algorithm runs in time polynomial in β, $1/\eta$ and $1/\lambda$, we can provide an arbitrary good approximations choosing sufficiently small values of η and λ, and a sufficiently large value for β. $\square$

6 | Random Valuations Setting

6.1. Introduction

Differently from Chapter 5 we assume now that each valuation vector $v_i \in [0,1]^d$, for each bidder $i \in \mathcal{N}$, is independently drawn from a probability distribution $\mathcal{V}_i$ and we indicate with $\mathcal{V} = \{\mathcal{V}_i\}_{i \in \mathcal{N}}$ the collection of such distributions. In addition we denote $V \sim \mathcal{V}$ to indicate that the matrix $V \in [0,1]^{n \times d}$ is built drawing each v_i independently from $\mathcal{V}_i$ while we let $\textsc{Rev}(\mathcal{V}, \xi) \coloneqq \mathbb{E}_{V \sim \mathcal{V}}[\textsc{Rev}(V, \xi)]$ be the expected revenue over the latter distribution. Finally, given a finite set $\Xi \subset \Delta_\Theta$ of posteriors, we formulate the problem of computing a revenue-maximizing signaling scheme supported in $\Xi \subset \Delta_\Theta$ with the following LP.

$$\max_{\gamma \geq 0} \sum_{\xi \in \Xi} \gamma(\xi) \, \textsc{Rev}(\mathcal{V}, \xi) \quad \text{s.t.} \tag{6.1a}$$

$$\sum_{\xi \in \Xi} \gamma(\xi) \, \xi(\theta) = \mu_\theta \qquad \forall \theta \in \Theta. \tag{6.1b}$$

In the following we let OPT_Ξ be the optimal value of LP 6.1, while we let OPT be the optimal expected revenue of the mechanism over all the possible signaling schemes γ. In this setting we assume that the auction mechanism has access to the distribution of bidders' valuations $\mathcal{V}$ only through a black-box sampling oracle. In particular, given $s \in \mathbb{N}_+$ *i.i.d.* samples of matrices of bidders' valuations, namely $V_1, \ldots, V_s \in [0,1]^{n \times d}$, we let $\mathcal{V}^s$ be their empirical distribution, given by:

$$\Pr_{V \sim \mathcal{V}^s}\left\{V = \hat{V}\right\} \coloneqq \frac{\sum_{t=1}^{s} \mathbb{I}\{V_t = \hat{V}\}}{s} \qquad \forall \hat{V} \in [0,1]^{n \times d}. \tag{6.2}$$

In this section we provide two preliminary useful lemmas. The first one (Lemma 6.1) works under the true distribution of bidders' valuations $\mathcal{V}$, and it establishes a connection between the optimal expected revenue and the optimal value of LP 6.1 for suitably-defined finite sets $\Xi \subset \Delta_\Theta$ of posteriors. In particular, we look at sets $\Xi \subset \Delta_\Theta$ for which the function $\textsc{Rev}(\mathcal{V}, \cdot)$ is stable according to the following definition.

Definition 6.1 ((α, ε)-stability)**.** *Given $\alpha, \varepsilon \geq 0$ and a finite set $\Xi \subset \Delta_\Theta$, we say that $\textsc{Rev}(\mathcal{V}, \cdot)$ is (α, ε)-stable for Ξ if, for every $\xi \in \Delta_\Theta$, there exists a distribution $\gamma_\xi \in \Delta_\Xi$ such that:*

$$\mathbb{E}_{\xi' \sim \gamma_\xi}\left[\textsc{Rev}(\mathcal{V}, \xi')\right] \geq (1-\alpha)\textsc{Rev}(\mathcal{V}, \xi) - \varepsilon. \tag{6.3}$$

Definition 6.1 plays a crucial role indeed, given a finite set $\Xi \subset \Delta_\Theta$ of posteriors such that $\textsc{Rev}(\mathcal{V}, \cdot)$ is (α, ε)-stable for it, an optimal solution of LP 6.1 restricted to $\Xi \subset \Delta_\Theta$ has value at least $(1-\alpha)\text{OPT} - \varepsilon$. Formally the following lemma holds.

Lemma 6.1. *Given $\alpha, \varepsilon \geq 0$ and $\Xi \subset \Delta_\Theta$ such that $\textsc{Rev}(\mathcal{V}, \cdot)$ is (α, ε)-stable for Ξ, then it holds $\text{OPT}_\Xi \geq (1-\alpha)\text{OPT} - \varepsilon$.*

Proof. Let γ be an optimal distribution over Δ_Θ satisfying the consistency constraints (see Equation 3.2), we define $\gamma^* \in \Delta_\Xi$ as follow:

$$\gamma^*(\tilde{\xi}) = \sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi)\gamma_\xi(\tilde{\xi}) \quad \forall \tilde{\xi} \in \Xi,$$

where γ_ξ is the distribution over $\Xi \subset \Delta_\Theta$ that satisfies Definition 6.1. It is easy to check that also $\gamma^* \in \Delta_\Xi$ satisfies the consistency constraints. Moreover, since $\textsc{Rev}(\mathcal{V}, \cdot)$ is (α, ϵ)-stable for $\Xi \subset \Delta_\Theta$, we get:

$$\begin{aligned}
\sum_{\tilde{\xi} \in \Xi} \gamma^*(\tilde{\xi})\textsc{Rev}(\mathcal{V}, \tilde{\xi}) &= \sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi) \sum_{\tilde{\xi} \in \Xi} \gamma_\xi(\tilde{\xi})\textsc{Rev}(\mathcal{V}, \tilde{\xi}) \\
&\geq \sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi)((1-\alpha)\textsc{Rev}(\mathcal{V}, \xi) - \varepsilon) \\
&= (1-\alpha) \sum_{\xi \in \text{supp}(\gamma)} \gamma(\xi)\textsc{Rev}(\mathcal{V}, \xi) - \varepsilon \\
&= (1-\alpha)\text{OPT} - \varepsilon,
\end{aligned}$$

proving the lemma. □

The second lemma (Lemma 6.4) deals with the approximation error introduced by using an empirical distribution of bidders' valuations $\mathcal{V}^s$, rather than the actual distribution $\mathcal{V}$ (see Equation 6.2 for its definition). Specifically, given a finite set $\Xi \subset \Delta_\Theta$ of posteriors we let $\gamma_{\mathcal{V}^s} \in \Delta_\Xi$ be an optimal solution of LP 6.1 for distribution $\mathcal{V}^s$ and set Ξ. Moreover, we define:

$$\text{OPT}_{\Xi,s} := \mathbb{E}\Big[\sum_{\xi \in \Xi} \gamma_{\mathcal{V}^s}(\xi)\textsc{Rev}(\mathcal{V}, \xi)\Big] \tag{6.4}$$

as the average expected[1] revenue of signaling schemes $\gamma_{\mathcal{V}^s}$ under the true distribution of valuations $\mathcal{V}$. Then, a concentration argument proves the following:

Lemma 6.2. *Given $\rho, \tau > 0$, let $\Xi \subseteq \Delta_\Theta$ be finite and $s = \left\lceil \frac{2(\lambda_1 m)^2}{\tau^2} \log \frac{2}{\rho} \right\rceil$ then* $\mathrm{OPT}_{\Xi,s} \geq (1 - \rho|\Xi|)\, \mathrm{OPT}_\Xi - \tau$.

Proof. We first observe that for each $\xi \in \Delta_\Theta$ we have:

$$\begin{aligned}\textsc{Rev}(\mathcal{V}, \xi) &\leq \sum_{j=1}^{m} j\ (\lambda_j - \lambda_{j+1}) \\ &\leq m \sum_{j=1}^{m} (\lambda_j - \lambda_{j+1}) = \lambda_1 m.\end{aligned}$$

So that by Hoeffding bound we have:

$$\Pr\Big(\big|\textsc{Rev}(\mathcal{V}_s, \xi) - \textsc{Rev}(\mathcal{V}, \xi)\big| \leq \tau/2\Big) \geq 1 - 2e^{\frac{-s\tau^2}{2(\lambda_1 m)^2}} = 1 - \rho$$

where the inequality is attained considering a number of samples $s = \left\lceil \frac{2(\lambda_1 m)^2}{\tau^2} \log(\frac{2}{\rho}) \right\rceil$. Moreover, by union bound and De Morgan's laws, we get:

$$\Pr\Big(\bigcap_{\xi \in \Xi} \Big\{\big|\textsc{Rev}(\mathcal{V}_s, \xi) - \textsc{Rev}(\mathcal{V}, \xi)\big| \leq \tau/2\Big\}\Big) \geq 1 - \rho|\Xi|.$$

Let $\gamma^* \in \Delta_\Xi$ be an optimal solution of LP 6.1 while let $\gamma_{\mathcal{V}_s} \in \Delta_\Xi$ an optimal solution of the same LP when the distribution is the empirical one. We observe that the expected revenue attained with $\gamma_{\mathcal{V}_s} \in \Delta_\Xi$ is greater or equal to the one prescribed by $\gamma^* \in \Delta_\Xi$ minus a fixed parameter with a probability of at least $1 - \rho|\Xi|$. Formally it holds:

$$\begin{aligned}\sum_{\xi \in \Xi} \gamma_{\mathcal{V}_s}(\xi) \textsc{Rev}(\mathcal{V}, \xi) &\geq \sum_{\xi \in \Xi} \gamma_{\mathcal{V}_s}(\xi) \textsc{Rev}(\mathcal{V}_s, \xi) - \tau/2 \\ &\geq \sum_{\xi \in \Xi} \gamma^*(\xi) \textsc{Rev}(\mathcal{V}_s, \xi) - \tau/2 \\ &\geq \sum_{\xi \in \Xi} \gamma^*(\xi) \textsc{Rev}(\mathcal{V}, \xi) - \tau.\end{aligned}$$

[1] The expectation is with respect to the sampling procedure that determines $\mathcal{V}^s$.

Finally, we have that:

$$
\begin{aligned}
\mathrm{OPT}_{\Xi,s} &= \mathbb{E}\Big[\sum_{\xi\in\Xi}\gamma_{\mathcal{V}_s}(\xi)\textsc{Rev}(\mathcal{V},\xi)\Big] \\
&\geq \Big(1-\rho|\Xi|\Big)\mathbb{E}\Big[\sum_{\xi\in\Xi}\gamma_{\mathcal{V}_s}(\xi)\textsc{Rev}(\mathcal{V},\xi)\,|\bigcap_{\xi\in\Xi}\Big\{\big|\textsc{Rev}(\mathcal{V}_s,\xi)-\textsc{Rev}(\mathcal{V},\xi)\big|\leq\tau/2\Big\}\Big] \\
&\geq \Big(1-\rho|\Xi|\Big)\sum_{\xi\in\Xi}\gamma^*(\xi)\textsc{Rev}(\mathcal{V},\xi)-\tau \\
&\geq \Big(1-\rho|\Xi|\Big)\mathrm{OPT}_{\Xi}-\tau
\end{aligned}
$$

proving the lemma. □

We also introduce a family of posteriors depending on a parameter $q \in \mathbb{N}_+$ which are those encoded as particular q-uniform probability distributions, as formally stated in the following definition.

Definition 6.2 (q-uniform posterior)**.** *A posterior $\xi \in \Delta_\Theta$ is* q-uniform *if it can be obtained by averaging the elements of a multi-set defined by $q \in \mathbb{N}_+$ canonical basis of $\mathbb{R}^d$.*

We denote the set of all q-uniform posteriors for a given $q \in \mathbb{N}_+$ as $\Xi_q \subset \Delta_\Theta$ and we provide a visual representation of the set Ξ_q for $q = 4$ and $d = 3$.

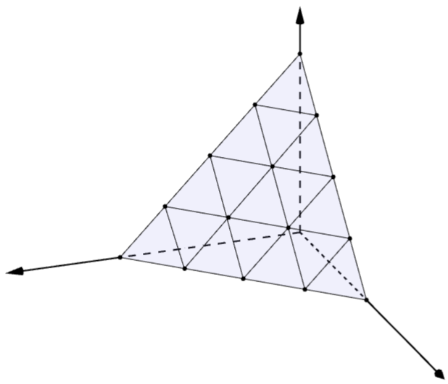

Figure 6.1: The set Ξ_q with $q = 4$ and $d = 3$.

Intuitively the set Ξ_q represents a uniform grid over the Simplex Δ_Θ and its cardinality depends both on the number of states of nature d and the parameter q as stated in the following lemma.

Lemma 6.3. *Given $q, d \in \mathbb{N}_+$ such that $d \geq 2$, $q \geq 1$ it holds:*

$$|\Xi_q| = \binom{q+d-1}{d-1}$$

Proof. Let $S(d,q) = |\Xi_q|$ such that $\Xi_q \subset \Delta_\Theta$ with $|\Theta| = d$. It is easy to check that $S(d,q)$ is recursively defined as follows: $S(d,q) = S(d,q-1) + S(d-1,q)$ with $d \geq 2$, $q \geq 1$ and board conditions: $S(2,q) = q+1$, $S(d,1) = d$. Thanks to the well known property of binomial coefficients[2] it is possible to check that a solution of the previous equation is given by $S(d,q) = \binom{q+d-1}{d-1}$. □

Finally we observe that Lemma 6.3 ensures that if the parameter $q \in \mathbb{N}_+$ is fixed then $|\Xi_q| = O(q^d)$ while, if the parameter $d \in \mathbb{N}_+$ is fixed, it holds $|\Xi_q| = O(d^q)$.

6.2. Parametrized Complexity

Analogously to what we have done before in the Known Valuation case, we study the paramatrized complexity of computing a revenue maximizing-signaling scheme. In particular in this section we show that the problem admits an additive FPTAS when the number of states of Nature d is fixed while the problem admits an additive QPTAS when the number of slots m is fixed.

6.2.1. Fixing the Number of States of Nature

First, we study the computational complexity of the problem of computing an optimal signaling scheme when the number of states d is fixed. We provide an (additive) FPTAS that works by performing the following two steps: (i) it collects a suitable number $s \in \mathbb{N}_{>0}$ of matrices of bidders' valuations, by invoking the sampling oracle; and (ii) it solves LP 6.1 for the resulting empirical distribution $\mathcal{V}^s$ and a suitably-defined set of q-uniform posteriors. In particular, given a desired (additive) error $\lambda > 0$, the algorithm works on the set Ξ_q for $q = \lceil \frac{md}{\lambda} \rceil$ and its approximation guarantees rely on the following Lemma 6.4, proved again by means of Lemma 6.1.

[2] $\binom{n}{k} = \binom{n}{k-1} + \binom{n-1}{k-1}$ for each $n, k \in \mathbb{N}_+$ such that $k \leq n$.

Lemma 6.4. *Given $\lambda > 0$ and $q = \lceil \frac{md}{\lambda} \rceil$, then* $\text{OPT}_{\Xi_q} \geq OPT - \lambda$.

Proof. First, we show that the revenue is a Lipschitz continuous function in the posterior probability $\xi \in \Delta_\Theta$ with respect to the infinity norm. In particular, it holds:

$$|\textsc{Rev}(\mathcal{V}, \xi) - \textsc{Rev}(\mathcal{V}, \xi')| \leq md||\xi - \xi'||_\infty \quad \forall \xi, \xi' \in \Delta_\Theta.$$

This follows from the fact that $\textsc{Rev}(\mathcal{V}, \xi)$ is a piecewise linear, continuous function in $\xi \in \Delta_\Theta$ and $||\nabla_\xi \textsc{Rev}(\mathcal{V}, \xi)||_1 \leq md$ almost everywhere. In addition we define:

$$I_\lambda(\xi) = \{\xi' \in \Delta_\Theta \mid ||\xi - \xi'||_\infty \leq \lambda/md\}$$

as the neighbourhood of a given posterior $\xi \in \Delta_\Theta$ and $\Xi(\xi) = I_\lambda(\xi) \cap \Xi_q$ its intersection with the set Ξ_q. It is easy to see that $\xi \in co(\Xi(\xi))$. Hence, by Caratheodory's theorem we can decompose each $\xi \in \Delta_\Theta$ as follow:

$$\sum_{\tilde{\xi} \in \Xi(\xi)} \gamma_\xi(\tilde{\xi})\tilde{\xi}(\theta) = \xi(\theta) \quad \forall \theta \in \Theta$$

with $\gamma_\xi \in \Delta_{\Xi(\xi)}$. We show now that such a decomposition decreases the expected revenue under $\xi \in \Delta_\Theta$ of at most a fixed parameter. Formally, we have that:

$$\begin{aligned}
\mathbb{E}_{\tilde{\xi} \sim \gamma_\xi}\left[\textsc{Rev}(\mathcal{V}, \tilde{\xi})\right] &= \sum_{\tilde{\xi} \in \Xi(\xi)} \gamma_\xi(\tilde{\xi}) \textsc{Rev}(\mathcal{V}, \tilde{\xi}) \\
&\geq \sum_{\tilde{\xi} \in \Xi(\xi)} \gamma_\xi(\tilde{\xi})(\textsc{Rev}(\mathcal{V}, \xi) - \lambda) \\
&= \textsc{Rev}(\mathcal{V}, \xi) - \lambda,
\end{aligned}$$

where the inequality comes from the Lipschitz continuity of $\textsc{Rev}(\mathcal{V}, \xi)$ and
$||\xi - \tilde{\xi}||_\infty \leq \lambda/md$ for all $\tilde{\xi} \in \Xi(\xi)$. This proves that the revenue is $(0, \lambda)$-stable over Ξ_q and by Lemma 6.1 we have that $\text{OPT}_{\Xi_q} \geq \text{OPT} - \lambda$ □

Thanks to Lemmas 6.2 and 6.4 (the former applied for suitable values $\rho, \tau > 0$), we can prove that the procedure described in steps (i) and (ii) above gives a signaling scheme achieving an expected revenue at most a function of λ lower than OPT, provided that the number of samples s is defined as in Lemma 6.4. Moreover, since d is fixed thanks to Lemma 6.3 we have that $|\Xi_q| = O(q^d) = O((\frac{1}{\lambda}md)^d)$ and then the overall procedure runs in time polynomial in the input size and in $\frac{1}{\lambda}$. Thus, we can conclude that:

Theorem 6.1. *In the Random Valuations setting, if the number of states d is fixed, then the problem of computing an optimal signaling scheme admits and additive FPTAS.*

Proof. Let $\eta > 0$ be the desired approximation and let τ, α, and q be three suitable values defined in the following. Applying Lemma 6.2 for $\Xi = \Xi_q$, $\rho = \frac{\alpha}{m}$, τ and $s = \lceil \frac{2m^2}{\tau^2} \log(\frac{2m}{\alpha}) \rceil$, we get:

$$\begin{aligned}
\text{OPT}_{\Xi_{q,s}} &= \mathbb{E}\Big[\sum_{\xi \in \Xi_q} \gamma_{\mathcal{V}_s}(\xi) \text{Rev}(\mathcal{V}, \xi) \Big] \\
&\geq \Big(1 - \frac{\alpha |\Xi_q|}{m}\Big) \text{OPT}_{\Xi_q} - \tau \\
&\geq \text{OPT}_{\Xi_q} - \tau - \alpha |\Xi_q|.
\end{aligned}$$

By Lemma 6.4, for a value λ defined in the following and $q = \lceil \frac{md}{\lambda} \rceil$ we have that:

$$\begin{aligned}
\text{OPT}_{\Xi_{q,s}} &= \mathbb{E}\Big[\sum_{\xi \in \Xi_q} \gamma_{\mathcal{V}_s}(\xi) \text{Rev}(\mathcal{V}, \xi) \Big] \\
&\geq \text{OPT}_{\Xi_q} - \tau - \alpha |\Xi_q| \\
&\geq \text{OPT} - \lambda - \tau - \alpha |\Xi_q| \\
&= \text{OPT} - \eta,
\end{aligned}$$

where the last equality holds taking $\lambda = \eta/3$, $\varepsilon = \eta/6$, $\alpha = \eta/(3|\Xi_q|)$. □

6.2.2. Fixing the Number of Slots

Next, we switch the attention to the case in which the number of slots m is fixed. We provide an additive PTAS that works as the FPTAS in Theorem 6.1, but whose approximation guarantees follow from Lemma 6.2 and Lemma 6.5 (rather than Lemma 6.4). Thus, the only difference with respect to the previous case is that the algorithm works on the set Ξ_q of q-uniform posteriors for q defined as in Lemma 6.5. As a result, since $|\Xi_q| = O(d^q)$ and q depends on a parameter $\eta > 0$ that is related to the quality of the obtained approximation, the algorithm is only a PTAS rather than an FPTAS. Formally, we can prove the following:

Lemma 6.5. *Given $\eta > 0$ and $q = \left\lceil \frac{1}{2\eta^2} \log \frac{m+1}{\eta} \right\rceil$, then $\text{OPT}_{\Xi_q} \geq \text{OPT} - 2\eta m$.*

Proof. We show that there exists a distribution $\gamma \in \Delta_{\Xi_q}$ over q-uniform posterior that provides an expected revenue that satisfies the conditions in the statement. For a $\xi \in \Delta_\Theta$, let $\xi^q \in \Xi_q$ be the empirical mean of q vectors built form q *i.i.d.* samples drawn from the given posterior $\xi \in \Delta_\Theta$. In particular, each sample is obtained by randomly drawing a state of nature, with each state $\theta \in \Theta$ having probability $\xi(\theta)$ of being selected, and, then, a d-dimensional vector is built by letting all its components equal to 0, except for that one corresponding to $\theta \in \Theta$, which is set to 1. Notice that ξ^q is a random vector supported on q-uniform posteriors, whose expected value is posterior $\xi \in \Delta_\Theta$. Then, we let $\gamma_\xi \in \Delta_{\Xi^q}$ be such that, for every $\tilde{\xi} \in \Xi^q$, it holds $\gamma_\xi(\tilde{\xi}) = \Pr\{\xi^q = \tilde{\xi}\}$. Moreover, in the following, given a $\xi \in \Delta_\Theta$ and a $j \in [m+1]$ we write the expected valuation $\xi^\top v_{i_j}$ without specifying that $\pi = (i_1, ..., i_{m+1}) \in \Pi_{m+1}$ is the tuple such that $\xi \in \Xi_\pi$. Then, by Hoeffding bound we have that:

$$\Pr_{\tilde{\xi}\sim\gamma_\xi}\Big(\tilde{\xi}^\top v_{i_j} \geq \xi^\top v_{i_j} - \eta\Big) \geq 1 - e^{-2q\eta^2} = 1 - \frac{\eta}{m+1} \quad \forall j \in [m+1]$$

where the equality follows from the definition of the number of posteriors q. Thanks to union bound and De Morgan's laws we get:

$$\Pr_{\tilde{\xi}\sim\gamma_\xi}\Big(\bigcap_{j=1}^{m+1}\{\tilde{\xi}^\top v_{i_j} \geq \xi^\top v_{i_j} - \varepsilon\}\Big) \geq 1 - \eta.$$

Now we prove that the revenue is $(0, 2\eta m)$-stable over Ξ_q.

$$\begin{aligned}
\mathbb{E}_{\tilde{\xi}\sim\gamma_\xi}\Big[\textsc{Rev}(\mathcal{V},\tilde{\xi})\Big] &\geq \mathbb{E}_{\tilde{\xi}\sim\gamma_\xi}\Big[\textsc{Rev}(\mathcal{V},\tilde{\xi}) | \bigcap_{j=1}^{m+1}\{\tilde{\xi}^\top v_{i_j} \geq \xi^\top v_{i_j} - \eta\}\Big]\Pr\Big(\bigcap_{j=1}^{m+1}\{\tilde{\xi}^\top v_{i_j} \geq \xi^\top v_{i_j} - \eta\}\Big) \\
&\geq (1-\eta)\mathbb{E}_{\tilde{\xi}\sim\gamma_\xi}\Big[\textsc{Rev}(\mathcal{V},\tilde{\xi}) | \bigcap_{j=1}^{m+1}\{\tilde{\xi}^\top v_{i_j} \geq \xi^\top v_{i_j} - \eta\}\Big] \\
&\geq (1-\eta)\Big(\sum_{j=1}^{m} j\,(\lambda_j - \lambda_{j+1})(\xi^\top v_{i_{j+1}} - \eta)\Big) \\
&\geq (1-\eta)\big(\textsc{Rev}(\mathcal{V},\xi) - \eta m\big) \\
&\geq \textsc{Rev}(\mathcal{V},\xi) - 2\eta m
\end{aligned}$$

he latter inequality shows that the revenue is $(0, 2\eta m)$-stable over Ξ_q with $q = \lceil \frac{1}{2\eta^2}\log\big(\frac{m+1}{\eta}\big)\rceil$. Hence, by Lemma 6.1, $\text{OPT}_{\Xi_q} \geq \text{OPT} - 2\eta m$ proving the lemma. □

Theorem 6.2. *In the Random Valuations setting, if the number of slots m is fixed, then the problem of computing an optimal signaling scheme admits and additive PTAS.*

Proof. Let $\nu > 0$ be the desired approximation and let τ, α, and q be three suitable values defined in the following. Applying Lemma 6.2 for $\Xi = \Xi_q$, $\rho = \frac{\alpha}{m}$, τ and $s = \lceil \frac{2m^2}{\tau^2}\log\frac{2m}{\alpha}\rceil$, we get:

$$\begin{aligned}\mathrm{OPT}_{\Xi_{q,s}} &= \mathbb{E}\Big[\sum_{\xi\in\Xi_q}\gamma_{\mathcal{V}_s}(\xi)\textsc{Rev}(\mathcal{V},\xi)\Big]\\ &\geq \Big(1-\frac{\alpha|\Xi_q|}{m}\Big)\mathrm{OPT}_{\Xi_q}-\tau\\ &\geq \mathrm{OPT}_{\Xi_q}-\tau-\alpha|\Xi_q|.\end{aligned}$$

By Lemma 6.5, we have that for a value η defined in the following and $q=\lceil\frac{1}{2\eta^2}\log\frac{m+1}{\eta}\rceil$ it holds:

$$\begin{aligned}\mathrm{OPT}_{\Xi_{q,s}} &= \mathbb{E}\Big[\sum_{\xi\in\Xi_q}\gamma_{\mathcal{V}_s}(\xi)\textsc{Rev}(\mathcal{V},\xi)\Big]\\ &\geq \mathrm{OPT}_{\Xi_q}-\tau-\alpha|\Xi_q|\\ &\geq \mathrm{OPT}-2\eta m-\tau-\alpha|\Xi_q|\\ &= \mathrm{OPT}-\nu.\end{aligned}$$

Where the last equality holds for $\eta=\nu/6m$, $\tau=\nu/3$, and $\alpha=\nu/\left(3|\Xi_q|\right)$. □

6.3. Valuations Bounded Away From Zero

We conclude the section by studying the case in which the bidders' valuations are bounded away from zero. This case is dealt with an algorithm identical to the one in Theorem 6.2, but relies on Lemma 6.2 and Lemma 6.6. Thus, since the value of q in Lemma 6.6 is related to the quality of the approximation thorough a parameter $\eta>0$ and also depends logarithmically on the number of slots m, we obtain:

Lemma 6.6. *Given $\eta>0$ and $q=\left\lceil\frac{1}{2\eta^2}\log\frac{m+1}{\eta}\right\rceil$ if, for some $\delta>0$, it is the case that $v_i(\theta)>\delta$ for all $i\in\mathcal{N}$ and $\theta\in\Theta$, then $\mathrm{OPT}_{\Xi_q}\geq(1-\eta/\delta)^2\,\mathrm{OPT}$.*

Proof. By Lemma 6.5 given $q=\left\lceil\frac{1}{2\eta^2}\log\frac{m+1}{\eta}\right\rceil$ we know that there exists a distribution $\gamma\in\Delta_{\Xi_q}$ such that:

$$\Pr_{\tilde\xi\sim\gamma_\xi}\Big(\tilde\xi^\top v_{i_j}\geq\xi^\top v_{i_j}-\eta\Big)\geq 1-e^{-2q\eta^2}=1-\frac{\eta}{m+1}\quad\forall j\in[m+1]$$

where $\pi=(i_1,...,i_{m+1})\in\Pi_{m+1}$ is the tuple such that $\xi\in\Xi_\pi$. Thanks to union bound and De Morgan's laws we get:

$$\Pr_{\tilde\xi\sim\gamma_\xi}\Big(\bigcap_{j=1}^{m+1}\{\tilde\xi^\top v_{i_j}\geq\xi^\top v_{i_j}-\varepsilon\}\Big)\geq 1-\eta.$$

Now, we prove that the revenue is $((1-\eta/\delta)^2, 0)$-stable over Ξ_q. In particular, it holds:

$$
\begin{aligned}
\mathbb{E}_{\tilde{\xi}\sim\gamma_\xi}\Big[\textsc{Rev}(\mathcal{V},\tilde{\xi})\Big] &\geq \mathbb{E}_{\tilde{\xi}\sim\gamma_\xi}\Big[\textsc{Rev}(\mathcal{V},\tilde{\xi}) \mid \bigcap_{j=1}^{m+1}\{\tilde{\xi}^\top v_{i_j} \geq \xi^\top v_{i_j} - \eta\}\Big]\Pr\Big(\bigcap_{j=1}^{m+1}\{\tilde{\xi}^\top v_{i_j} \geq \xi^\top v_{i_j} - \eta\}\Big) \\
&\geq (1-\eta)\mathbb{E}_{\tilde{\xi}\sim\gamma_\xi}\Big[\textsc{Rev}(\mathcal{V},\tilde{\xi}) \mid \bigcap_{j=1}^{m+1}\{\tilde{\xi}^\top v_{i_j} \geq \xi^\top v_{i_j} - \eta\}\Big] \\
&\geq (1-\eta)\Big(\sum_{j=1}^{m} j\,(\lambda_j - \lambda_{j+1})(\xi^\top v_{i_{j+1}} - \frac{\eta}{\delta}\xi^\top v_{i_{j+1}})\Big) \\
&\geq (1-\eta)\Big(\Big(1-\frac{\eta}{\delta}\Big)\sum_{j=1}^{m} j\,(\lambda_j - \lambda_{j+1})\xi^\top v_{i_{j+1}}\Big) \\
&\geq \Big(1-\frac{\eta}{\delta}\Big)^2 \textsc{Rev}(\mathcal{V},\xi).
\end{aligned}
$$

The latter inequality shows that the revenue is $((1-\eta/\delta)^2, 0)$-stable over Ξ_q with $q = \lceil \frac{1}{2\eta^2}\log\big(\frac{m+1}{\eta}\big)\rceil$. Thus, by Lemma 6.1, $\text{OPT}_{\Xi_q} \geq (1-\eta/\delta)^2\text{OPT}$ proving the lemma. $\square$

Theorem 6.3. *In the Random Valuations setting, if $v_i(\theta) \geq \delta$ for all $i \in \mathcal{N}$ and $\theta \in \Theta$ for some $\delta > 0$, then the problem of computing an optimal signaling scheme admits a (multiplicative) QPTAS.*

Proof. Let β be the desired approximation. Moreover, let α and η be values defined in the following, $\tau = \nu\delta\lambda_1$, $q = \lceil \frac{1}{2\eta^2}\log\big(\frac{m+1}{\eta}\big)\rceil$, and $s = \lceil \frac{2m^2}{(\delta\nu)^2}\log(\frac{2}{\alpha})\rceil$. By Lemma 6.2, it holds

$$
\begin{aligned}
\text{OPT}_{\Xi_{q,s}} &= \mathbb{E}\Big[\sum_{\xi\in\Xi_q}\gamma_{\mathcal{V}_s}(\xi)\textsc{Rev}(\mathcal{V},\xi)\Big] \\
&\geq \Big(1-\alpha|\Xi_q|\Big)\text{OPT}_{\Xi_q} - \nu\delta\lambda_1 \\
&\geq \Big(1-\nu-\alpha|\Xi_q|\Big)\text{OPT}_{\Xi_q}.
\end{aligned}
$$

By Lemma 6.6 we have:

$$
\begin{aligned}
\text{OPT}_{\Xi_{q,s}} &= \mathbb{E}\Big[\sum_{\xi\in\Xi_q}\gamma_{\mathcal{V}_s}(\xi)\textsc{Rev}(\mathcal{V},\xi)\Big] \\
&\geq \Big(1-\nu-\alpha|\Xi_q|\Big)\text{OPT}_{\Xi_q} \\
&\geq (1-\nu-\alpha|\Xi_q|)\Big(1-\frac{\eta}{\delta}\Big)^2\text{OPT} \\
&= (1-\beta)\text{OPT},
\end{aligned}
$$

where the last inequality holds for $\nu = \beta/8$, $\alpha = \beta/(8|\Xi_q|)$, and $\eta = \delta\beta/4$. $\square$

The following theorem shows that the result is tight.

Theorem 6.4. *Assuming the* ETH, *there exists a constant* $\omega > 0$ *such that finding a signaling scheme that provides an expected revenue at least of* $(1-\omega)$OPT *requires* $I^{\tilde{\Omega}(\log I)}$ *time, where* I *is the size of the problem instance. This holds even when* $v_i(\theta) > \frac{1}{3}$ *for all* $i \in \mathcal{N}$ *and* $\theta \in \Theta$.[3]

Proof. We reduce from public signaling in elections with a k-voting rule. In particular, each receiver $i \in \mathcal{N}$ has an utility $u^i_\theta \in [-1, 1]$ in a state $\theta \in \Theta$, where u^i_θ represent the difference between the utility of receiver $i \in \mathcal{N}$ in voting c_0 with respect to c_1. The sender's utility is 1 if the at least k-voters vote for c_0, *i.e.*, the induced posterior $\xi \in \Delta_\Theta$ is such that $\sum_{\theta\in\Theta} \xi(\theta)u^i_\theta \geq 0$. Otherwise, the sender's utility is 0. See Castiglioni et al. [2020] for a more detailed description of the problem. Castiglioni et al. [2020] show that assuming the ETH[4], there exists a constant $\varepsilon > 0$ such that distinguish between this two cases requires $n^{\tilde{\Omega}(log(n))}$ time:

1. there exists a signaling scheme such that in any induced posterior $\xi \in \Delta_\Theta$ at least k voters have $\sum_{\theta\in\Theta} \xi(\theta)u^i_\theta \geq 0$;
2. in all the posteriors there are strictly less than k receivers with $\sum_{\theta\in\Theta} \xi(\theta)u^i_\theta \geq -\varepsilon$.

To prove the theorem, we show how to reduce the k-voting problem to our revenue maximization problem. In particular, given an instance of k-voting, we build an instance of our problem with the same number of receivers. The valuation of receiver $i \in \mathcal{N}$ in a state $\theta \in \Theta$ is $v_i(\theta) = \frac{u^i_\theta}{3} + \frac{2}{3}$, Moreover, there are $m = k - 1$ slots with $\lambda_j = 1$ for each $j \in [m]$. Finally, we set the required approximation $\omega = \varepsilon/2$. We show that when the first case holds, the revenue is at least $(k-1)\frac{2}{3}$, while in the second case it is strictly less than $(k-1)\frac{2-\varepsilon}{3}$. Hence, a $\frac{2-\varepsilon}{3}/\frac{2}{3} = 1-\varepsilon/2 = 1-\omega$ approximation to the signaling problem can be used to provide an ε-approximation to k-voting. Since we provide a polynomial time reduction from k-voting to the revenue maximization problem, this is sufficient to prove the theorem.

soundness. Suppose that there exists a signaling scheme such that in any induced posterior $\xi \in \Delta_\Theta$ at least k voters have $\sum_{\theta\in\Theta} \xi(\theta)u^i_\theta \geq 0$. Consider the same signaling scheme in the revenue maximization problem. Then, in all the induced posteriors $\xi \in \Delta_\Theta$

[3]The $\tilde{\Omega}$ notation hides poly-logarithmic factors.
[4]Exponential Time Hypothesis.

there are at least k receivers with expected valuation at least $\frac{2}{3}$ and the total revenue is at least $(k-1)\frac{2}{3}$.

completeness. Suppose that in all the posteriors $\xi \in \Delta_\Theta$ there are strictly less than k receivers with $\sum_{\theta\in\Theta} \xi(\theta) u^i_\theta \geq -\varepsilon$. Notice that the revenue of a posterior is given by $(k-1)x$, where x is the smallest expected valuation. Hence, the maximum revenue is strictly less than $(k-1)\frac{2-\varepsilon}{3}$. $\square$

7 | Experimental Results

In this chapter, we provide some experimental results showing how committing to a signaling scheme gives the mechanism higher expected revenue. In particular, we analyze different settings of a Bayesian ad auction, showing how varying a parameter of a given instance can increase or decrease the persuasiveness power of committing to a signaling scheme. In all the experiments we consider that the bidders' valuations matrix $V \in [0,1]^{n \times d}$ is built drawing independently each $V(i,\theta) \sim \text{Be}(p)$ from a Bernoulli distribution of parameter $p \in [0,1]$ for each $\theta \in \Theta$ and $i \in \mathcal{N}$. Once the bidders' valuation matrix is built, we assume that the auction mechanism has access to it, and then it computes an (almost) optimal signaling scheme by solving LP 5.1 instantiated for the set Ξ_q with $q = 30$ (see Definition 6.2). Moreover, to reduce the bias of our analysis, we repeat 30 different times the procedure described above. Finally, for each instance of the problem I, we considered the persuasiveness ratio given by:

$$r(I) = \frac{\mathbb{E}_{\theta \sim \mu, \xi \sim \gamma}[\text{Rev}(\xi, V)]}{\mathbb{E}_{\theta \sim \mu}[\text{Rev}(\mathbb{I}_\theta, V)]},$$

with γ optimal solution of the previously discussed LP. Notice that $r(I) \geq 1$ for each possible instance of the problem I, since the solution in which the mechanism simply reveals the state of nature is feasible for LP 5.1. In addition, the ratio $r(I)$ has an intuitive meaning since the numerator represents the expected revenue prescribed by an optimal solution of the previously discussed LP, while the denominator represents the expected revenue the mechanism would get without committing to it. So, higher values of such a ratio indicate that the persuasiveness power of a signaling scheme is higher for that specific instance. On the contrary, when the ratio is equal to 1, the persuasiveness power of committing to a signaling scheme is negligible.

7.0.1. Varying the Bidders' Interest

In the first scenario we consider an ad auction with $n = 30$ bidders, $m = 4$ slots and $d = 3$ states of nature with a uniform prior distribution $\mu \in \Delta_\Theta$. In addition, we assume that the set of CTRs is given by $\Lambda = \{0.7, 0.6, 0.5, 0.4\}$ and we vary the parameter p in the range between 0.05 and 0.25, which defines the bidders' valuation matrix as previously discussed. The latter parameter has an intuitive meaning since greater values of p are associated to bidders' valuations matrices with a higher expected number of bidders interested in displaying their ad. Figure 7.1 shows that a larger persuasiveness power is associated with lower values of $p \in [0, 1]$, while when $p = 0.25$ the persuasiveness ratio is equal to 1. Due to that, we can conclude that in auctions in which the average interest in displaying an ad is greater, the effect of committing to a signaling scheme is negligible, while in contexts of lower interest, a persuasive mechanism guarantees the auctioneer a higher revenue.

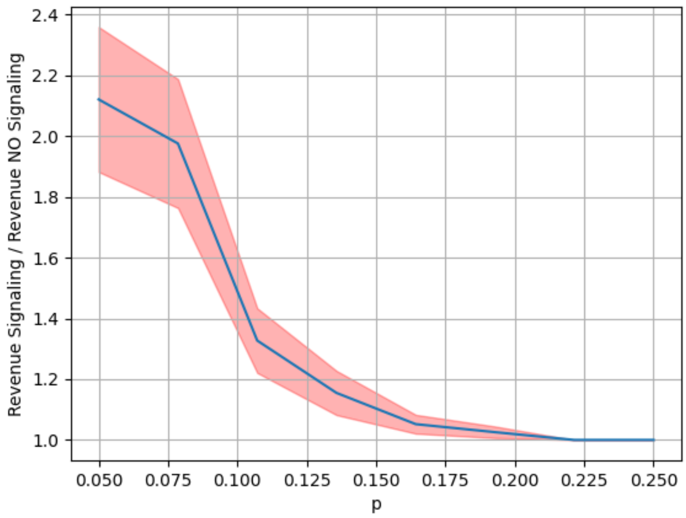

Figure 7.1: Varying the bidders' interest.

7.0.2. Varying the Number of Slots

In the second scenario, we consider an ad auction with $n = 30$ bidders, $m = 4$ slots and $d = 3$ states of nature with a uniform prior distribution $\mu \in \Delta_\Theta$. Moreover we assume that the bidders' valuation matrix is built in a way that $V(i, \theta) \sim \text{Be}(p)$ for each $\theta \in \Theta, i \in \mathcal{N}$ taking $p = 0.1$. Differently from above, we now consider different ad auctions with a number of slots varying from 2 to 10. In particular, for each instance with m items we assume that the set of CTRs is given by $\Lambda = \{2/10, ..., m/10\}$. So that, for example, when the number of slots is equal to $m = 4$ we have that $\Lambda = \{0.1, 0.2, 0.3, 0.4\}$. In Figure 7.2 we report how the persuasiveness ratio varies in auction with a different number of slots. In this case, we observe that, since the number of bidders is fixed, with a progressive increase in the number of slots, the persuasive power of a signaling scheme grows. In particular, in an ad auction with two or three slots, the persuasiveness ratio equals the unitary value, while for an ad auction with ten slots, the ratio is almost equal to 2.

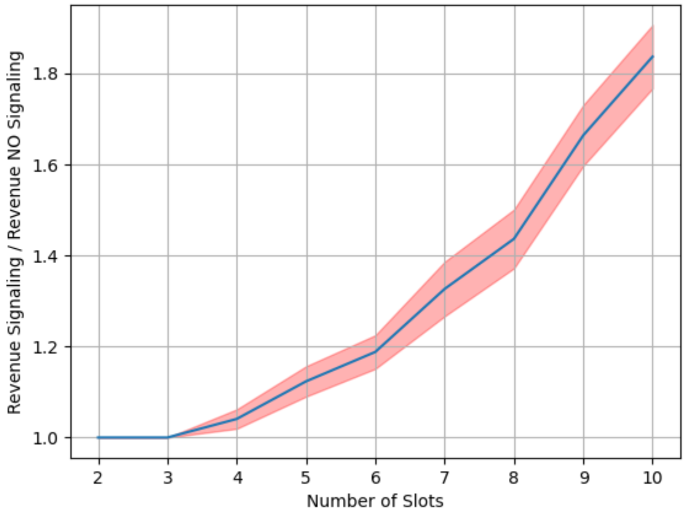

Figure 7.2: Varying the number of slots.

7.0.3. Varying the Number of Bidders

In the third setting, analogously to what we have done so far, we consider ad auctions with $m = 4$ slots and $d = 3$ states of nature with a uniform prior distribution $\mu \in \Delta_\Theta$. Moreover, we assume that the bidders' valuation matrix is built in a way that $V(i, \theta) \sim \text{Be}(p)$ for each $\theta \in \Theta, i \in \mathcal{N}$ taking $p = 0.1$ and we vary the number of bidders participating in the auction from $n = 15$ to $n = 45$. Figure 7.3 shows that the persuasiveness ratio is higher for instances with fewer bidders, while it results lower when the number of bidders increases. Such a behaviour proves that if the number of slots does not change, increasing the number of bidders also increases the probability of finding some interested advertisers in displaying their ad. Specifically, as we can see from Figure 7.3, when we consider an instance with 45 bidders, the persuasiveness ratio is almost equal to 1.2, while when we allow for a setting with only 15 bidders, the persuasiveness ratio equals the value of 2.3.

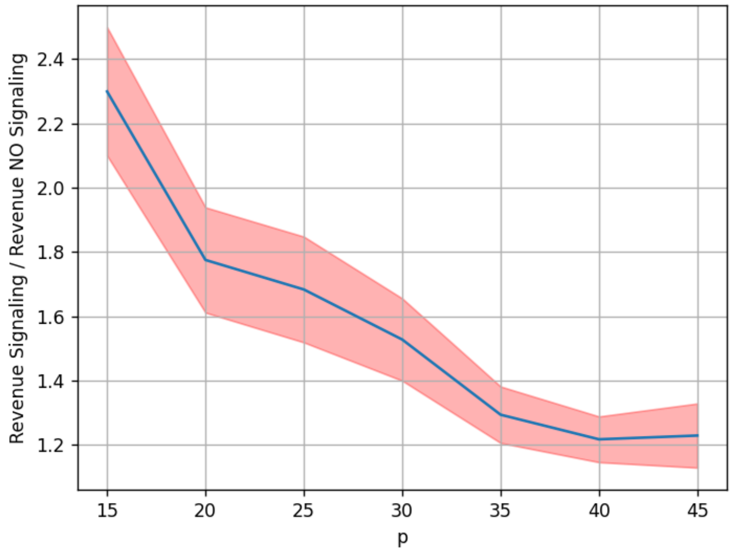

Figure 7.3: Varying the number of bidders.

7.0.4. Economic Interpretation

These experiments show that in some ad auctions, a persuasive mechanism guarantees a higher revenue, while in other cases, the effectiveness of committing to a signaling scheme is almost null. In fact, in instances where the mechanism has receives sufficiently high average evaluations for each state of nature, the persuasive power of a signal scheme is negligible since the solution in which the mechanism simply reveals the state of nature turns out to be the winning one. On the contrary, in situations where the mechanism does not have this guarantee, such as in the case of an ad auction with many slots or few participants, the persuasive power of a public signaling scheme proves to be of fundamental importance. For this reason, in instances where the mechanism has no guarantee to receive sufficiently high valuations, a partial disclosure of the information on the state of nature provides the mechanism with higher revenue since it increases the competition among bidders.

8 | Conclusions and Future Works

This thesis studies how an ad auction mechanism should commit to a public signaling scheme to maximize revenue. To the best of our knowledge, we extended the study of public signaling to auctions with one more slot, which was limited to the study of second-price auctions. In particular, we first proved that there is no PTAS for computing a revenue-maximizing signaling scheme unless P $=$ NP; due to that, we studied several scenarios in which such a negative result could be circumvented. All the proposed algorithms solve suitably-defined linear programs of polynomial size and provide an optimal or an ε-approximation of an optimal revenue-maximizing signaling scheme.

In this work, we focused only on public signaling schemes in which the mechanism draws a signal and publicly reports it to all the bidders in the same way. A new possible research direction could be the study of a private signaling scheme characterized by the fact that the mechanism could send different signals to different receivers. In this scenario, it has been shown that bidders may not report their true valuation (since it does not represent an equilibrium) also in the simpler case of a second-price auction. Moreover, one of the most limiting assumptions in our model is that the auctioneer is required to have access to the bidders' valuation distribution through a black-box sampling oracle. Due to that, a new possible research track could relax this assumption through an online learning framework, as first proposed by Castiglioni et al. [2020], in which the auctioneer repeatedly faces bidders whose valuations are unknown and chosen adversarially. This study should focus on no-regret algorithms prescribing a signaling scheme at each round close to the best-in-hindsight one. Finally, also the assumption that bidders may not share information could be relaxed. Indeed, if bidders are allowed to communicate, they could agree to report a lower value to the mechanism getting a better utility. In this setting, it is possible to study how an ad auction mechanism should commit to a signaling scheme depending on the possible coalitions formed by bidders.

Bibliography

[1] P. Alimonti and V. Kann. Some apx-completeness results for cubic graphs. *Theoretical Computer Science*, 237:123–134, 04 2000. doi: 10.1016/S0304-3975(98)00158-3.

[2] R. Alonso and O. Câmara. Persuading voters. *American Economic Review*, 106(11): 3590–3605, 2016.

[3] Y. Babichenko and S. Barman. Algorithmic aspects of private Bayesian persuasion. In *Innovations in Theoretical Computer Science Conference*, 2017.

[4] F. Bacchiocchi, M. Castiglioni, A. Marchesi, G. Romano, and N. Gatti. Public signaling in bayesian ad auctions. *CoRR*, abs/2201.09728, 2022. URL `https://arxiv.org/abs/2201.09728`.

[5] A. Badanidiyuru, K. Bhawalkar, and H. Xu. Targeting and signaling in ad auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2545–2563. SIAM, 2018.

[6] U. Bhaskar, Y. Cheng, Y. K. Ko, and C. Swamy. Hardness results for signaling in bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 479–496, 2016.

[7] P. Bro Miltersen and O. Sheffet. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 234–247, 2012.

[8] O. Candogan. Persuasion in networks: Public signals and k-cores. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 133–134, 2019.

[9] M. Castiglioni and N. Gatti. Persuading voters in district-based elections. In *The AAAI Conference on Artificial Intelligence, AAAI*. AAAI Press, 2021.

[10] M. Castiglioni, A. Celli, and N. Gatti. Persuading voters: It's easy to whisper, it's hard to speak loud. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 1870–1877, 2020.

[11] M. Castiglioni, A. Celli, and N. Gatti. Public bayesian persuasion: Being almost optimal and almost persuasive. *CoRR*, abs/2002.05156, 2020. URL `https://arxiv.org/abs/2002.05156`.

[12] M. Castiglioni, A. Celli, A. Marchesi, and N. Gatti. Online bayesian persuasion. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16188–16198. Curran Associates, Inc., 2020. URL `https://proceedings.neurips.cc/paper/2020/file/ba5451d3c91a0f982f103cdbe249bc78-Paper.pdf`.

[13] M. Castiglioni, A. Celli, A. Marchesi, and N. Gatti. Signaling in bayesian network congestion games: the subtle power of symmetry. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence*, 2021.

[14] M. Castiglioni, G. Romano, A. Marchesi, and N. Gatti. Signaling in posted price auctions. *CoRR*, abs/2201.12183, 2022. URL `https://arxiv.org/abs/2201.12183`.

[15] Y. Cheng, H. Y. Cheung, S. Dughmi, E. Emamjomeh-Zadeh, L. Han, and S.-H. Teng. Mixture selection, mechanism design, and signaling. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1426–1445. IEEE, 2015.

[16] eMarketer. Worldwide digital ad spending year-end update. `https://www.emarketer.com/content/worldwide-digital-ad-spending-year-end-update`, 2021. Accessed: 2022-01-09.

[17] Y. Emek, M. Feldman, I. Gamzu, R. PaesLeme, and M. Tennenholtz. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation (TEAC)*, 2(2):1–19, 2014.

[18] E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[19] Y. Mansour, A. Slivkins, V. Syrgkanis, and Z. S. Wu. Bayesian exploration: Incentivizing exploration in bayesian games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 661–661, 2016.

[20] Z. Rabinovich, A. X. Jiang, M. Jain, and H. Xu. Information disclosure as a means to security. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 645–653, 2015.

[21] S. Vasserman, M. Feldman, and A. Hassidim. Implementing the wisdom of waze. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, pages 660–666, 2015.

[22] H. Xu. On the tractability of public persuasion with no externalities. In S. Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2708–2727. SIAM, 2020. doi: 10.1137/1.9781611975994.165. URL `https://doi.org/10.1137/1.9781611975994.165`.

[23] H. Xu, R. Freeman, V. Conitzer, S. Dughmi, and M. Tambe. Signaling in bayesian stackelberg games. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems*, pages 150–158, 2016.