

Signaling Schemes for Revenue Maximization

YUVAL EMEK, ETH Zurich

MICHAL FELDMAN, Hebrew University and Harvard University

IFTAH GAMZU, Yahoo Research

RENATO PAES LEME, Microsoft Research

MOSHE TENNENHOLTZ, Microsoft Research and Technion

Signaling is an important topic in the study of asymmetric information in economic settings. In particular, the transparency of information available to a seller in an auction setting is a question of major interest. We introduce the study of signaling when conducting a second price auction of a probabilistic good whose actual instantiation is known to the auctioneer but not to the bidders. This framework can be used to model impressions selling in display advertising. We establish several results within this framework. First, we study the problem of computing a signaling scheme that maximizes the auctioneer's revenue in a Bayesian setting. We show that this problem is polynomially solvable for some interesting special cases, but computationally hard in general. Second, we establish a tight bound on the minimum number of signals required to implement an optimal signaling scheme. Finally, we show that at least half of the maximum social welfare can be preserved within such a scheme.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—Complexity of proof procedures; J.4 [Social and Behavioral Sciences]: Economics

General Terms: Algorithms, Economics

Additional Key Words and Phrases: Signaling, asymmetric information, probabilistic auctions

ACM Reference Format:

Emek, Y., Feldman, M., Gamzu, I., Paes Leme, R., and Tennenholtz, M. 2014. Signaling schemes for revenue maximization. *ACM Trans. Econ. Comp.* 2, 2, Article 5 (June 2014), 19 pages.

DOI: <http://dx.doi.org/10.1145/2594564>

1. INTRODUCTION

A major concern in market design is to ensure that the markets are *thick* in the sense that there is a sufficient volume of participants to produce the necessary level of competition for the market to work well. Another concern is to design a practical language that is sufficiently expressive to allow players to specify how much they value the goods in the market. In the market for diamonds described by Levin and Milgrom [2010] for example, the auctioneers could elicit bids for each individual stone. However, the enormous effort required for the players to learn the value of each individual stone and to submit individual bids would make the auction impractical. Moreover, bidding on each stone separately can lead to the *cherry-picking* phenomenon, where very few customers are interested in any one stone. This may lead to a situation where little revenue is

Michal Feldman's work was partially supported by the Israel Science Foundation (grant number 1219/09), by the Leon Recanati Fund of the Jerusalem School of Business Administration, the Google Inter-university center for Electronic Markets and Auctions, and the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 274919. Renato Paes Leme's work was partially supported by NSF grants CCF-0729006 and by a Microsoft Research Fellowship.

Author's address: Michal Feldman; email: michal.feldman@cs.tau.ac.il.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

© 2014 held by the Owner/Author. Publication rights licensed to ACM. 2167-8375/2014/06-ART5 \$15.00

DOI: <http://dx.doi.org/10.1145/2594564>

generated although the goods are valuable. In practice, stones are categorized into *deals* and then auctioned. This method of treating different goods as identical is called *conflation*.

Milgrom [2010] and Levin and Milgrom [2010] provide a comprehensive analysis of the phenomenon of conflation in various markets, with particular emphasis on online advertisement. In these markets, auctioning each good individually is usually not an option and conflation must be used. One particular online market that we will focus on is the multibillion-dollar display advertisement market, where publishers (such as MSN and Yahoo) attempt to maximize the revenue they collect from the advertisers (say, Nike or Coca-Cola) for wisely targeting their ads at the right users. For example, an ad referring to the surfing lifestyle on the sunny beaches of the Pacific Ocean may be most valuable when targeted at a teenager from California; perhaps less so when targeted at a 10 year old from Oregon; and even less when targeted at older folks in areas that are far from the ocean. However, it would be impossible for advertisers to decide how to bid on each individual impression. Instead, the impressions are categorized based on attributes such as the time when the impression was made, cookies in the user's browser, certain demographic properties, geographic location, and so on; impressions with similar attributes are then treated as instances of the same good.

In the context of display advertisement, the main question we deal with is: how should those impressions be categorized in order to maximize the publisher's revenue? The high-level idea behind our model is to explore the natural asymmetry of information between the publisher and the advertisers: while advertisers may know the distribution of users visiting a particular site, the publisher usually has much more accurate information about each individual impression.¹ Upon receiving an impression, the publisher may choose to reveal to the advertisers certain attributes of this impression (say, age and gender), while concealing other attributes (say, geographic location). One may argue that concealing information from the advertisers might generate inefficiencies in the market, but the amount of information is typically so large that it would be impossible for the advertisers to grasp everything anyhow. More importantly, ensuring that revenue is generated is essential for the proper functioning of markets. As described by Muthukrishnan [2009], in Ad Exchanges, which are systems that bring together publishers and advertisers in a common marketplace, ensuring good revenue is vital to keeping publishers in the market.

More generally, our goal in this article is to cope with the undesired effects of competition deficiency on some items in an auction. To achieve this goal, we will exploit an inherent information asymmetry between the auctioneer and the bidders that exists in many market settings. We model² the auctioneer/bidders asymmetry by considering a framework termed a *probabilistic single-item auction*, in which n bidders participate in an auction for a single item, which is chosen randomly from a set of m indivisible goods according to a commonly known probability distribution $p \in \Delta(m)$. In contrast to the bidders, who know only the probability distribution over the possible goods, the auctioneer knows its actual realization, and can use this informational superiority to increase the collected revenue.

Specifically, the auctioneer may choose to reveal partial information to the bidders by means of a *signaling scheme*. A signaling scheme is a (possibly randomized) policy that specifies some signal σ revealed to the bidders upon the choice (made by nature) of each good $j \in [m]$. This policy is known to the bidders who can therefore induce the

¹In reality, the additional information about the visitor to a site is often handled by third party demand side platform (usually referred as DSP). For simplicity, we abstract away this distinction.

²For the formal exposition of our model, see Section 2.

revealed signal σ to update their perceived probability for the chosen good j from $p(j)$ to the “more accurate” $p(j | \sigma)$.

One approach would be design an optimal auction from scratch for the problem of maximizing revenue in a probabilistic single item auction setting. Since we are in a setting with very correlated values, this is likely to produce contorted and impractical auctions in the style of the auction of Cremer and McLean [1985, 1988], where full surplus extraction is possible. Instead, we focus on the standard second price auction mechanism which is the de facto standard for the sale of online advertisement [Edelman et al. 2007; Muthukrishnan 2009; Varian 2007]. We believe this will generate an auction that is more relevant to practical applications and that can be easily integrated with the current implementations.

In this auction, after the bidders receive the signal, they submit their bids, and the winner and the payment are determined according to the second-price auction; namely, the winner is the bidder with the highest bid and the payment is the second-highest bid. The goal of the auctioneer, which is the subject of this article, is to design a signaling scheme that maximizes her expected revenue.

A simple but crucial observation that facilitates our analysis is that, similar to the classical setting of second-price auctions, here too, it is a dominant strategy for the bidders to reveal their true expected valuations, where the expectation in this context is taken with respect to the conditional probability $p(j | \sigma)$. Therefore, the problem, termed *revenue maximization by signaling*, reduces to finding a signaling scheme that maximizes the expected second-highest bid (amounting to the expected revenue).

Two trivial signaling schemes are the one that reveals no information to the bidders and the one that reveals the actual realization (all the information). Interestingly, there are instances in which an appropriate signaling scheme provides a substantial improvement over the two trivial ones. This can be demonstrated through a special case of a signaling scheme, termed *clustering*: the auctioneer a priori partitions the set of goods into disjoint clusters, and the signal is the cluster that contains the chosen good. Consider the case in which there are m bidders and m types of goods, an item is chosen uniformly at random, and each bidder i is only interested in good i with a unit valuation. If no information is revealed then the expected revenue is $1/m$ as the expected valuation of each bidder is $1/m$. If the actual realization is revealed, no revenue is collected since for every realization, the second-highest valuation is 0. However, if the goods are partitioned into clusters of size 2, then the expected revenue is $1/2$, providing an improvement of a linear factor over the best trivial scheme.

Note that clustering schemes can be thought of as restricting the auctioneer to deterministic policies. The class of signaling schemes considered in this article is more general than clustering, as we allow the auctioneer to toss coins when deciding on the revealed signal.

1.1. Our Results

We begin our analysis assuming that the valuations of the bidders are known to the auctioneer. In this somewhat less realistic case, the problem of revenue maximization by signaling can be formalized as a concise linear program and, as such, solved to optimality in polynomial time. A natural question is to what extent the increase in revenue comes at the expense of social welfare. Notably, we prove that a signaling scheme that obtains the optimal revenue can preserve at least half of the optimal social welfare. In addition, it is shown that if the auctioneer is restricted to invoking a signaling scheme by means of clustering, at least half of the optimal revenue can be achieved, and this is tight.

Up until now we have assumed that the valuations of the bidders are known to the auctioneer. However, in practice, the auctioneer rarely knows the bidders' valuations.

This motivates the main technical contribution of this article, namely, the study of a Bayesian setting, in which the auctioneer holds probabilistic knowledge on the bidders' valuations. We show that in this case the revenue maximization by signaling problem becomes NP-hard. Still, in several cases of interest the problem remains tractable even in the Bayesian setting. Finally, we show that m signals are always sufficient to extract the optimal revenue. It is an interesting open problem how to find good approximation algorithms for the cases where the revenue maximization is NP-hard or to prove hardness of approximation in those cases.

Notice that our model captures the Bayesian knowledge on behalf of the auctioneer by assuming a probability distribution over finitely many valuation matrices. This representation can capture complicated dependencies between the different valuations, however it may exhibit plenty of redundancy when the valuations are assumed to be independent. As such, it will be interesting to study our framework under more concise representations—for example, where each entry of the matrix is sampled independently from some distribution.

1.2. Related Work

There is a rich theory on markets with information asymmetry. In such markets, agents on one side have more (or better) information than those on the other side. The foundation of this theory dates back to the work of Akerlof, Spence, and Stiglitz on the analysis of markets with asymmetric information, which earned them the 2001 Nobel Prize. In particular, Akerlof [1970] introduced the first formal analysis of markets in which sellers have more information than buyers regarding the quality of products. Spence [1973, 2002] demonstrated that in certain settings, well-informed agents can improve their outcome by signaling their private information to poorly informed agents.

There is also a vast literature on the nature and effects of information revelation in auctions. One of the most fundamental results in auction theory, namely the “Linkage Principle” of Milgrom and Weber [1982], states that the expected revenue of an auctioneer is enhanced when bidders are provided with more information. While this work advocates transparency in various markets, later work observed that such transparency may not be optimal in general (see, e.g., Perry and Reny [1999], Weber [2003], Thierry and Stefano [2003], and Feinberg and Tennenholtz [2005]). More recent work [Levin and Milgrom 2010; Milgrom 2010] advocated the need for careful grouping of goods as an important market design principle. Our work may be viewed as a study of information revelation through an optimization lens, since we seek to maximize the expected revenue of an auctioneer by designing an effective information revelation scheme.

Notice that Myerson's classic result on revenue maximization [Myerson 1981] does not apply to our model due to the asymmetry of information. Also, Myerson's mechanism works only for single parameter settings. Our Bayesian models are multiparameter and are typically highly correlated. Revenue maximization results for correlated valuations [Papadimitriou and Pierrakos 2011; Dobzinski et al. 2011] also do not apply here due to asymmetric information. One could try to reveal all information and then apply one of those mechanisms, but we would get a rather contorted auction, with no guarantees, as contrasted with our auction. In fact, it is easy to construct examples where this generates arbitrary less revenue than our signaling scheme. Our auction, on the other hand, is very practical and close to what is actually implemented in online advertising markets.

Closer to our work, is that of Ghosh et al. [2007], which studies revenue-maximizing clustering schemes under a second-price auction in a setting with full information and additive valuations. While this setting is different from our framework of signaling in a

probabilistic item auction, the mathematical formulation of the optimization problem in their setting is a special case of our optimization problem: the case where the valuation matrix is known to the auctioneer and the signaling scheme is restricted to take the form of a clustering scheme. Our focus, though, is on the more realistic Bayesian case, which is not treated in Ghosh et al. [2007]. In addition, our framework relies on signaling that can be viewed as a fractional clustering, which is more powerful. Indeed, while Ghosh et al. show that it is strongly NP-hard to compute the optimal clustering scheme, an optimal signaling scheme can be computed in polynomial time. Results of the similar flavour of the ones in Ghosh et al. [2007] were rederived independently in a previous version of the current article [Emek et al. 2011].

Independently of our work, Miltersen and Sheffet [2012] also analyze the problem of obtaining optimal signaling schemes for revenue maximization using linear programming, obtaining a result similar to our Theorem 3.2.

2. THE MODEL

In this section, we introduce the auctioning model on top of which our signaling schemes are defined. Our focus in this article is on a Bayesian setting, treating the uncertainty of the auctioneer regarding the bidders' valuations in a probabilistic manner. For clarity of the exposition, we shall first consider the less realistic *known-valuations* setting, where no such uncertainty is assumed.

2.1. Known-Valuations Probabilistic Single-Item Auctions

A *known-valuations probabilistic single-item auction (KPSA)* \mathcal{A} is formally depicted by the four-tuple

$$\mathcal{A} = \langle n, m, p, V \rangle ,$$

where $n \in \mathbb{Z}_{>0}$ stands for the number of *bidders*, $m \in \mathbb{Z}_{>0}$ stands for the number of distinct indivisible *goods*, $p \in \Delta(m)$ is a probability distribution over the goods, and $V \in \mathbb{R}_{\geq 0}^{n \times m}$ is a non-negative real matrix capturing the *valuation* $V(i, j)$ of bidder i for good j . A single good $j \in [m]$ is chosen by nature according to the distribution p , which is a common knowledge.

The auction is conducted according to the *second-price* rule: each player i places her bid b_i and the chosen good j is sold to the bidder that placed the highest bid, $\max_{i \in [n]} \{b_i\}$, (ties are broken arbitrarily) for the price of the second highest bid, $\max_{2i \in [n]} \{b_i\}$.

2.2. Signaling Schemes

Although the bidders know the distribution p , they do not know its actual realization, which is observed only by the auctioneer. In an attempt to increase her expected revenue, the auctioneer may partially reveal the realization $j \in [m]$ of p to the bidders. This partial revelation is carried out by means of *signaling*: given that the chosen good is j (recall that this choice is made by nature), the auctioneer sends the bidders some signal σ ; the bidders then hold a “more accurate picture” of the chosen good that corresponds to the probability distribution p conditioned on σ . The policy that dictates the signal that the auctioneer reveals to the bidders for each good $j \in [m]$ is referred to as a *signaling scheme*. It is important to point out that this policy is decided by the auctioneer and reported to the bidders prior to nature's random choice of item j .

More formally, a signaling scheme is given by a set of $s \in \mathbb{Z}_{>0}$ *signals* and a *signaling function* $S : [s] \times [m] \rightarrow [0, 1]$, which satisfies

$$\sum_{\sigma \in [s]} S(\sigma, j) = 1 \quad \forall j \in [m] . \tag{1}$$

Given that nature chose good $j \in [m]$, the auctioneer reveals signal $\sigma \in [s]$ to the bidders with probability $\mathcal{S}(\sigma, j)$. It will be convenient to use the notation \mathcal{S} to address the signaling scheme as well as its inherent signaling function.

Once again, it is assumed that \mathcal{S} (and s) are decided by the auctioneer and reported to the bidders prior to the random choice of j ; it is the actual signal σ (determined according to \mathcal{S}) that is revealed to the bidders after the choice of j . This can be thought of as a commitment of the auctioneer to stick to the signaling scheme that it previously reported. Upon receiving signal σ , the bidders, knowing p and \mathcal{S} , update their belief from $\mathbb{P}(\text{chosen good is } j) = p(j)$ to

$$\begin{aligned} \mathbb{P}(\text{chosen good is } j \mid \text{signal is } \sigma) &= \frac{\mathbb{P}(\text{signal is } \sigma \mid \text{chosen good is } j) \cdot \mathbb{P}(\text{chosen good is } j)}{\mathbb{P}(\text{signal is } \sigma)} \\ &= \frac{\mathcal{S}(\sigma, j) \cdot p(j)}{\sum_{j' \in [m]} \mathcal{S}(\sigma, j') \cdot p(j')} . \end{aligned}$$

Example 2.1. Let the set of goods be composed of {squash, spinach, banana, grapes} occurring with respective probabilities $\frac{1}{3}, \frac{1}{3}, \frac{1}{6}, \frac{1}{6}$. Let the set of signals be {fruit, vegetable, green, yellow}, and consider the signaling function given by the following table.

	fruit	vegetable	green	yellow
squash	0	.5	0	.5
spinach	0	.4	.6	0
banana	.1	0	0	.9
grapes	.5	0	.5	0

Suppose the realized good is $j = \text{spinach}$; then, the auctioneer releases the signal “vegetable” with probability 0.4 and the signal “green” with probability 0.6. If the released signal was $\sigma = \text{“green”}$, then the agents perform the Bayesian updates depicted in the preceding and will form the following beliefs accordingly:

$$\begin{aligned} \mathbb{P}(\text{spinach} \mid \text{green}) &= \frac{0.6 \cdot (1/3)}{0.6 \cdot (1/3) + 0.5 \cdot (1/6)} = \frac{12}{17} \\ \mathbb{P}(\text{grapes} \mid \text{green}) &= \frac{0.5 \cdot (1/6)}{0.6 \cdot (1/3) + 0.5 \cdot (1/6)} = \frac{5}{17} . \end{aligned}$$

For succinctness, we will subsequently denote the events “chosen good is j ” and “signal is σ ” by j and σ , respectively (our intention will be clear from the context).

Before we proceed, let us consider the restricted variant of a *deterministic auctioneer*, which is not allowed to use randomness when determining which signal to reveal. This is equivalent to imposing an additional integrality requirement on the signaling scheme: $\mathcal{S}(\sigma, j) \in \{0, 1\}$ for every $\sigma \in [s]$ and $j \in [m]$. In other words, each signal $\sigma \in [s]$ now corresponds to a *cluster* $C_\sigma \subseteq [m]$ so that the clusters are pairwise disjoint and $\bigcup_{\sigma \in [s]} C_\sigma = [m]$. Following this view, the general case (under which the auctioneer may use randomness when determining the signal) can be interpreted as a fractional clustering of the goods, where $\mathcal{S}(\sigma, j)$ is the fraction of good j in cluster C_σ .

It is well known that in the classical setting of second-price single-item auctions, it is a dominant strategy for the bidders to be truthful: to bid their true valuations [Vickrey 1961]. It turns out that this remains valid in probabilistic single-item auctions under signaling as well, as the following observation demonstrates.

OBSERVATION 2.2. For every $i \in [n]$ and $\sigma \in [s]$, bidding $b_i(\sigma) = \mathbb{E}[V(i,j) \mid \sigma]$ in response to the signal σ is a dominant strategy for bidder i .

PROOF. Consider the known-valuations ex-ante game defined by setting the strategy space of each bidder $i \in [n]$ to be the collection of all possible functions $b_i : [s] \rightarrow \mathbb{R}_{\geq 0}$ and the utility of bidder i from strategy profile $b = (b_1, \dots, b_n)$ to be the expected utility of bidder i in the KPSA \mathcal{A} assuming that each bidder adheres to b . Fix some arbitrary strategies $b_{i'} : [s] \rightarrow \mathbb{R}_{\geq 0}$ for all bidders $i' \neq i$ and consider the strategy b_i of bidder i that bids

$$b_i(\sigma) = \mathbb{E}[V(i,j) \mid \sigma] = \sum_{j \in [m]} \mathbb{P}(j \mid \sigma) \cdot V(i,j) = \sum_{j \in [m]} \frac{\mathcal{S}(\sigma,j) \cdot p(j)}{\sum_{j' \in [m]} \mathcal{S}(\sigma,j') \cdot p(j')} \cdot V(i,j)$$

in response to each signal $\sigma \in [s]$.

Fix some signal σ . From bidder i 's point of view, the expected valuation of the chosen good is $b_i(\sigma)$, whereas each other bidder $i' \neq i$, bids $b_{i'}(\sigma)$. If bidder i does not win the chosen good, which happens only if $\max_{i' \neq i} \{b_{i'}(\sigma)\} \geq b_i(\sigma)$, then her utility in the ex-ante game is 0. This can be changed only if bidder i increases her bid so that it exceeds $\max_{i' \neq i} \{b_{i'}(\sigma)\}$, but this imposes a negative utility on i . So, assume that $\max_{i' \neq i} \{b_{i'}(\sigma)\} \leq b_i(\sigma)$ and that bidder i does win the chosen good. By the definition of the second-price rule, the utility of i must be non-negative. Clearly, bidder i has no incentive to increase her bid. Decreasing her bid does not change her utility as long as it still exceeds $\max_{i' \neq i} \{b_{i'}(\sigma)\}$; decreasing her bid further resets her utility to zero. The assertion follows. \square

2.3. Optimization Problems

Consider some KPSA $\mathcal{A} = \langle n, m, p, V \rangle$ and signaling scheme \mathcal{S} . In light of Observation 2.2, we subsequently assume that the bidders are indeed truthful, that is, bidder i bids $\mathbb{E}[V(i,j) \mid \sigma] = \sum_{j \in [m]} \mathbb{P}(j \mid \sigma) \cdot V(i,j)$ in response to the signal σ . Therefore, the expected *revenue* of the auctioneer, denoted $\mathbf{Rev}_{\mathcal{A}}(\mathcal{S})$, is given by

$$\mathbf{Rev}_{\mathcal{A}}(\mathcal{S}) = \sum_{\sigma \in [s]} \mathbb{P}(\sigma) \cdot \max_{i \in [n]} \left\{ \sum_{j \in [m]} \mathbb{P}(j \mid \sigma) \cdot V(i,j) \right\}.$$

This raises the following optimization problem, referred to as the *revenue maximization by signaling (RMS) problem*: given a KPSA \mathcal{A} , construct the signaling scheme \mathcal{S} that maximizes $\mathbf{Rev}_{\mathcal{A}}(\mathcal{S})$. One may also be interested in the (expected) *social welfare* resulting from signaling scheme \mathcal{S} , defined as

$$\mathbf{SW}_{\mathcal{A}}(\mathcal{S}) = \sum_{\sigma \in [s]} \mathbb{P}(\sigma) \cdot \max_{i \in [n]} \left\{ \sum_{j \in [m]} \mathbb{P}(j \mid \sigma) \cdot V(i,j) \right\}.$$

When \mathcal{A} is clear from the context, we may omit it from the subscript and write simply $\mathbf{Rev}(\mathcal{S})$ and $\mathbf{SW}(\mathcal{S})$.

Notice that the revenue of the auctioneer can be rewritten as

$$\begin{aligned}
\mathbf{Rev}(S) &= \sum_{\sigma \in [s]} \mathbb{P}(\sigma) \cdot \max_{i \in [n]} \left\{ \sum_{j \in [m]} \frac{\mathbb{P}(\sigma | j) \cdot \mathbb{P}(j)}{\mathbb{P}(\sigma)} \cdot V(i, j) \right\} \\
&= \sum_{\sigma \in [s]} \max_{i \in [n]} \left\{ \sum_{j \in [m]} S(\sigma, j) \cdot p(j) \cdot V(i, j) \right\} \\
&= \sum_{\sigma \in [s]} \max_{i \in [n]} \left\{ \sum_{j \in [m]} S(\sigma, j) \cdot \Psi(i, j) \right\}, \tag{2}
\end{aligned}$$

where $\Psi(i, j) = p(j) \cdot V(i, j)$ is referred to as the *normalized valuation* of bidder i for item j . Following the same line of arguments, we can also rewrite the social welfare as

$$\mathbf{SW}(S) = \sum_{\sigma \in [s]} \max_{i \in [n]} \left\{ \sum_{j \in [m]} S(\sigma, j) \cdot \Psi(i, j) \right\}.$$

Under the deterministic auctioneer requirement, the RMS problem turns into the following clustering problem. Given the normalized valuation matrix $\Psi \in \mathbb{R}_{\geq 0}^{n \times m}$, devise a pairwise disjoint partition of $[m]$ into clusters $\{C_\sigma\}_{\sigma \in [s]}$ that maximizes

$$\mathbf{Rev}(\{C_\sigma\}_{\sigma \in [s]}) = \sum_{\sigma \in [s]} \max_{i \in [n]} \left\{ \sum_{j \in C_\sigma} \Psi(i, j) \right\}.$$

2.4. A Bayesian Setting

Recall that up until now, we assumed that the valuations of the bidders are known to the auctioneer.³ However, in many practical scenarios the auctioneer does not know the exact valuation of each bidder. To tackle this obstacle, we assume a *Bayesian* setting, treating the state of knowledge that the auctioneer holds on the bidders' valuations in a probabilistic manner. This is captured in our model by considering $k \in \mathbb{Z}_{>0}$ distinct valuation matrices $V_1, \dots, V_k \in \mathbb{R}_{\geq 0}^{n \times m}$ and a probability distribution $q \in \Delta(k)$ associating each valuation matrix V_ℓ with the probability $q(\ell)$ that it occurs. A *probabilistic single-item auction* (PSA) is then depicted by the 6-tuple

$$\mathcal{A} = \langle n, m, k, p, q, \{V_\ell\}_{\ell \in [k]} \rangle,$$

where $n \in \mathbb{Z}_{>0}$, $m \in \mathbb{Z}_{>0}$, and $p \in \Delta(m)$ have the same role as in the known-valuations case; and $k \in \mathbb{Z}_{>0}$, $q \in \Delta(k)$, and $\{V_\ell \in \mathbb{R}_{\geq 0}^{n \times m}\}_{\ell \in [k]}$ capture the aforementioned Bayesian angle.

³In some sense, we also assumed that the valuations of each bidder are known to the other bidders. However, Observation 2.2 implies that this does not matter: a bidder is better off bidding its true (expected) valuation regardless of the strategies of the other bidders.

The expected revenue of the auctioneer from the signaling scheme \mathcal{S} is now defined to be

$$\begin{aligned} \mathbf{Rev}_{\mathcal{A}}(\mathcal{S}) &= \sum_{\ell \in [k]} q(\ell) \sum_{\sigma \in [s]} \mathbb{P}(\sigma) \cdot \max_{i \in [n]} \left\{ \sum_j \mathbb{P}(j \mid \sigma) \cdot V_{\ell}(i, j) \right\} \\ &= \sum_{\ell \in [k]} q(\ell) \sum_{\sigma \in [s]} \max_{i \in [n]} \left\{ \sum_j \mathcal{S}(\sigma, j) \cdot \Psi_{\ell}(i, j) \right\}, \end{aligned}$$

where $\Psi_{\ell}(i, j) = p(j) \cdot V_{\ell}(i, j)$. The last equation follows from the same line of arguments that was used to establish (2).

3. OPTIMAL SIGNALING SCHEMES—THE KNOWN-VALUATIONS CASE

Let us start our technical treatment of signaling schemes with the simpler known-valuations setting, considering a KPSA $\mathcal{A} = \langle n, m, p, V \rangle$. We show that an optimal signaling scheme for \mathcal{A} can be obtained by solving an LP with $O(n^2m)$ variables and $O(n^2 + m)$ constraints (excluding the non-negativity constraints).

Given an s -signal signaling scheme \mathcal{S} for \mathcal{A} and a signal $\sigma \in [s]$, let $h_1^{\mathcal{S}}(\sigma)$ and $h_2^{\mathcal{S}}(\sigma)$ denote the bidders i that realize $\max_{i \in [n]} \left\{ \sum_{j \in [m]} \mathcal{S}(\sigma, j) \cdot \Psi(i, j) \right\}$ and $\max_{i \in [n]} \left\{ \sum_{j \in [m]} \mathcal{S}(\sigma, j) \cdot \Psi(i, j) \right\}$, respectively. (When the signaling scheme \mathcal{S} is clear from the context, we may omit the superscripts.) Our concise LP relies on the following observation.

OBSERVATION 3.1. *There exists an optimal s -signal signaling scheme \mathcal{S} for \mathcal{A} such that given $\sigma, \sigma' \in [s]$, if $h_1^{\mathcal{S}}(\sigma) = h_1^{\mathcal{S}}(\sigma')$ and $h_2^{\mathcal{S}}(\sigma) = h_2^{\mathcal{S}}(\sigma')$, then $\sigma = \sigma'$.*

PROOF. Consider an optimal s -signal signaling scheme \mathcal{S} that minimizes s . We argue that \mathcal{S} must satisfy the assertion. To that end, assume by contradiction that there are two distinct signals $\sigma, \sigma' \in [s]$ such that $h_1^{\mathcal{S}}(\sigma) = h_1^{\mathcal{S}}(\sigma') = i_1$ and $h_2^{\mathcal{S}}(\sigma) = h_2^{\mathcal{S}}(\sigma') = i_2$. Let \mathcal{S}^* be the $(s - 1)$ -signal signaling scheme obtained from \mathcal{S} by replacing both signals σ and σ' by a new signal σ^* defined by setting $\mathcal{S}(\sigma^*, j) = \mathcal{S}(\sigma, j) + \mathcal{S}(\sigma', j)$ for every $j \in [m]$.

It is easy to verify that \mathcal{S}^* is valid in terms of (1). Moreover, since $h_1^{\mathcal{S}^*}(\sigma^*) = i_1$ and $h_2^{\mathcal{S}^*}(\sigma^*) = i_2$, we can use (2) to conclude that the combined contributions of σ and σ' to $\mathbf{Rev}(\mathcal{S})$ are

$$\sum_{j \in [m]} (\mathcal{S}(\sigma, j) + \mathcal{S}(\sigma', j)) \cdot \Psi(i_2, j) = \sum_{j \in [m]} \mathcal{S}^*(\sigma^*, j) \cdot \Psi(i_2, j),$$

which is precisely the contribution of σ^* to $\mathbf{Rev}(\mathcal{S}^*)$. Thus, $\mathbf{Rev}(\mathcal{S}) = \mathbf{Rev}(\mathcal{S}^*)$, in contradiction to the minimality of s . \square

A direct corollary of Observation 3.1 is that it suffices to consider signaling schemes with $s = n(n - 1)$ signals—each signal σ is uniquely identified by $h_1(\sigma)$ and $h_2(\sigma)$. This turns out to be asymptotically tight, as there are examples showing that $\Omega(n^2)$ signals are required to implement an optimal signaling scheme (see Section 5). Based on

Observation 3.1 and on the formulation of revenue in (2), we can construct an optimal signaling scheme \mathcal{S} by solving the following linear program, denoted **LP1**:

$$\begin{aligned}
\max \quad & \sum_{i_1, i_2 \in [n], i_1 \neq i_2} R(\sigma_{i_1, i_2}) \text{ s.t.} \\
R(\sigma_{i_1, i_2}) \leq & \sum_{j \in [m]} \mathcal{S}(\sigma_{i_1, i_2}, j) \cdot \Psi(i_1, j) \quad \forall i_1, i_2 \in [n], i_1 \neq i_2 \\
R(\sigma_{i_1, i_2}) = & \sum_{j \in [m]} \mathcal{S}(\sigma_{i_1, i_2}, j) \cdot \Psi(i_2, j) \quad \forall i_1, i_2 \in [n], i_1 \neq i_2 \\
\sum_{i_1, i_2 \in [n], i_1 \neq i_2} & \mathcal{S}(\sigma_{i_1, i_2}, j) = 1 \quad \forall j \in [m] \\
\mathcal{S}(\sigma_{i_1, i_2}, j) \geq 0 & \quad \forall i_1, i_2 \in [n], i_1 \neq i_2, \forall j \in [m],
\end{aligned}$$

where $R(\sigma_{i_1, i_2})$ is the contribution of signal σ_{i_1, i_2} to the revenue. Notice that in the optimal solution it is equal to $\max\{2 \sum_{j \in [m]} \mathcal{S}(\sigma_{i_1, i_2}, j) \cdot \Psi(i_1, j), \sum_{j \in [m]} \mathcal{S}(\sigma_{i_1, i_2}, j) \cdot \Psi(i_2, j)\}$.

Since LP1 consists of $O(n^2 m)$ variables and $O(n^2 + m)$ constraints (excluding the non-negativity constraints), it can be solved in polynomial time.

THEOREM 3.2. *Under the known-valuations setting, the RMS problem can be solved in polynomial time.*

3.1. Signaling versus Clustering

A clustering scheme is a special case of a signaling scheme, where the auctioneer cannot use randomness (see Section 2). This restricted case has been studied in Ghosh et al. [2007], and is equivalent to imposing the requirement that $\mathcal{S}(\sigma, j) \in \{0, 1\}$ for every σ and j in our framework.

THEOREM 3.3. *The optimal revenue that can be extracted by a signaling scheme is at most twice the optimal revenue that can be extracted by a clustering scheme, and this is tight.*

PROOF. The algorithm in Ghosh et al. [2007] produces a clustering scheme that extracts revenue that is greater or equal to half of

$$\min_{i'} \sum_j \max_{i \neq i'} \Psi(i, j). \quad (3)$$

Therefore, in order to establish the upper bound, it is sufficient to show that the revenue extracted by any signaling scheme is bounded by Equation (3). For every $i' \in [n]$, one can express the revenue of a signaling scheme \mathcal{S} as

$$\begin{aligned}
\mathbf{Rev}(\mathcal{S}) &= \sum_{\sigma} \max_{i'} \sum_j \mathcal{S}(\sigma, j) \Psi(i, j) \\
&\leq \sum_{\sigma} \max_{i \neq i'} \sum_j \mathcal{S}(\sigma, j) \Psi(i, j) \\
&\leq \sum_{\sigma} \sum_j \mathcal{S}(\sigma, j) \max_{i \neq i'} \Psi(i, j) \\
&= \sum_j \max_{i \neq i'} \Psi(i, j),
\end{aligned}$$

where the last equality holds since $\sum_{\sigma} \mathcal{S}(\sigma, j) = 1$. The upper bound follows.

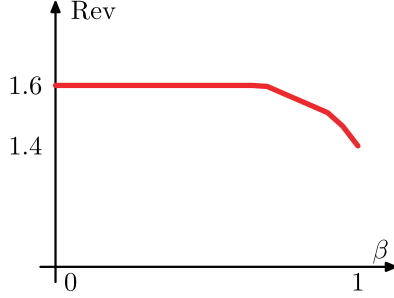


Fig. 1. Maximum revenue that can be extracted while guaranteeing a β -fraction of the social welfare.

To establish the lower bound, consider Example 5.3. While the optimal signaling scheme extracts revenue $\frac{n}{(n+1)}$, it is not difficult to verify that the optimal clustering scheme partitions items $1, \dots, m$ into pairs and leaves item 0 as a singleton. This clustering scheme extracts revenue $\frac{n}{2(n+1)}$, which is half of the revenue extracted by the optimal signaling scheme. \square

We remark that while we used Equation (3) as our benchmark, a better benchmark would be to compare the clustering revenue to the solution of the LP in Theorem 3.2. However, Example 5.3 demonstrates that $\frac{1}{2}$ is tight with respect to this benchmark as well.

3.2. Social Welfare versus Revenue

Increasing the revenue by signaling usually comes at the expense of degrading the social welfare. We show, however, that it is easy to calculate the best revenue one can get without degrading the social welfare by much. For every $j \in [m]$, let $\mu(j)$ denote the bidder i that maximizes the normalized valuation $\Psi(i, j)$ (which means that i also maximizes $V(i, j)$). Then, the optimal social welfare is given by $W^* = \sum_{j \in [m]} \Psi(\mu(j), j)$. By augmenting **LP1** with the constraint

$$\sum_{i_1, i_2 \in [n], i_1 \neq i_2} \sum_{j \in [m]} S(\sigma_{i_1, i_2}, j) \cdot \Psi(i_1, j) \geq \beta W^*,$$

we guarantee the highest possible revenue conditioned on preserving at least a β -fraction of the social welfare.

Example 3.4. Consider an instance with 4 agents and 5 goods with the uniform probability distribution ($p_j = \frac{1}{5}$ for every j) and the following valuations:

$$v = \begin{bmatrix} 1 & 0 & 1 & 1 & 0 \\ 0 & 2 & 0 & 0 & 1 \\ 3 & 0 & 2 & 0 & 3 \\ 0 & 4 & 0 & 1 & 2 \end{bmatrix}.$$

The curve in Figure 1 depicts the revenue that is obtained by solving **LP1** augmented with the preceding constraint for various values of β .

One might wonder whether the flat curve for small values of β is an artifact of this particular example or a general property. Theorem 3.5 shows that values of β that are smaller than $1/2$ do not affect the solution of **LP1**. Note that this theorem can be viewed as a signaling analogue of Theorem 2 in Ghosh et al. [2007] and its proof essentially follows similar arguments.

THEOREM 3.5. *There exists a revenue-optimal signaling scheme \mathcal{S} with $\mathbf{SW}(\mathcal{S}) \geq W^*/2$.*

PROOF. Consider an optimal signaling scheme and let σ be a signal and j an item such that $\mathcal{S}(\sigma, j) > 0$. If $\mu(j) \notin \{h_1^{\mathcal{S}}(\sigma), h_2^{\mathcal{S}}(\sigma)\}$, then we construct the new signaling scheme $\widehat{\mathcal{S}}$ obtained from \mathcal{S} by replacing signal σ with the two new signals σ', σ'' such that $\widehat{\mathcal{S}}(\sigma', j) = \mathcal{S}(\sigma, j)$, $\widehat{\mathcal{S}}(\sigma', j') = 0$ for $j' \neq j$; and $\widehat{\mathcal{S}}(\sigma'', j) = 0$, $\widehat{\mathcal{S}}(\sigma'', j') = \mathcal{S}(\sigma, j')$ for $j' \neq j$.

We argue that $\mathbf{Rev}(\widehat{\mathcal{S}}) \geq \mathbf{Rev}(\mathcal{S})$ which, by the optimality of \mathcal{S} , implies that $\mathbf{Rev}(\widehat{\mathcal{S}}) = \mathbf{Rev}(\mathcal{S})$. To that end, note that the contribution of σ to $\mathbf{Rev}(\mathcal{S})$ is

$$\begin{aligned} \mathbf{Rev}(\sigma, \mathcal{S}) &= \mathcal{S}(\sigma, j) \cdot \Psi(h_2^{\mathcal{S}}(\sigma), j) + \sum_{j' \neq j} \mathcal{S}(\sigma, j') \cdot \Psi(h_2^{\mathcal{S}}(\sigma), j') \\ &\leq \mathcal{S}(\sigma, j) \cdot \Psi(h_1^{\mathcal{S}}(\sigma), j) + \sum_{j' \neq j} \mathcal{S}(\sigma, j') \cdot \Psi(h_1^{\mathcal{S}}(\sigma), j'), \end{aligned}$$

whereas the contributions of σ' and σ'' to $\mathbf{Rev}(\widehat{\mathcal{S}})$ are

$$\mathbf{Rev}(\sigma', \widehat{\mathcal{S}}) = \mathcal{S}(\sigma, j) \cdot \max_{i \in [n]} \{\Psi(i, j)\}$$

and

$$\mathbf{Rev}(\sigma'', \widehat{\mathcal{S}}) = \max_{i \in [n]} \left\{ \sum_{j' \neq j} \mathcal{S}(\sigma, j') \cdot \Psi(i, j') \right\},$$

respectively. The argument follows since

$$\mathbf{Rev}(\sigma', \widehat{\mathcal{S}}) \geq \max \left\{ \mathcal{S}(\sigma, j) \cdot \Psi(h_2^{\mathcal{S}}(\sigma), j), \mathcal{S}(\sigma, j) \cdot \Psi(h_1^{\mathcal{S}}(\sigma), j) \right\}$$

and

$$\mathbf{Rev}(\sigma'', \widehat{\mathcal{S}}) \geq \min \left\{ \sum_{j' \neq j} \mathcal{S}(\sigma, j') \cdot \Psi(h_2^{\mathcal{S}}(\sigma), j'), \sum_{j' \neq j} \mathcal{S}(\sigma, j') \cdot \Psi(h_1^{\mathcal{S}}(\sigma), j') \right\}.$$

It follows that there exists a revenue-optimal signaling scheme \mathcal{S} , such that $\mathcal{S}(\sigma, j) > 0$ only if $\mu(j) \in \{h_1(\sigma), h_2(\sigma)\}$. Therefore, the social welfare of \mathcal{S} satisfies

$$\begin{aligned} \mathbf{SW}(\mathcal{S}) &= \sum_{\sigma \in [s]} \sum_{j \in [m]} \mathcal{S}(\sigma, j) \cdot \Psi(h_1(\sigma), j) \\ &\geq \sum_{\sigma \in [s]} \sum_{j \in [m]} \mathcal{S}(\sigma, j) \cdot \frac{\Psi(h_1(\sigma), j) + \Psi(h_2(\sigma), j)}{2} \\ &\geq \frac{1}{2} \sum_{\sigma \in [s]} \sum_{j \in [m]} \mathcal{S}(\sigma, j) \cdot \Psi(\mu(j), j) = \frac{1}{2} W^*. \end{aligned}$$

The assertion follows. □

4. OPTIMAL SIGNALING SCHEMES—THE BAYESIAN CASE

We now turn to discuss the more interesting Bayesian setting, considering a PSA $\mathcal{A} = \langle n, m, k, p, q, \{V_\ell\}_{\ell \in [k]} \rangle$, where q is a probability distribution over the valuation matrices $V_1, \dots, V_k \in \mathbb{R}_{\geq 0}^{n \times m}$. Our goal in this section is twofold: (1) proving that the

RMS problem under the Bayesian setting is NP-hard; and (2) presenting poly-time algorithms when k or m are fixed. Note that the RMS problem remains NP-hard if n is fixed, as long as both k and m are free parameters.

4.1. Tractable Special Cases

Let us start with developing an efficient algorithm for the RMS problem, assuming that $k = O(1)$ (without any restriction on n or m). Consider some s -signal signaling scheme \mathcal{S} for \mathcal{A} . Given a Bayesian outcome $\ell \in [k]$ and a signal $\sigma \in [s]$, let $h_1^{\mathcal{S}}(\ell, \sigma)$ and $h_2^{\mathcal{S}}(\ell, \sigma)$ denote the bidders i that realize $\max_{i \in [n]} \left\{ \sum_{j \in [m]} \mathcal{S}(\sigma, j) \cdot \Psi_{\ell}(i, j) \right\}$ and $\max_{2i \in [n]} \left\{ \sum_{j \in [m]} \mathcal{S}(\sigma, j) \cdot \Psi_{\ell}(i, j) \right\}$, respectively. (When the signaling scheme \mathcal{S} is clear from the context, we may omit the superscripts.) Using this notation, we can now state the following observation, which is established by repeating the line of arguments that led to Observation 3.1.

OBSERVATION 4.1. *There exists an optimal s -signal signaling scheme \mathcal{S} for \mathcal{A} such that given $\sigma, \sigma' \in [s]$, if $h_1^{\mathcal{S}}(\ell, \sigma) = h_1^{\mathcal{S}}(\ell, \sigma')$ and $h_2^{\mathcal{S}}(\ell, \sigma) = h_2^{\mathcal{S}}(\ell, \sigma')$ for every $\ell \in [k]$, then $\sigma = \sigma'$.*

Observation 4.1 implies that it is sufficient to consider $O(n^{2k})$ signals σ , each uniquely identified by $h_1(1, \sigma), h_2(1, \sigma), \dots, h_1(k, \sigma), h_2(k, \sigma)$. In order to formulate it as a concise linear program, we fix

$$\Lambda = \left\{ (i_1^1, i_2^1), \dots, (i_1^k, i_2^k) \mid i_h^{\ell} \in [n] \forall h \in \{1, 2\}, \ell \in [k] \wedge i_1^{\ell} \neq i_2^{\ell} \forall \ell \in [k] \right\}.$$

An optimal signaling scheme \mathcal{S} can now be constructed by solving the following linear program, denoted **LP2**:

$$\begin{aligned} \max \quad & \sum_{\lambda \in \Lambda} R(\sigma_{\lambda}) \text{ s.t.} \\ R(\sigma_{\lambda}) \leq \quad & \sum_{\ell \in [k]} q(\ell) \sum_{j \in [m]} \mathcal{S}(\sigma_{\lambda}, j) \cdot \Psi_{\ell}(\lambda(\ell, 1), j) \quad \forall \lambda \in \Lambda \\ R(\sigma_{\lambda}) = \quad & \sum_{\ell \in [k]} q(\ell) \sum_{j \in [m]} \mathcal{S}(\sigma_{\lambda}, j) \cdot \Psi_{\ell}(\lambda(\ell, 2), j) \quad \forall \lambda \in \Lambda \\ \sum_{\lambda \in \Lambda} \mathcal{S}(\sigma_{\lambda}, j) = \quad & 1 \quad \forall j \in [m] \\ \mathcal{S}(\sigma_{\lambda}, j) \geq \quad & 0 \quad \forall \lambda \in \Lambda, \forall j \in [m]. \end{aligned}$$

Since **LP2** consists of $O(n^{2k}m)$ variables and $O(n^{2k} + m)$ constraints (excluding the non-negativity constraints), it can be solved in polynomial time as long as k is constant.

THEOREM 4.2. *If k is fixed, then the RMS problem can be solved in polynomial time.*

Next, we show how to compute an optimal s -signal signaling scheme \mathcal{S} when $m = O(1)$ (without any restrictions on k and n). The main ingredient for this will be the following lemma (refer to Stanley [2004] for a proof).

LEMMA 4.3. *The number of distinct regions with nonempty interior⁴ defined by $t \geq m$ hyperplanes in \mathbb{R}^m is bounded from above by the Whitney number $W(m, t) = \sum_{i=0}^m \binom{t}{i} = O(t^m)$.*

Given some $\lambda \in \Lambda$, we define X_λ to be the region that contains every vector $x \in \mathbb{R}_{\geq 0}^m - \{0\}$, such that

$$\sum_{j \in [m]} x(j) \cdot \Psi_\ell(\lambda(\ell, 1), j) \geq \sum_{j \in [m]} x(j) \cdot \Psi_\ell(\lambda(\ell, 2), j) \geq \sum_{j \in [m]} x(j) \cdot \Psi_\ell(i, j) \quad (4)$$

for every $\ell \in [k]$ and $i \notin \{\lambda(\ell, 1), \lambda(\ell, 2)\}$. The key observation here is that if two signals $\sigma, \sigma' \in [s]$ are such that their corresponding vectors $\mathcal{S}(\sigma, \cdot), \mathcal{S}(\sigma', \cdot) \in \mathbb{R}_{\geq 0}^m$ fall into the same region X_λ , $\lambda \in \Lambda$, then we can merge them without decreasing the revenue. Therefore, if we can come up with a poly-size subset $\Lambda' \subseteq \Lambda$ so that the regions in $\{X_\lambda \mid \lambda \in \Lambda'\}$ cover the entire $\mathbb{R}_{\geq 0}^m - \{0\}$, then we can construct an optimal signaling scheme by picking one signal (the right one) for each region X_λ , such that $\lambda \in \Lambda'$.

So, how can we come up with such a subset $\Lambda' \subseteq \Lambda$? It turns out that although there are many regions X_λ , only a polynomially small subset of them have a nonempty interior. Indeed, the total number of linear constraints (4) involved in the definition of the regions X_λ , $\lambda \in \Lambda$, is n^2k (each linear constraint is of the form $\sum_{j \in [m]} x(j) \cdot \Psi_\ell(i, j) \geq \sum_{j \in [m]} x(j) \cdot \Psi_\ell(i', j)$ for some $i, i' \in [n]$ and $\ell \in [k]$). Since those linear constraints correspond to hyperplanes in \mathbb{R}^m , Lemma 4.3 guarantees that there are $O((n^2k)^m)$ regions X_λ with nonempty interiors.

Once the subset $\Lambda' \subseteq \Lambda$ of regions with nonempty interior has been identified, providing the linear constraints of (4) for each such region, we can rewrite **LP2**, dedicating a single signal σ to each region X_λ , such that $\lambda \in \Lambda'$ (the vector $\mathcal{S}(\sigma, \cdot)$ takes the role of the vector x in (4)). The resulting linear program consists of $O(|\Lambda'| \cdot m)$ variables and $O(|\Lambda'| \cdot kn)$ constraints, thus it can be solved in polynomial time.

It remains to show that we can efficiently enumerate the collection of regions with nonempty interiors. This is carried out by recursion on k : when $k = 1$, we can simply iterate through all the regions and check if their interior is nonempty. For the recursive step, observe that if the region X_λ has an empty interior, then clearly, so does the region $X_{\lambda \circ (i_1^{k+1}, i_2^{k+1})}$ for every $i_1^{k+1}, i_2^{k+1} \in [n]$. Therefore, we can keep iterating only through those regions that had a nonempty interior in the previous recursive level. By Lemma 4.3, the whole process requires checking $n(n-1) \cdot O((n^2(k-1))^m)$ regions.

THEOREM 4.4. *If m is fixed, then the RMS problem can be solved in polynomial time.*

4.2. Hardness of the General Case

Finally, we establish the NP-completeness of (the decision version of) the RMS problem in the Bayesian setting for the case $n = 3$. The inclusion of this problem in NP follows from Theorem 5.2, which ensures that it suffices to consider signaling schemes with at most m signals (which also implies that the number of bits required to represent the solutions of the LP is polynomial).

The remainder of the section is dedicated to proving that the RMS problem in the Bayesian setting is NP-hard. This is done by a reduction from MAX-CUT (problem ND16 in Garey and Johnson [1990]). Given a graph $G = (V, E)$ and two vertices $x, y \in V$, the MAX-CUT problem asks for the maximum integer k , such that there exists a

⁴A region $R \in \mathbb{R}^m$ is said to have a nonempty interior if it contains an m -dimensional open set.

1	K_1	t	1	K_2	u	1	K_2	u	1	1	v
2		K_1	2		K_2	2		K_2	2		1
3	K_1	K_1	3			3			3	1	1
(a) $\ell = 1$			(b) $2 \leq \ell \leq n - 1$			(c) $n \leq \ell \leq 2n - 3$			(d) $2n - 2 \leq \ell \leq 2n + m - 3$		

Fig. 2. Representation of the mapping from MAX-CUT to Bayesian signaling. The tables represent $\Phi_\ell(i, v)$ for $v \in V$ and $i \in \{1, 2, 3\}$. Values not specified in the tables are zero.

vertex subset $U \subseteq V$, $|\{x, y\} \cap U| = 1$, with at least k edges crossing between U and $V - U$. Given such an instance of MAX-CUT, assuming that $|V| = n$ and $|E| = m$, we construct a PSA \mathcal{A} with 3 bidders, n items (associated with the vertices in V), and $2n + m - 3$ Bayesian outcomes. It will be convenient to associate the Bayesian outcomes as follows: each $2 \leq \ell \leq n - 1$ is associated with some vertex $u \in V - \{x, y\}$; each $n \leq \ell \leq 2n - 3$ is also associated with some vertex $u \in V - \{x, y\}$; each $2n - 2 \leq \ell \leq 2n + m - 3$ is associated with some edge $(u, v) \in E$. Figure 2 depicts the values of $\Phi_\ell(i, j) = q(\ell) \cdot \Psi_\ell(i, j)$ for every $i \in \{1, 2, 3\}$, $j \in [n]$, and $\ell \in [2n + m - 3]$, where $K_1 \gg K_2 \gg 1$ are integers that will be determined in the course of the proof. This specifies everything we need for the reduction.

Suppose that the solution to the MAX-CUT instance is C^* and that this is realized by the vertex subset $X \subseteq V$, where $x \in X$ and $y \in Y = V - X$. We design a signaling scheme

\mathcal{S} with two signals σ_x, σ_y such that $\mathcal{S}(\sigma_x, u) = \begin{cases} 1 & \text{if } u \in X \\ 0 & \text{if } u \in Y \end{cases}$ and $\mathcal{S}(\sigma_y, u) = \begin{cases} 0 & \text{if } u \in X \\ 1 & \text{if } u \in Y \end{cases}$.

It can be checked that the revenue generated by this signaling scheme is

$$\sum_{\ell \in [2n+m-3]} \sum_{\sigma \in \{\sigma_x, \sigma_y\}} \max_{2_{i \in \{1,2,3\}}} \left\{ \sum_{u \in V} \Phi_\ell(i, u) \cdot \mathcal{S}(\sigma, u) \right\} = 2K_1 + (n - 2)K_2 + m + C^* .$$

The reduction is completed by showing that this is an upper bound on the revenue generated by any signaling scheme.

Consider some s -signal signaling scheme \mathcal{S} . The revenue of \mathcal{S} is

$$\begin{aligned} \mathbf{Rev}(\mathcal{S}) &= \sum_{\ell \in [2n+m-3]} \sum_{\sigma \in [s]} \max_{2_{i \in \{1,2,3\}}} \left\{ \sum_{u \in V} \Phi_\ell(i, u) \cdot \mathcal{S}(\sigma, u) \right\} \\ &= K_1 \sum_{\sigma \in [s]} \max\{\mathcal{S}(\sigma, x), \mathcal{S}(\sigma, y)\} \\ &\quad + K_2 \sum_{u \in V - \{x, y\}} \sum_{\sigma \in [s]} [\min\{\mathcal{S}(\sigma, x), \mathcal{S}(\sigma, u)\} + \min\{\mathcal{S}(\sigma, y), \mathcal{S}(\sigma, u)\}] \\ &\quad + \sum_{(u, v) \in E} \sum_{\sigma \in [s]} \max\{\mathcal{S}(\sigma, u), \mathcal{S}(\sigma, v)\} . \end{aligned} \tag{5}$$

We argue that if \mathcal{S} is optimal, then for each signal σ , either $\mathcal{S}(\sigma, x) = 0$ or $\mathcal{S}(\sigma, y) = 0$. Indeed, if there is a signal σ with $\mathcal{S}(\sigma, x)$ and $\mathcal{S}(\sigma, y)$ simultaneously positive, then we can split this signal into two signals σ', σ'' such that $\mathcal{S}(\sigma', x) = \mathcal{S}(\sigma, x)$ and $\mathcal{S}(\sigma', u) = 0$ for $u \neq x$; $\mathcal{S}(\sigma'', x) = 0$ and $\mathcal{S}(\sigma'', u) = \mathcal{S}(\sigma, u)$ for $u \neq x$. Taking K_1 to be sufficiently large ensures that the revenue increases following this transformation.

Now, consider a signal $\sigma \in [s]$ with $\mathcal{S}(\sigma, x) > 0$. If there exists some $u \neq x$ such that $\mathcal{S}(\sigma, u) > \mathcal{S}(\sigma, x)$, then there must exist a signal $\sigma' \in [s]$ such that $\mathcal{S}(\sigma', u) < \mathcal{S}(\sigma', x)$ since $\sum_{\sigma \in [s]} \mathcal{S}(\sigma, x) = \sum_{\sigma \in [s]} \mathcal{S}(\sigma, u) = 1$. Therefore, we can increase the value of

$\mathcal{S}(\sigma', u)$ and decrease the value of $\mathcal{S}(\sigma, u)$ by a small value δ , obtaining a valid signaling scheme with larger revenue (this is due to the fact that K_2 is large compared to 1). Similarly, we can claim that in an optimal signaling scheme, a signal $\sigma \in [s]$ with $\mathcal{S}(\sigma, y) > 0$ has $\mathcal{S}(\sigma, u) \leq \mathcal{S}(\sigma, y)$ for all $u \in V$. The same argument also implies that for each signal σ that has positive probability, either $\mathcal{S}(\sigma, x) > 0$ or $\mathcal{S}(\sigma, y) > 0$.

Consider a signal σ and let $x = u_0, u_1, \dots, u_k \in V$ be the items with positive $\mathcal{S}(\sigma, u_i)$. Assume without loss of generality that $\mathcal{S}(\sigma, u_0) \geq \mathcal{S}(\sigma, u_1) \geq \dots \geq \mathcal{S}(\sigma, u_k) > \mathcal{S}(\sigma, u_{k+1}) = 0$. Now, split signal σ into $k+1$ signals $\sigma^0, \sigma^1, \dots, \sigma^k$ such that $\mathcal{S}(\sigma^i, s) = \mathcal{S}(\sigma^i, u_1) = \dots = \mathcal{S}(\sigma^i, u_i) = \mathcal{S}(\sigma, u_i) - \mathcal{S}(\sigma, u_{i+1})$. By substituting the split signals into (5), we conclude that the revenue is kept unchanged.

Therefore, PSA instances produced by our reduction always admit an optimal s -signal signaling scheme \mathcal{S} such that for every signal $\sigma \in [s]$, there exists a vertex subset $U_\sigma \subseteq V$ and a real $0 < p_\sigma \leq 1$ satisfying:

- (1) $|U_\sigma \cap \{x, y\}| = 1$;
- (2) $\mathcal{S}(\sigma, u) = p_\sigma$ for every $u \in U_\sigma$;
- (3) $\mathcal{S}(\sigma, u) = 0$ for every $u \notin U_\sigma$;
- (4) $\sum_{\sigma: v \in U_\sigma} p_\sigma = 1$ for every $v \in V$; and
- (5) $\sum_{\sigma \in [s]} p_\sigma = \sum_{\sigma: x \in U_\sigma} p_\sigma + \sum_{\sigma: y \in U_\sigma} p_\sigma = 2$. This is employed in order to prove that $\mathbf{Rev}(\mathcal{S}) \leq 2K_1 + (n-2)K_2 + m + C^*$. From (5), we see that $\mathbf{Rev}(\mathcal{S}) = 2K_1 + (n-2)K_2 + \sum_{(u,v) \in E} \sum_{\sigma \in [s]} \max\{\mathcal{S}(\sigma, u), \mathcal{S}(\sigma, v)\}$, so it remains to show that

$$\sum_{(u,v) \in E} \sum_{\sigma \in [s]} \max\{\mathcal{S}(\sigma, u), \mathcal{S}(\sigma, v)\} \leq m + C^*. \quad (6)$$

To that end, note that every edge $(u, v) \in E$ and signal $\sigma \in [s]$ contributes p_σ to the left-hand side in (6) if $|\{u, v\} \cap U_\sigma| \geq 1$; and 0 otherwise. Therefore, the left-hand side in (6) is equal to

$$\sum_{(u,v) \in E} \left[1 + (1/2) \sum_{\sigma: |\{u,v\} \cap U_\sigma| = 1} p_\sigma \right] = m + \sum_{\sigma \in [s]} \frac{p_\sigma}{2} |\partial(U_\sigma)| \leq m + C^*,$$

where $\partial(U)$ is the set of edges with exactly one endpoint in U , and the last inequality follows, since $\sum_{\sigma \in [s]} \frac{p_\sigma}{2} = 1$, hence $\sum_{\sigma \in [s]} \frac{p_\sigma}{2} |\partial(U_\sigma)|$ can be viewed as the average size of cuts corresponding to the vertex subsets U_σ , $\sigma \in [s]$. This establishes the following theorem.

THEOREM 4.5. *The decision version of the RMS problem in the Bayesian setting is NP-complete. The problem is hard already for $n = 3$.*

5. BOUNDING THE NUMBER OF SIGNALS

In a market with n bidders, the algorithm described in Section 3 generates a signaling scheme with $O(n^2)$ signals. In fact, some instances might require that many signals in order to produce the optimal signaling.

Example 5.1. Consider a KPSA with $n^2 - n$ items appearing with uniform probability. Each item $I_{i,i'}$ is labeled with an ordered pair (i, i') of bidders, where $V(i, I_{i,i'}) = 1$, $V(i', I_{i,i'}) = \frac{1}{2}$, and $V(i'', I_{i,i'}) = 0$ for every $i'' \in [n] - \{i, i'\}$. The optimal signaling generates revenue of $\frac{3}{4}$, by emitting a signal $\sigma_{i,i'}$ when either $I_{i,i'}$ or $I_{i',i}$ is chosen. Notice that $s = \binom{n}{2}$ signals are required to implement this signaling scheme. It is not hard to see that with fewer signals, it is impossible to achieve this revenue.

The KPSA described in Example 5.1 has $m = \Omega(n^2)$ items. Can we construct a similar bad example with much fewer items? More generally, can we bound the minimum number of signals required to implement an optimal signaling scheme as a function depending only on m ?

THEOREM 5.2. *Every PSA has an optimal signaling scheme with m signals.*

PROOF. For brevity, we establish the assertion assuming the known-valuations setting; the proof of the Bayesian setting follows from the same line of arguments. Consider some KPSA $\mathcal{A} = \langle n, m, p, V \rangle$ and let \mathcal{S} be an s -signal signaling scheme for \mathcal{A} that minimizes s . Assume by contradiction that $s > m$. Associate with each signal $\sigma \in [s]$ a vector $\vec{q}_\sigma \in \mathbb{R}_{\geq 0}^m$, where $q_\sigma(j) = \mathcal{S}(\sigma, j)$ for every $j \in [m]$. We can rewrite the contribution of σ to the revenue of \mathcal{S} as $\mathbf{Rev}(\sigma, \mathcal{S}) = \max_{i \in [n]} \{ \sum_{j \in [m]} q_\sigma(j) \cdot \Psi(i, j) \}$, so $\mathbf{Rev}(\cdot, \mathcal{S})$ is a homogeneous, but not necessarily linear, operator.

Since $s > m$, the vector collection $\{ \vec{q}_\sigma \mid \sigma \in [s] \}$ must exhibit linear dependencies. Therefore, there must exist some reals $x_1, \dots, x_r, x_{r+1}, \dots, x_t > 0$ and signals $\sigma_1, \dots, \sigma_r, \sigma_{r+1}, \dots, \sigma_t$ such that

$$x_1 \vec{q}_{\sigma_1} + \dots + x_r \vec{q}_{\sigma_r} = x_{r+1} \vec{q}_{\sigma_{r+1}} + \dots + x_t \vec{q}_{\sigma_t} .$$

Consider the s -signal signaling scheme \mathcal{S}' defined by the modified signals σ' obtained from \mathcal{S} by setting

$$\vec{q}_{\sigma'_z} = \begin{cases} (1 + \epsilon x_z) \vec{q}_{\sigma_z} & \text{if } 1 \leq z \leq r \\ (1 - \epsilon x_z) \vec{q}_{\sigma_z} & \text{if } r + 1 \leq z \leq t \\ \vec{q}_{\sigma_z} & \text{otherwise} \end{cases} .$$

For any $\epsilon \in \left[-\frac{1}{\max_{1 \leq z \leq r} \{x_z\}}, \frac{1}{\max_{r+1 \leq z \leq t} \{x_z\}} \right]$, the resulting signaling scheme \mathcal{S}' is valid, as $\sum_{\sigma' \in [s]} \vec{q}_{\sigma'} = \sum_{\sigma \in [s]} \vec{q}_\sigma = \mathbf{1}$ and $\vec{q}_{\sigma'} \geq 0$.

Now, notice that the revenue of the new signaling scheme satisfies

$$\mathbf{Rev}(\mathcal{S}') = \mathbf{Rev}(\mathcal{S}) + \epsilon \left[\sum_{z=1}^r x_z \mathbf{Rev}(\sigma_z, \mathcal{S}) - \sum_{z=r+1}^t x_z \mathbf{Rev}(\sigma_z, \mathcal{S}) \right] .$$

Since \mathcal{S} is optimal, it must be the case that $\sum_{z=1}^r x_z \mathbf{Rev}(\sigma_z, \mathcal{S}) - \sum_{z=r+1}^t x_z \mathbf{Rev}(\sigma_z, \mathcal{S}) = 0$ (recall that ϵ can be taken to be positive or negative). Thus, we can take ϵ to be either $-\frac{1}{\max_{1 \leq z \leq r} \{x_z\}}$ or $\frac{1}{\max_{r+1 \leq z \leq t} \{x_z\}}$ and get an optimal signaling scheme with less than s signals, in contradiction to the choice of \mathcal{S} . \square

If one can achieve revenue R with s signals, then it is trivial to see that with $s' \leq s$ signals, one can achieve revenue $\left\lfloor \frac{s'}{s} \right\rfloor R$. This turns out to be the best possible outcome in some cases.

Example 5.3. Consider $n + 1$ items $\{0, 1, \dots, n\}$, each chosen with probability $\frac{1}{n+1}$ and $n + 1$ bidders with valuations $V(i, i) = 1$ for $i = 1, \dots, n$, $V(0, 0) = n$, and $V(i, j) = 0$ otherwise. The optimal signaling scheme \mathcal{S} uses n signals $\sigma_1, \dots, \sigma_n$ such that

$S(\sigma_i, i) = 1$ and $S(\sigma_i, 0) = \frac{1}{n}$ for $i = 1, \dots, n$ and zero otherwise. The revenue of this scheme is $\frac{n}{n+1}$. Now, for any $s < n$ and an s -signal signaling scheme S' , we have

$$\begin{aligned} \mathbf{Rev}(S') &= \sum_{\sigma' \in [s]} \max_{i \in [n]} \left\{ \sum_{j \in [m]} S'(\sigma', j) \cdot \Psi(i, j) \right\} \\ &\leq \sum_{\sigma' \in [s]} \max_{i \in [n] - \{0\}} \left\{ \sum_{j \in [m]} S'(\sigma', j) \cdot \Psi(i, j) \right\} \\ &\leq \sum_{\sigma' \in [s]} \frac{1}{n+1} = \frac{s}{n+1}. \end{aligned}$$

REFERENCES

- George A. Akerlof. 1970. The market for 'lemons': Quality uncertainty and the market mechanism. *Quarterly J. Econ.* 84, 3, 488–500.
- Jacques Cremer and Richard P. McLean. 1985. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* 53, 2, 345–361.
- Jacques Cremer and Richard P. McLean. 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56, 6, 1247–1257.
- Shahar Dobzinski, Hu Fu, and Robert D. Kleinberg. 2011. Optimal auctions with correlated bidders are easy. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 129–138.
- Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *Amer. Econ. Rev.* 97, 1, 242–259.
- Yuval Emek, Michal Feldman, Iftah Gamzu, and Moshe Tennenholtz. 2011. Signaling schemes for revenue maximization. *AdAuctions Workshop*.
- Yossi Feinberg and Moshe Tennenholtz. 2005. Anonymous bidding and revenue maximization. *B.E. J. Theoret. Econ. Topics* 5, 1, 1–12.
- Michael R. Garey and David S. Johnson. 1990. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY.
- Arpita Ghosh, Hamid Nazerzadeh, and Mukund Sundararajan. 2007. Computing optimal bundles for sponsored search. In *Proceedings of the 3rd International Conference on Internet and Network Economics (WINE)*. Springer-Verlag, Berlin, 576–583.
- Jonathan Levin and Paul Milgrom. 2010. Online advertising: Heterogeneity and conflation in market design. *Amer. Econ. Rev.* 100, 2, 603–607.
- Paul Milgrom. 2010. Simplified mechanisms with an application to sponsored-search auctions. *Games Econ. Behav.* 70, 1, 62–70.
- Paul R. Milgrom and Robert J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* 50, 5, 1089–1122.
- Peter Bro Miltersen and Or Sheffet. 2012. Send mixed signals—Earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 234–247.
- S. Muthukrishnan. 2009. Ad exchanges: Research issues. In *Internet and Network Economics*, Stefano Leonardi (Ed.), Lecture Notes in Computer Science, Vol. 5929, Springer Berlin, Chapter 1, 1–12.
- Roger B. Myerson. 1981. Optimal auction design. *Math. Oper. Res.* 6, 1, 58–73.
- Christos H. Papadimitriou and George Pierrakos. 2011. On optimal single-item auctions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 119–128.
- Motty Perry and Philip J. Reny. 1999. On the failure of the linkage principle in multi-unit auctions. *Econometrica* 67, 4, 895–900.
- Michael Spence. 1973. Job market signaling. *Quarterly J. Econ.* 87, 3, 355–374.
- Michael Spence. 2002. Signaling in retrospect and the informational structure of markets. *Amer. Econ. Rev.* 92, 3, 434–459.
- Richard P. Stanley. 2004. An introduction to hyperplane arrangements. In *Lecture Notes*, IAS/Park City Mathematics Institute.

- Foucault Thierry and Lovo Stefano. 2003. *Linkage principle, Multi-Dimensional Signals and Blind Auctions*. Les Cahiers de Recherche 785, HEC Paris.
- Hal R. Varian. 2007. Position auctions. *Intern. J. Indus. Org.* 25, 6, 1163–1178.
- William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16, 1, 8–37.
- Robert J. Weber. 2003. Auction Theory: By Vijay Krishna. Academic Press, 2002. *Games Econ. Behav.* 45, 2, 488–497.

Received December 2012; accepted May 2013