

How to Sell Information Optimally: An Algorithmic Study

Yang Cai 

Yale University, New Haven, CT, USA
yang.cai@yale.edu

Grigoris Velegkas 

Yale University, New Haven, CT, USA
grigoris.velegkas@yale.edu

Abstract

We investigate the algorithmic problem of selling information to agents who face a decision-making problem under uncertainty. We adopt the model recently proposed by Bergemann et al. [4], in which information is revealed through signaling schemes called *experiments*. In the single-agent setting, any mechanism can be represented as a menu of experiments. Our results show that the computational complexity of designing the revenue-optimal menu depends heavily on the way the model is specified. When all the parameters of the problem are given explicitly, we provide a polynomial time algorithm that computes the revenue-optimal menu. For cases where the model is specified with a succinct implicit description, we show that the tractability of the problem is tightly related to the efficient implementation of a *Best Response Oracle*: when it can be implemented efficiently, we provide an additive FPTAS whose running time is independent of the number of actions. On the other hand, we provide a family of problems, where it is computationally intractable to construct a best response oracle, and we show that it is NP-hard to get even a constant fraction of the optimal revenue. Moreover, we investigate a generalization of the original model by Bergemann et al. [4] that allows multiple agents to compete for useful information. We leverage techniques developed in the study of auction design (see e.g. [5, 1, 6, 7, 8]) to design a polynomial time algorithm that computes the revenue-optimal mechanism for selling information.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory

Keywords and phrases Mechanism Design, Algorithmic Game Theory, Information Design

Digital Object Identifier 10.4230/LIPIcs.ITCS.2021.81

Related Version A full version of the paper is available at <https://arxiv.org/abs/2011.14570>.

Funding *Yang Cai*: Supported by a Sloan Foundation Research Fellowship and the NSF Award CCF-1942583 (CAREER).

Grigoris Velegkas: Partially supported by a PhD Scholarship from the Onassis Foundation and a PhD Scholarship from the Bodossaki Foundation.

Acknowledgements We would like to thank Dirk Bergemann for helpful discussions and the anonymous reviewers for their valuable comments and suggestions.

1 Introduction

Decision-making heavily relies on information availability. The uneven distribution of information thus enables markets for trading information. Imagine a bank reviewing a loan application. Information about the borrower's financial status clearly influences the bank's lending decisions. In this setting, the bank already has some private knowledge about the borrower, i.e., through prior interactions, but may still be willing to pay to acquire supplemental information to guide its decision-making. Indeed, Equifax, the credit report



© Yang Cai and Grigoris Velegkas;

licensed under Creative Commons License CC-BY

12th Innovations in Theoretical Computer Science Conference (ITCS 2021).

Editor: James R. Lee; Article No. 81; pp. 81:1–81:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

agency, provides its business consumers, e.g., banks and credit card companies, a product called Undisclosed Debt Monitoring, that tracks negative information about individual borrowers.¹

How should the information owner reveal and price the information? We adopt the model introduced by Bergemann et al. [4] to study this problem. Their model involves a data buyer and a data seller. The data buyer faces a decision under uncertainty, and his payoff depends on the *action* he decides to take and the underlying *state* of the world. Initially, the buyer only has some *imperfect* knowledge about the state, i.e. a prior distribution over the possible states. This piece of information is private to the buyer. On the contrary, the data seller knows the state of the world and can sell supplemental information to the buyer. Since the buyer's willingness to pay for supplemental information is determined by the precision of his own prior belief, we refer to the buyer's prior belief as his type. How does the seller optimize her revenue if the buyer's type is assumed to be drawn from a known distribution? To screen heterogeneous buyer types, the seller offers a menu of information products. Each product has the form of a *statistical experiment*, whose result reveals a signal that is correlated with the underlying state. Bergemann et al. [4] investigate what experiments should be included and how to price them. They obtain analytic solution of the revenue-optimal menu in two special cases: (i) the case with only two possible buyer types and (ii) the case with two states and two possible actions for the buyer to choose from. However, general characterization of the revenue-optimal menu remains elusive. In this paper, we initiate the algorithmic study of this problem and investigate the computational complexity for finding the revenue-optimal menu.

Our first result considers an explicit representation of the problem, where the input contains the buyer's type distribution and his payoff for each action and state pair.

Result 1: We design an algorithm that computes the revenue-optimal menu in time polynomial in the number of buyer types, the number of buyer actions, and the number of underlying states. (Theorem 3)

For many settings of interest, the model is too expensive to be specified explicitly but has a natural succinct implicit description. Consider the following motivating example. Suppose there is a traffic network G , and a binary state ω that indicates the level of congestion on the edges of G . A driver wants to go from a vertex s to another vertex t , and his payoff is H minus the expected travel time from s to t using the path he picks.² Of course, which path is the fastest depends on the driver's belief of the state ω . Suppose Waze is offering a service that provides the driver supplemental information about the congestion. How should Waze price its service? In this setting, the driver is the data buyer and the number of actions available to him is exactly the number of $s - t$ paths, which can be exponential in the size of the network G . Applying our first result is thus computationally inefficient in this setting. Our second result concerns exactly these settings with succinct implicit descriptions. We show that the key to tractability is the existence of a computationally efficient *Best Response Oracle*, that is, an oracle that accepts a distribution over the underlying state as input and outputs an action with the highest expected payoff. Clearly, in the example above, it is straightforward to construct a computationally efficient best response oracle – simply assign the expected congestion on each edge as its length and run any shortest path algorithm on the graph G . We show that finding the revenue-optimal menu is tractable as long as there exists a best response oracle.

¹ <https://www.equifax.com/business/undisclosed-debt-monitoring/>

² H is a sufficiently large constant so that the payoff is always nonnegative.

Result II: For any setting with a constant number of possible underlying states, we design an FPTAS to compute an *up-to- ε optimal menu*, i.e., a menu whose revenue is at most ε less than the optimum, given access to a best response oracle.³ (Theorem 9)

Without the best response oracle, we show that it is NP-hard to even find a constant factor approximation to the optimal revenue for a family of succinctly describable instances with only two underlying states and one buyer type. (Theorem 12)

We also investigate extensions of the basic model studied in Bergemann et al. [4]. First, we consider the setting where multiple data buyers are competing with each other to receive an informative signal from the data seller. We obtain the following generalization of Result I.

Result III: We design an algorithm that computes the revenue-optimal mechanism in time polynomial in the *total number of buyer types*, the number of buyer actions, and the number of underlying states. (Theorem 19)

Note that the straightforward generalization of Result I only gives an algorithm that runs in time polynomial in the *total number of buyer type profiles*, which is exponential in the number of buyers. Our algorithm in Result III runs in time polynomial in the *total number of buyer types*, which is the description size for specifying each buyer’s type distribution.

In Section 5, we discuss another natural extension. We can treat a buyer’s payoffs for taking actions in various states as the buyer’s private information and may be different across buyers. In other words, a buyer type is no longer just the buyer’s prior belief of the underlying state but also his payoff function. We show that all our results (c.f. Result I, II, and III) can be easily extended to handle this case, and the modification is summarized in Section 5.

Our Approach

In the explicit model, we first show that the revenue-optimal menu can be captured by a LP with polynomially many decision variables but exponentially many constraints. We leverage a technique introduced by Chen et al. [15] to transform it to an equivalent LP with polynomially many constraints. In the implicit model with a best response oracle, the main difficulty is that the optimal menu may contain experiments that use an arbitrary number of signals, and as a result the menu cannot even be represented in polynomial time. To overcome this difficulty, we first argue that there always exists an up-to- ε optimal menu that only contains experiments that use a small number of signals, then apply an algorithm similar to the one in the explicit model to find such a menu. For the multi-agent setting, one can use a LP similar to the one in the explicit model to capture the optimal mechanism. The issue with this approach is that both of the number of variables and constraints are exponential in the number of agents, which is too large to solve. Our solution is inspired by an approach used to computing the revenue-optimal multi-item auctions [5, 1, 6, 7, 8]. The key idea is to first represent mechanisms in a succinct way known as the “reduced forms” and use a LP to search for the revenue-optimal reduced form. However, as a reduced form is only a succinct description, given a reduced form, it is not obvious what mechanism it corresponds to. Indeed, it is not even clear whether it corresponds to any mechanism. The main technical barrier we overcome is to design efficient algorithms to (i) check the feasibility of a reduced form and (ii) to implement a feasible reduced form as an actual mechanism.

³ Our algorithm runs in time polynomial in the total number of buyer types and $1/\varepsilon$.

1.1 Related Work

Relationship with Monopoly Pricing

A well-studied problem from mechanism design, the *monopoly pricing problem*, bears some resemblance to our problem. The monopoly pricing problem asks what is the revenue-optimal mechanism to sell one or more items to a buyer, whose valuation/willingness to pay for the items is drawn from a known distribution. In single-dimensional settings, the problem is completely resolved by [35, 36]. In multi-dimensional settings, complete characterizations are known only in several special cases [27, 28, 30, 18, 26, 22, 21], but simple and approximately optimal mechanisms [13, 14, 31, 2, 34, 37, 9, 12], as well as algorithmic characterizations [5, 1, 6, 7, 8] have been discovered in fairly general settings. The nature of information goods adds an extra layer of difficulty to the pricing problem. The value of information is determined by how much such information can improve the quality of decision-making. Buyers with different beliefs do not simply have different values for different experiments, but they may even disagree on their ranking. This richness in buyer valuations does not happen in single-dimensional monopoly pricing, but already exists in single-dimensional information pricing, where there are only two possible underlying states.⁴ As a result, the optimal menu in information pricing has a more complex structure. For example, Bergemann et al. [4] showed that even in the two-state case, the seller sometimes needs to use randomized experiments to maximize her revenue. In this work, we show that despite the new challenges of selling information products, some of the algorithmic ideas from monopoly pricing are still useful.

Relationship with Information Design

Similar to our problem, the designer constructs a signaling scheme that reveals partial information about the underlying state to influence the action of the agents in information design [32]. There has been growing interest in the algorithmic study of information design [24, 23, 25, 16, 19, 10]. The fundamental difference between our setting and information design is that, in our setting, the buyer's action does not have direct effect on the seller's utility. The seller only derives utility from the monetary transfers received from the buyer.

Relationship with other Information-Selling Models

The work that is most related to the problem we are studying is by Babaioff et al. [3]. Similar to us, they consider a seller who knows the state of the world ω and wants to sell information to a buyer whose prior is drawn from some distribution Θ . However, there is a subtle but crucial difference between our models. In their work, the seller's information disclosure strategy and the price are allowed to be dependent on the realized state of the world ω . By contrast, our model requires the seller to commit to a mechanism before the realization of the state ω . Their main results state that (i) when ω and the buyer's type are independently distributed, revelation principle holds, i.e., a single-round interaction suffices; and (ii) when ω and the buyer's type are correlated, a full surplus extraction mechanism, similar to Crémer and McLean [17], exists and can be computed efficiently. These results differ quite substantially from the structural results for the model we study here [4]. For instance, full surplus extraction is in general impossible in our model. In a recent work, Chen et al. [15] extends the results by Babaioff et al. [3] to the setting where the buyer is budget-constrained, and improves upon some of the algorithmic results.

⁴ In this case, the buyer's prior belief can be represented using a single real number in $[0, 1]$.

2 Preliminaries

Model and Notation

The data buyer faces a decision problem under uncertainty. The state of the world ω is drawn from a state space Ω . The buyer chooses an action a from an action space A . Throughout the paper, we use m to denote the size of A . The buyer's **ex-post utility** for choosing action a under state ω is defined to be $u_{\omega,a}$ and is assumed to lie in $[0, 1]$. The buyer has some prior information about the state of the world which is denoted by θ and comes from a set $\Theta \subseteq \Delta\Omega$. We call θ the *type* of the buyer, and use θ_ω to denote the probability that the buyer assigns to the event that the state of the world is ω . The type of the buyer is distributed according to F . Apart from the buyer, there is also a *seller* who observes the state of the world and is willing to sell supplemental information to the buyer. We refer to the buyer as he and to the seller as she.

Experiment

The seller provides supplemental information to the buyer via a *signaling scheme* which we call *experiment*. A signaling scheme is a commitment to $|\Omega|$ probability distributions over a set of different signals S , such that when the state of the world is realized, the seller draws a signal from the corresponding distribution and sends it to the buyer. We denote such an experiment by $E = (S, \pi(E))$, where $\pi(E) : \Omega \rightarrow \Delta S$ denotes the distributions that experiment E is using. We denote the probability that experiment E sends signal s_k when the state of the world is ω by $\pi_{\omega,k}(E) = \Pr[s_k|\omega]$. In this work, it is useful to think of the experiment $\pi(E)$ as a matrix whose rows are indexed by the states of the world and columns are indexed by the signals. A *menu* of experiments is a collection $\mathcal{M} = \{(E, t(E))\}$, where $t(E) \in [0, 1]$ is the payment the buyer has to make when he purchases experiment E . In the multi-agent setting, a menu is insufficient to generate the optimal revenue, and our goal there is to compute the revenue-optimal *mechanism*. In the single-agent setting, the interaction between the seller and the buyer works as follows:

1. The seller posts a menu \mathcal{M} .
2. The state of the world ω and the type of the buyer θ are realized.
3. The buyer chooses some experiment E from the menu based on his type and pays $t(E)$.
4. The seller sends the buyer a signal s that is drawn from $\pi_{\omega,\cdot}(E)$.
5. The buyer chooses an action a , based on his original belief θ and the signal s , and receives utility $u_{\omega,a}$.

The Value of an Experiment

To understand the behavior of the buyer, we first explain how the buyer evaluates an experiment. We first explain how the buyer would act if the only information available to him was his type θ . Since the buyer picks an action that maximizes his expected utility, his best move without receiving any additional information from the seller is $a(\theta) = \arg \max_a \sum_{\omega} \theta_{\omega} u_{\omega,a}$ and his **base utility** following that move is $u(\theta) = \max_a \sum_{\omega} \theta_{\omega} u_{\omega,a}$. If he receives extra information from the seller, he updates his beliefs and may choose a new action that induces higher expected value based on his posterior distribution over the states. After receiving signal s_k from experiment E his belief about the state of the world is

$$\Pr[\omega|s_k, \theta] = \frac{\theta_{\omega} \pi_{\omega,k}(E)}{\sum_{\omega' \in \Omega} \theta_{\omega'} \pi_{\omega',k}(E)}.$$

Hence, the best action is

$$a(s_k|\theta) \in \arg \max_a \sum_{\omega} \left(\frac{\theta_{\omega} \pi_{\omega,k}(E)}{\sum_{\omega' \in \Omega} \theta_{\omega'} \pi_{\omega',k}(E)} \right) u_{\omega,a},$$

which yields **conditional expected utility**

$$u(s_k|\theta) := \max_a \sum_{\omega} \left(\frac{\theta_{\omega} \pi_{\omega,k}(E)}{\sum_{\omega'} \theta_{\omega'} \pi_{\omega',k}(E)} \right) u_{\omega,a}.$$

When it is clear from the context, we might drop θ in the previous expressions. Notice that after computing his posterior, the buyer's conditional expected utility is linear in the actions so we can assume w.l.o.g. that he picks a single action and not a distribution over actions. Taking the expectation over the signal he will receive, we denote the **value of the experiment** E for type θ to be

$$V_{\theta}(E) = \sum_{s_k \in S} \max_a \left\{ \sum_{\omega} \theta_{\omega} \pi_{\omega,k}(E) u_{\omega,a} \right\}.$$

We assume that the buyer is quasilinear. We denote the **expected net utility** of type θ for experiment E as $V_{\theta}(E) - t(E)$.

Notice that two different types θ, θ' may have different favorite actions under the same signal. As a result, $V_{\theta}(E)$ is not a linear function over θ for a fixed experiment E even when the type θ is single-dimensional, i.e., $|\Omega| = 2$. This is in sharp contrast to standard single-dimensional auction design settings, where the agent's value is always a linear function over the type when the allocation is fixed.⁵

IC and IR Menu

The buyer chooses the experiment that gives him the highest net utility, and we slightly abuse notation and denote by $E(\theta)$ the experiment of the menu \mathcal{M} that type θ prefers

$$E(\theta) = \arg \max_E V_{\theta}(E) - t(E).$$

If all experiments give him expected net utility smaller than his base utility $u(\theta)$, he will not purchase any experiment. For convenience, we assume that there is a null experiment that provides no information offered at price 0, so now without loss of generality every type takes an experiment from the menu. We sometimes abuse notation to call a menu *Incentive Compatible* (IC) and *Individually Rational* (IR). By that, we mean that every buyer type θ selects the experiment $(E(\theta), t(\theta))$ that maximizes his expected net utility. Formally, these constraints are captured by the following two sets of inequalities:

$$\begin{aligned} V_{\theta}(E(\theta)) - t(\theta) &\geq V_{\theta}(E(\theta')) - t(\theta'), \quad \forall \theta, \theta' \in \Theta, \\ V_{\theta}(E(\theta)) - t(\theta) &\geq u(\theta), \quad \forall \theta \in \Theta. \end{aligned}$$

Once we have fixed the parameters of the model, we denote by $\text{REV}(\mathcal{M})$ the revenue that menu \mathcal{M} generates and by OPT the optimal revenue in this setting.

⁵ Indeed, this linearity even holds for quite general multi-dimensional settings in auction design, for example, when the buyer has additive valuations.

3 Optimal Menu for a Single Agent

In this section, we consider the problem of computing the optimal menu for a single agent. As we will show, the computational complexity of the problem is tightly related to the way that the problem is specified. We consider three different models.

- **Explicit model:** the distribution F and the ex-post utility matrix U are given explicitly in the input.
- **Implicit model with a best response (BR) oracle:** the distribution F is given along with a best response oracle. The best response oracle accepts a distribution over the states as input and outputs the action that generates the highest expected utility w.r.t. the input distribution.
- **Implicit model with succinct description:** the distribution F is given along with the description of a Turing Machine that computes the ex-post utility of any pair of state and action in time polynomial in the description size.

We summarize our results with respect to the three different models: (i) we show that computing the optimal menu in the explicit model is captured by a polynomial size LP; (ii) even though the number of actions in the implicit model with BR oracle may be arbitrarily large, we provide an FPTAS to find an up-to- ε optimal menu for settings with a constant number of underlying states; (iii) we construct a succinctly representable instance of the problem such that computing a constant factor approximation or an up-to- ε optimal menu is NP-hard.

3.1 Explicit Model

In this section, we discuss the basic setting where the model is explicitly given. First, we state a structural result of the optimal menu that allows us to restrict the number of signals.

► **Definition 1** (Responsive Experiment). *A buyer type θ is responsive to an experiment E if every signal s of E leads to a different optimal choice of action and, in particular $a(s_k | \theta) = a_k$ for all $s_k \in S$.*

Note that if an experiment E is responsive to any buyer type, E has exactly m signals. An intuitive way to think about Definition 1 is that signal s_i of E recommends the buyer to take an action a_i , and type θ is responsive to E if θ always follows the recommendation. Importantly, a different type θ' may not be responsive to experiment E and does not follow the recommendations. We now state the structural result from [4] that states that it is without loss of generality to consider a menu where each buyer type purchases a responsive experiment.

► **Lemma 2** (Adapted from Proposition 1 from [4]). *A menu \mathcal{M} is responsive if for every buyer type θ , it chooses a responsive experiment $E(\theta)$ from the menu. We define the outcome of a menu as the joint distribution of states, actions, and monetary transfers resulting from every buyer type's optimal choice of experiment and subsequent choice of action. The outcome of every menu can be attained by a responsive menu.*

The proof of Lemma 2 is based on a revelation-principle type of argument. More precisely, assume that some type θ prefers an experiment $E(\theta)$ for which he is not responsive. Then, we can merge all the signals $s_{i_1}, \dots, s_{i_\ell}$ of $E(\theta)$ that lead the type to take the same action a_k and just send the merged signal as the new k -th signal s_k . Equipped with the structural result, we are ready to show how to capture the design of the optimal menu as a LP. Since

there exists an optimal menu that is responsive, every experiment in this menu has at most m signals. For every buyer type θ , we use $\pi(\theta)$ to denote the experiment type θ purchases, where $\pi_{\omega,i}(\theta)$ is the probability to send signal s_i when the state is ω . $\{\pi_{\omega,i}(\theta)\}_{\omega \in \Omega, i \in [m], \theta \in \Theta}$ is the first set of variables. We also have the prices for each experiment $\{t(\theta)\}_{\theta \in \Theta}$ as variables. The main difference between selling experiments and selling items is that different types may interpret the same experiment differently. More specifically, for a responsive experiment $(\pi(\theta), t(\theta))$, the optimal choice of action for type θ after receiving signal s_i is simply action a_i ; while for a different type θ' , the best action after receiving signal s_i can be a completely different action a_j due to a different induced posterior. As our LP is designed to compute the optimal and responsive menu, we need to guarantee that for any type θ , purchasing experiment $(\pi(\theta), t(\theta))$ and following the recommendation is better than purchasing any experiment $(\pi(\theta'), t(\theta'))$ for a different type θ' and subsequently choosing the optimal action based on the induced posterior. We remark that the straightforward formulation of the previous constraints requires one to consider all the possible mappings from signals to actions, resulting in a LP that has exponentially many constraints in m . Although this huge LP can be solved using the ellipsoid algorithm, we use a technique inspired by Chen et al. [15] that allows us to formulate it using only polynomially many constraints in m .

See Figure 1 for our LP. Besides the variables representing the experiments and prices, we introduce a new set of variables $\{z_i(\theta, \theta')\}_{i \in [m], \theta, \theta' \in \Theta}$. $z_i(\theta, \theta')$ serves as an upper bound of the conditional expected utility that θ gets when he considers misreporting as θ' and receives signal s_i , which is guaranteed by the second set of constraints. The first set of constraints make sure that for any type θ , the net utility of purchasing experiment $E(\theta)$ and following its recommendations is no worse than the utility of purchasing any other experiment and subsequently choosing the best action under each signal. Note that we allow θ' to be the same θ in the first set of constraints, which guarantees that the menu is responsive. The third set of constraints guarantees that purchasing experiment $(\pi(\theta), t(\theta))$ is no worse than θ 's base utility. We refer to these constraints as the individual rationality (IR) constraints. The last two constraints guarantee that $\pi(\theta)$ is indeed an experiment.

► **Theorem 3.** *The LP in Figure 1 can be solved in time polynomial in $m = |A|$, $|\Omega|$, and $|\Theta|$, and its optimal solution is the revenue-optimal menu.*

Proof of Theorem 3. The number of constraints and the number of variables of the LP in Figure 1 is polynomial in $m, |\Omega|, |\Theta|$, so it is clear that it can be solved in polynomial time. We now argue that its solution is indeed the optimal responsive menu, hence the optimal menu as guaranteed by Lemma 2. Observe that every responsive menu corresponds to a feasible solution of the LP. This is because if $\{(\pi(\theta), t(\theta))\}_{\theta \in \Theta}$ is a responsive menu, then we can satisfy all the constraints by setting $z_i(\theta, \theta') = \max_{a_j} \sum_{\omega} \theta_{\omega} \pi_{\omega,i}(\theta') u_{\omega, a_j}$. Conversely, note that every feasible solution of the LP induces a responsive menu that simply chooses the experiment $E(\theta)$ as $(\pi(\theta), t(\theta))$ for each type $\theta \in \Theta$. This is because (i) the first and second sets of constraints guarantee that $E(\theta)$ is a responsive experiment for θ (this follows from setting θ' to be θ in both the first and second set of constraints) and $V_{\theta}(E(\theta)) - t(\theta) \geq V_{\theta}(E(\theta')) - t(\theta')$; (ii) the IR constraints ensure that $V(E(\theta)) - t(\theta) \geq u(\theta)$. Hence, the optimal solution corresponds to the revenue-optimal menu. ◀

3.2 Implicit Model with a Best Response Oracle

In this section, we study the case where the model is provided implicitly. We prove that, given access to a best response oracle, our algorithm can compute an up-to- ε optimal menu in time $\text{poly}\left(|\Theta|, \frac{|\Omega|^{|\Omega|^2}}{\varepsilon^{|\Omega|^2 + |\Omega|}}\right)$, which is independent of the number of actions.

Variables:

- $\{\pi_{\omega,i}(\theta)\}_{\omega \in \Omega, i \in [m], \theta \in \Theta}$, denoting the experiments in the menu.
- $\{t(\theta)\}_{\theta \in \Theta}$, denoting the prices of the experiments.
- $\{z_i(\theta, \theta')\}_{i \in [m], \theta, \theta' \in \Theta}$, helper variables. $z_i(\theta, \theta')$ represents an upper bound of the conditional expected utility of signal s_i from experiment $E(\theta')$ for type θ .

Linear Program:

$$\begin{aligned}
& \max \sum_{\theta \in \Theta} F(\theta) t(\theta) \\
& \text{s.t.} \quad \sum_{i \in [m]} \sum_{\omega \in \Omega} \theta_{\omega} \pi_{\omega,i}(\theta) u_{\omega,a_i} - t(\theta) \geq \sum_{i \in [m]} z_i(\theta, \theta') - t(\theta'), \quad \forall \theta, \theta' \in \Theta \quad (\text{IC}) \\
& \quad z_i(\theta, \theta') \geq \sum_{\omega} \theta_{\omega} \pi_{\omega,i}(\theta') u_{\omega,a_j}, \quad \forall \theta, \theta' \in \Theta, \forall i, j \in [m] \\
& \quad \sum_{i \in [m]} \sum_{\omega \in \Omega} \theta_{\omega} \pi_{\omega,i}(\theta) u_{\omega,a_i} - t(\theta) \geq u(\theta), \quad \forall \theta \in \Theta \quad (\text{IR}) \\
& \quad \sum_{i \in [m]} \pi_{\omega,i}(\theta) = 1, \quad \forall \theta \in \Theta, \omega \in \Omega \\
& \quad \pi_{\omega,i}(\theta) \geq 0, \quad \forall \theta \in \Theta, \forall \omega \in \Omega, \forall i \in [m]
\end{aligned}$$

■ **Figure 1** A linear program to find the revenue-optimal menu in the explicit model.

We first prove a structural result which shows that no matter how large the actual action set A may be, there is some $A' \subseteq A$ such that a menu \mathcal{M} that recommends actions only from A' has negligible revenue loss compared to OPT. Moreover, $|A'|$ depends only on $|\Omega|$ and the additive approximation error. We also show that these sets can be computed efficiently.

► **Theorem 4.** *For any constant $\varepsilon > 0$, given access to a BR oracle we can compute for each type $\theta \in \Theta$ a set of actions A_{θ} by querying $O\left(\frac{|\Omega|^{|\Omega|^2}}{\varepsilon^{|\Omega|^2+|\Omega|}}\right)$ times the BR oracle, so that there exists an IC and IR menu \mathcal{M} , whose experiments all contain no more than $O\left(\frac{|\Omega|^{|\Omega|^2}}{\varepsilon^{|\Omega|^2+|\Omega|}}\right)$ many signals. Moreover, every type θ only uses actions from A_{θ} upon receiving any of these signals generated by his experiment. Finally, the revenue of \mathcal{M} is at least $\text{OPT} - O(\sqrt{\varepsilon})$.*

We present some lemmas that are used in the proof of Theorem 4. Firstly, we show that given a menu \mathcal{M} and ε , we can create a menu \mathcal{M}' with the following two properties: (i) the number of signals that \mathcal{M}' uses depends only on $|\Omega|$ and ε , and every type θ values the new experiment he gets at most $O(\varepsilon)$ less than his original experiment. The proof is based on the idea that merging signals of an experiment E that are close does not decrease the value of the experiment by much.

► **Lemma 5.** *Let $\varepsilon > 0$ be some given constant and let $E = \{E(\theta)\}_{\theta \in \Theta}$ be a set of experiments, where each $E(\theta)$ uses an arbitrary number of signals. Then, we can create a set of experiments $E' = \{E'(\theta)\}_{\theta \in \Theta}$ that uses $O\left(\left(\frac{|\Omega|}{\varepsilon}\right)^{|\Omega|}\right)$ signals per player such that $V_{\theta}(E'(\theta)) \geq V_{\theta}(E(\theta)) - 2\varepsilon$.*

We now argue that when we merge two signals, no type can value any experiment more than he did before.

▷ **Claim 6.** Let $E(\theta)$ be the initial experiment that is offered to type θ and $E'(\theta)$ the experiment that is offered to θ after the merge. Then, for any type $\theta' \in \Theta$ it holds that $V_{\theta}(E'(\theta')) \leq V_{\theta}(E(\theta'))$.

So far we have established the existence of a menu \mathcal{M}' whose number of signals is significantly smaller than the initial one. However, if we do not have access to \mathcal{M} we cannot compute \mathcal{M}' . Lemma 7 shows that we can overcome this issue. By “rounding” the entries of the experiment so that for any type the value this experiment generates does not change much. Now, since there is a small number of signals, and the size of every experiment depends only on ε and $|\Omega|$, we can do an exhaustive search over the discretized entries.

► **Lemma 7.** *Let $\{E(\theta)\}_{\theta \in \Theta}$ be a set of experiments, where each $E(\theta)$ uses signals from S . We also let $0 < \delta < 1$ be a given number such that $1/\delta \in \mathbb{N}$. Then, we can create a set of experiments $\{E'(\theta)\}_{\theta \in \Theta}$ that uses signals from S' with $|S'| = |S|$, such that $E'(\theta)$ is a valid experiment with $\pi_{\omega, s'}(E'(\theta))$ being a multiple of δ for all $\omega \in \Omega, s' \in S', \theta \in \Theta$, and $|V_{\theta}(E'(\theta')) - V_{\theta}(E(\theta'))| \leq \delta|S|$, for all $\theta, \theta' \in \Theta$.*

One construction that will be useful in our proofs is the ε -IC to IC transformation. Lemma 8 shows that if we have a menu whose IC, IR constraints are violated by at most ε , we can modify the prices so that it becomes IC, IR and has negligible $O(\sqrt{\varepsilon})$ revenue loss. The construction is based on a technique developed [20, 11] and frequently used in the Mechanism Design literature. In the single agent setting, the idea is to offer a small multiplicative discount to all types to make sure that if they want to deviate to some other experiment, this will not be much cheaper than the one they were buying in the initial ε -IC menu.

► **Lemma 8.** *Let $\mathcal{M} = \{E_i, t(E_i)\}_{i \in [k]}$ be a menu with k experiments. Suppose that the IC, IR constraints are violated by at most ε . Then, we can compute a new set of prices $\{\tilde{t}(E_i)\}_{i \in [k]}$ such that the menu $\tilde{\mathcal{M}} = \{E_i, \tilde{t}(E_i)\}_{i \in [k]}$ is IR, IC and $REV(\tilde{\mathcal{M}}) \geq (1 - \sqrt{\varepsilon})REV(\mathcal{M}) - \sqrt{\varepsilon} - \varepsilon$, in time $O(k)$.*

We are now ready to present a sketch of the proof for Theorem 4. Assume that we start with the optimal menu \mathcal{M}^* . By Lemma 5, Claim 6, and Lemma 7, we know that we can modify the experiments in \mathcal{M}^* so that they use only discretized signals. Moreover, the new menu is approximately IC and IR. We then apply Lemma 8 to obtain a menu that is IC and IR by sacrificing a negligible amount of revenue. Finally, to compute the collection of action sets that the types will choose after receiving the signals, we query BR oracle on all the possible discretized distributions where the signals are drawn from.

So far we have only shown a structural result about the existence of a menu that uses $O\left(\frac{|\Omega||\Omega|^2}{\varepsilon^{|\Omega|^2+|\Omega|}}\right)$ signals and gives an additive $O(\sqrt{\varepsilon})$ -approximation to the optimal revenue. We now argue that we can use a LP (Figure 2) to find such a menu. This is done by combining the result of Theorem 4 and modifying the LP we used in Section 3.1.

► **Theorem 9.** *For any $\varepsilon > 0$, any set Θ of types, given access to a BR oracle, we can use the LP in Figure 2 to compute a menu \mathcal{M} that achieves $REV(\mathcal{M}) \geq OPT - O(\sqrt{\varepsilon})$. The number of queries to the BR oracle is poly $\left(|\Theta|, \frac{|\Omega||\Omega|^2}{\varepsilon^{|\Omega|^2+|\Omega|}}\right)$ and the running time is poly $\left(|\Theta|, \frac{|\Omega||\Omega|^2}{\varepsilon^{|\Omega|^2+|\Omega|}}\right)$. Moreover, each experiment contains at most $O\left(\frac{|\Omega||\Omega|^2}{\varepsilon^{|\Omega|^2+|\Omega|}}\right)$ many signals, and each type θ only chooses actions from a set of actions A_{θ} with size $O\left(\frac{|\Omega||\Omega|^2}{\varepsilon^{|\Omega|^2+|\Omega|}}\right)$.*

So far, the number of experiments in our constructions depends on the number of types $|\Theta|$. We show that the number of experiments needed for an up-to- ε optimal menu is independent of $|\Theta|$. We achieve this by dropping experiments that are offered to types who are close in

Variables:

- $\{\pi_{\omega,i}(\theta)\}_{\omega \in \Omega, i \in [|A_\theta|], \theta \in \Theta}$, denoting the experiments in the menu.
- $\{t(\theta)\}_{\theta \in \Theta}$, denoting the prices of the experiments.
- $\{z_i(\theta, \theta')\}_{i \in [|A_{\theta'}|], \theta, \theta' \in \Theta}$, helper variables. $z_i(\theta, \theta')$ represents an upper bound of the conditional expected utility of signal s_i from experiment $E(\theta')$ for type θ .

Linear Program:

$$\begin{aligned}
& \max \quad \sum_{\theta \in \Theta} F(\theta) t(\theta) \\
& \text{s.t.} \quad \sum_{i \in [|A_\theta|]} \sum_{\omega \in \Omega} \theta_\omega \pi_{\omega,i}(\theta) u_{\omega, a_{\tau_\theta(i)}} - t(\theta) \geq \sum_{i \in [|A_{\theta'}|]} z_i(\theta, \theta') - t(\theta'), \quad \forall \theta, \theta' \in \Theta \quad (\text{IC}) \\
& \quad \quad z_i(\theta, \theta') \geq \sum_{\omega} \theta_\omega \pi_{\omega,i}(\theta') u_{\omega, a_j}, \quad \forall \theta, \theta' \in \Theta, \forall i \in [|A_{\theta'}|], \forall a_j \in A \\
& \quad \quad \sum_{i \in [|A_\theta|]} \sum_{\omega \in \Omega} \theta_\omega \pi_{\omega,i}(\theta) u_{\omega, a_{\tau_\theta(i)}} - t(\theta) \geq u(\theta), \quad \forall \theta \in \Theta \quad (\text{IR}) \\
& \quad \quad \sum_{i \in [|A_\theta|]} \pi_{\omega,i}(\theta) = 1, \quad \forall \theta \in \Theta, \omega \in \Omega \\
& \quad \quad \pi_{\omega,i}(\theta) \geq 0, \quad \forall \theta \in \Theta, \forall \omega \in \Omega, \forall i \in [|A_\theta|]
\end{aligned}$$

■ **Figure 2** A linear program to find an approximately revenue-optimal menu in the implicit model.

TV-distance. We show that this leads to a menu that preserves the revenue, is $O(|\Omega|\varepsilon)$ -IC, IR and has $O\left(\frac{|\Omega|^{2|\Omega|}}{\varepsilon^{|\Omega|}}\right)$ different experiments. Finally, we apply the ε -IC to IC transformation to the modified menu.

► **Lemma 10.** *Let $\mathcal{M} = \{E(\theta), t(\theta)\}_{\theta \in \Theta}$ be a menu of experiments. Then, for any $\theta, \theta' \in \Theta$ with $d_{TV}(\theta, \theta') \leq \varepsilon$ it holds that $V_\theta(E(\theta')) - t(\theta') \geq V_\theta(E(\theta)) - t(\theta) - 2|\Omega|\varepsilon$.*

We are now ready to prove that we can create a menu that offers a small number of experiments and loses negligible revenue compared to OPT. We do that in two steps, since we are dealing with an action space and type space that are arbitrary. The first step is to shrink the action space that we are considering. In order to do that, we use Theorem 4 that guarantees the existence of a menu which loses negligible revenue and only considers actions from smaller action spaces. The next step is to divide the state space into regions in which all the types are within ε in TV-distance. Lemma 10 shows us that if we consider offering a single experiment to all the types in the same region, their values for the new experiment will not change much compared to the one they were getting. Finally, we apply Lemma 8 to solve the issue that the menu resulting from dropping experiments might not be IC, IR.

► **Theorem 11.** *Consider an environment with a type space Θ , action space A and state space Ω . Then, given some $\varepsilon > 0$ and access to a BR oracle we can find a menu \mathcal{M} that generates revenue at least $\text{OPT} - O(\sqrt{\varepsilon})$ and offers at most $O\left(\frac{|\Omega|^{2|\Omega|}}{\varepsilon^{|\Omega|}}\right)$ experiments, in time $\text{poly}\left(|\Theta|, \frac{|\Omega|^{|\Omega|^2}}{\varepsilon^{|\Omega|^2+|\Omega|}}\right)$.*

3.3 Implicit Model with Succinct Description

In this section, we consider a setting where the model has a succinct implicit description. We show that no algorithm can obtain even a constant factor approximation to the optimal revenue for this setting, unless $P = NP$. To be more precise, we consider the following problem.

Information Pricing SAT (IP-SAT): find the revenue-optimal menu in the following setting:

- State space Ω , type space Θ .
- for each state ω , Φ_ω is a boolean formula in CNF over variables in $X = \{x_1, \dots, x_n\}$
- Action space A : all the possible truth assignments of the variables in X
- $u_{\omega,a} = \frac{\# \text{ satisfied clauses of } \Phi_\omega \text{ with assignment } a}{\# \text{ clauses of } \Phi_\omega}$

IP-SAT is a hard problem for the buyer in general, as maximizing his net utility requires solving an NP-hard problem. We show here that designing an approximately revenue-optimal menu for IP-SAT is also computationally intractable for the seller. Of course, it is not even clear what the optimal menu looks like in general for IP-SAT as we only have a limited characterization of the optimal menu. In Theorem 12, we construct a special family of IP-SAT instances with 2 states and 1 buyer type, and show how to reduce SAT to it.

► **Theorem 12.** *For any constant $\varepsilon > 0$, there does not exist a polynomial time algorithm \mathcal{A} that computes an menu with revenue at least $(1/2 + \varepsilon)OPT - \frac{\varepsilon}{2m+4}$ in the IP-SAT problem with $m + 2$ clauses, unless $P = NP$.*

Proof of Theorem 12. Let $\Phi = C_1 \wedge \dots \wedge C_m$ be any SAT instance over variables x_1, \dots, x_n . We show that given \mathcal{A} we can decide whether Φ is satisfiable. We create the following IP-SAT instance: there are two states ω_1, ω_2 , a single type $\theta = (\frac{1}{2}, \frac{1}{2})$ and we set $\Phi_{\omega_1} = (C_1 \vee y) \wedge \dots \wedge (C_m \vee y) \wedge (x_1 \vee y) \wedge (\neg x_1 \vee y)$, $\Phi_{\omega_2} = (C_1 \vee \neg y) \wedge \dots \wedge (C_m \vee \neg y) \wedge (x_1 \vee \neg y) \wedge (\neg x_1 \vee \neg y)$ to be the two SAT instances with $m + 2$ clauses that the buyer faces. At each state, the actions are the possible assignments of the variables x_1, \dots, x_n, y . Let Ψ be some boolean formula in CNF and a some assignment of its variables. We define $z_a(\Psi)$ to be the number of clauses in Ψ that a satisfies and $w(\Psi)$ the total number of clauses in Ψ . Then, at each state ω when the agent chooses assignment a we define his ex-post utility to be $u_{\omega,a} = \frac{z_a(\Phi_\omega)}{w(\Phi_\omega)}$.

From Bergemann et al. [4], we know the optimal menu should contain only the fully informative experiment E^* , and the price for this experiment is $V_\theta(E^*) - u(\theta)$.⁶ Clearly, $V_\theta(E^*) = 1$, because the buyer can pick y to be T and F in states ω_1, ω_2 respectively and satisfy all clauses. We now focus our attention on $u(\theta)$. Assume that without receiving any information the buyer decides to set $y = T$. This is w.l.o.g. since it is symmetric with the case he decides to set $y = F$. When the state is ω_1 , he satisfies all the clauses. According to his prior, this happens $1/2$ of the time, so we see that so far $u(\theta) \geq 1/2$. Let us consider which assignment he should pick when the state is ω_2 . Observe that no matter which value he picks for x_1 , he will always satisfy exactly one of the last two clauses in Φ_2 . Hence, for variables x_1, \dots, x_n he better pick the assignment a that maximizes $z_a(\Phi)$. Let $k = \max_a z_a(\Phi)$. Then $u(\theta) = 1/2 + (k + 1)/(2m + 4) = \frac{m+k+3}{2m+4}$. Hence, the optimal revenue is $V_\theta(E^*) - u(\theta) = 1 - \frac{m+k+3}{2m+4} = \frac{m-k+1}{2m+4}$. If Φ is satisfiable, then $k = m$, so $OPT = \frac{1}{2m+4}$. If Φ is not satisfiable then $k \leq m - 1$ so $OPT \geq \frac{2}{2m+4}$. Now assume that there is such an algorithm \mathcal{A} and denote $REV(\mathcal{A})$ the revenue generated by the mechanism output by \mathcal{A} . If Φ is not satisfiable, since $OPT \geq \frac{2}{2m+4}$, it must be that $REV(\mathcal{A}) > \frac{1}{2m+4}$. On the other hand, if Φ is not satisfiable we have that $REV(\mathcal{A}) \leq \frac{1}{2m+4}$. Hence, the existence of \mathcal{A} allows us to distinguish between satisfiable and unsatisfiable SAT formulas. ◀

⁶ The fully informative experiment simply sends out a signal to reveal the state.

► Remark 13. Since m is linear in the description length of the problem, given access to a BR oracle, for any $\varepsilon > 0$ we can set $\varepsilon' = c(m\varepsilon)^2$, for some appropriate constant $c > 0$, and then apply Theorem 9. Since $|\Omega| = 2$ and $|\Theta| = 1$, the revenue we get is at least $\text{OPT} - \varepsilon$ and the running time is $O\left(\frac{1}{(m\varepsilon)^{10}}\right)$. However, this does not contradict with the result of Theorem 12, since in this setting the BR oracle solves an NP-hard problem.

4 Multi-Agent Setting

In this section, we consider a multi-agent generalization of the model by Bergemann et al. [4]. More specifically, we assume that there are n buyers who are interested in acquiring extra information from the seller and each buyer's ex-post utility only depends on the state of the world and his own action. We further assume that the types of the buyers are drawn independently from their own type distributions. If there is no competition among them, the solution to the problem follows immediately from the single-agent setting, since the seller can offer each agent his optimal menu separately. Thus, we focus on a more interesting case where the buyers are competitors and only *one* of them can receive an informative signal.

Input Model and New Notation

We first need to introduce some new notation. We use Θ^i to denote the type space of buyer i and $F^i(\theta^i)$ to denote the probability that buyer i 's type is θ^i . We use Θ to denote the set of all type profiles and $F(\theta)$ to denote $\times_{i \in [n]} F^i(\theta^i)$. We assume the action space A is the same for each buyer i , but the ex-post utility $u_{\omega,a}^i$ for choosing action a under state ω may be different for different buyers. We consider the *explicit model*, that is, for each buyer i , both F^i and the ex-post utility matrix $U^i = \{u_{\omega,a}^i\}_{\omega \in \Omega, a \in A}$ are given as input. We use $u^i(\theta^i)$ to denote the base utility of buyer i for choosing the best action under distribution θ^i .

Interaction between the Seller and Buyers

The interaction happens in the following order:

1. The seller commits to a mechanism $\{(\Pi(\theta) = (\Pi^1(\theta), \dots, \Pi^n(\theta))), \{t(\theta) = (t^1(\theta), \dots, t^n(\theta))\}_{\theta \in \Theta}\}$, and announces the mechanism to all buyers.
2. The types of the buyers $\theta = (\theta^1, \dots, \theta^n)$ are realized.
3. Each buyer i **privately** submits his type θ^i to the seller.
4. The seller chooses buyer i as the winner with probability $p^i(\theta)$.
5. The seller observes the state of the world ω and sends buyer i a signal s according to the signaling scheme $\Pi^i(\theta)$ and charges buyer i price $t^i(\theta)$.
6. Each buyer i chooses an action a^i and receives ex-post utility u_{ω,a^i}^i .

There are some subtle issues in our model that require further clarification. The most important of them being the following. After the winner has been chosen, does he observe the signaling scheme $\Pi^i(\theta)$ that the seller uses to generate the signal s ? In this work, we consider the setting where the signaling scheme $\Pi^i(\theta)$ is **not revealed** and the winner only observes the realized signal.⁷ Some remarks are in order. Firstly, the seller may want to

⁷ One may worry that the winner can obtain extra information from the price $t^i(\theta)$. To avoid this, we will design a mechanism so that the price for buyer i only depends on i 's type θ^i . This is without loss of generality, as we can simply set the price to be $\mathbb{E}_{\theta^i}[t^i(\theta)]$.

preserve the privacy of the buyers, and revealing $\Pi^i(\theta)$ allows the winner i to infer the other buyers' priors. Secondly, hiding the implemented signaling scheme $\Pi^i(\theta)$ from the winner allows the seller to design a mechanism with less stringent IC constraints and thus generates higher revenue for the seller. This is because the winner does not know the exact experiment he is getting, if he wants to deviate from the recommendation he must map the same signal to the same action for all potential experiments that he may win. Therefore, he would map the signal to an action that induces the highest *expected utility*, where the expectation is over the other bidders' types and the chosen experiment. On the other hand, if the buyer knew which experiment the signal is drawn from, he could use a mapping that is the best for each particular experiment.

Our goal in this section is to design a polynomial time algorithm to find the Bayesian Incentive Compatible (BIC) and Interim Individually Rational (IIR) mechanism that achieves the highest revenue among all BIC and IIR mechanisms for our model. It is not hard to see that Lemma 2 generalizes to our multi-agent setting. We begin by introducing an extension of the LP in Figure 1 to the multi-agent setting. Define $\Pi^i(\theta) = (\pi_{\omega,j}^i(\theta))_{\omega \in \Omega, j \in [m]}$, $p^i(\theta)$, and $t^i(\theta)$ as the decision variables for every buyer i and type profile θ . Recall that $m = |A|$.

$$\begin{aligned}
& \max \sum_{\theta \in \Theta} F(\theta) \sum_{i \in [n]} t^i(\theta) \\
& \text{s.t.} \quad \sum_{\theta^{-i}} F^{-i}(\theta^{-i}) \left(\sum_{\substack{\omega \in \Omega, \\ j \in [m]}} \theta_{\omega}^i \pi_{\omega,j}^i(\theta^i, \theta^{-i}) u_{\omega, a_j}^i + (1 - p^i(\theta)) u^i(\theta^i) - t^i(\theta^i, \theta^{-i}) \right) \geq \\
& \quad \sum_{\theta^{-i}} F^{-i}(\theta^{-i}) \left(\sum_{j \in [m]} z_j^i(\theta^i, \tilde{\theta}^i, \theta^{-i}) + (1 - p^i(\tilde{\theta}^i, \theta^{-i})) u^i(\theta^i) - t^i(\tilde{\theta}^i, \theta^{-i}) \right), \quad \forall i, \forall \theta^i, \tilde{\theta}^i \text{ (BIC)} \\
& \quad \sum_{\theta^{-i}} F^{-i}(\theta^{-i}) z_j^i(\theta^i, \tilde{\theta}^i, \theta^{-i}) \geq \sum_{\theta^{-i}} F^{-i}(\theta^{-i}) \sum_{\omega} \theta_{\omega}^i \pi_{\omega,j}^i(\tilde{\theta}^i, \theta^{-i}) u_{\omega, a_k}, \quad \forall i, \forall j, k, \forall \theta^i, \tilde{\theta}^i \\
& \quad \sum_{\theta^{-i}} F^{-i}(\theta^{-i}) \left(\sum_{\substack{\omega \in \Omega, \\ j \in [m]}} \theta_{\omega}^i \pi_{\omega,j}^i(\theta) u_{\omega, a_j} + (1 - p^i(\theta)) u^i(\theta^i) - t^i(\theta) \right) \geq u^i(\theta^i), \quad \forall i, \theta^i \text{ (IIR)} \\
& \quad \sum_{j \in [m]} \pi_{\omega,j}^i(\theta) = p^i(\theta), \quad \forall i, \forall \theta, \forall \omega \text{ (feasibility)} \\
& \quad \sum_{i \in [n]} p^i(\theta) \leq 1, \quad \forall \theta \text{ (feasibility)} \\
& \quad \pi_{\omega,j}^i(\theta) \geq 0, \quad \forall i, \forall \theta, \forall \omega, \forall j \text{ (feasibility)}
\end{aligned}$$

Observe that the number of variables is exponential in n and the number of constraints is exponential in both n and m . There is no hope to solve this LP in polynomial time. The main challenge is how to remove the exponential dependence on n . To overcome this obstacle, we use a method that is powerful in the study of multi-item auctions, that is, rewriting the LP using a more succinct representation of the mechanism known as the reduced form [5, 1, 6, 7, 8]. We first define the reduced form of a mechanism.

► **Definition 14 (Reduced Form).** *Given a mechanism $\mathcal{M} = (\{\Pi(\theta)\}_{\theta \in \Theta}, \{t(\theta)\}_{\theta \in \Theta})$, we define its reduced form $\{\hat{\Pi}^i(\theta^i)\}_{i \in [n], \theta^i \in \Theta^i}$, where $\hat{\Pi}^i(\theta^i) = \{\hat{\pi}_{\omega,j}^i(\theta^i)\}_{\omega \in \Omega, j \in [m]}$ and $\hat{\pi}_{\omega,j}^i(\theta^i) = \mathbb{E}_{\theta^{-i}}[\pi_{\omega,j}^i(\theta)]$ for each state ω and $j \in [m]$, and its interim prices $\{\hat{t}^i(\theta^i)\}_{i \in [n], \theta^i \in \Theta^i}$, where $\hat{t}^i(\theta^i) = \mathbb{E}_{\theta^{-i}}[t^i(\theta)]$. We use $\mathcal{P}(F)$ to denote the set of all reduced forms for a particular type distribution F .*

It is not hard to see that $\mathcal{P}(F)$ is a closed convex set, as the set of all mechanisms is clearly closed and convex, and $\mathcal{P}(F)$ is simply a linear transformation of that set. Intuitively, the reduced form is the “expected experiment and price” that the each buyer believes he will be allocated when his type is realized, and the expectation is taken over the randomness of the other buyers’ types.

► **Lemma 15.** *For any type distribution F , $\mathcal{P}(F)$ is a closed convex set.*

The LP in Figure 3 searches for the reduced form of the revenue-optimal mechanism. Notice that the size of the reduced-form LP is substantially smaller than the original one, and the number of variables is polynomial in the number of agents. With these new variables we can still express the BIC and IIR constraints of the initial LP. However, it is not yet clear how to check whether these variables correspond to an actual feasible mechanism. In the next section, we show how to design a separation oracle that checks the feasibility efficiently.

Variables:

- $\{\hat{\pi}_{\omega,j}^i(\theta^i)\}_{\omega \in \Omega, i \in [n], \theta^i \in \Theta^i, j \in [m]}$, denoting the reduced form of the mechanism.
- $\{\hat{t}^i(\theta^i)\}_{i \in [n], \theta^i \in \Theta^i}$, denoting the interim prices.
- $\{\hat{p}^i(\theta^i)\}_{i \in [n], \theta^i \in \Theta^i}$, denoting the allocation probabilities of the experiment
- $\{\hat{z}_j^i(\theta^i, \tilde{\theta}^i)\}_{i \in [n], j \in [m], \theta^i, \tilde{\theta}^i \in \Theta^i}$, helper variables. $\hat{z}_j^i(\theta^i, \tilde{\theta}^i)$ represents an upper bound of the conditional expected utility of signal s_j for type θ^i .

Linear Program:

$$\begin{aligned}
& \max && \sum_{i \in [n]} \sum_{\theta^i \in \Theta^i} F^i(\theta^i) \hat{t}^i(\theta^i) \\
& \text{subject to} && \sum_{j \in [m]} \sum_{\omega \in \Omega} \theta_{\omega}^i \hat{\pi}_{\omega,j}^i(\theta^i) u_{\omega, a_j}^i + (1 - \hat{p}^i(\theta^i)) u^i(\theta^i) - \hat{t}^i(\theta^i) \geq \\
& && \sum_{j \in [m]} \hat{z}_j^i(\theta^i, \tilde{\theta}^i) + (1 - \hat{p}^i(\tilde{\theta}^i)) u^i(\theta^i) - \hat{t}^i(\tilde{\theta}^i), && \forall i, \forall \theta^i, \tilde{\theta}^i \text{ (BIC)} \\
& && \hat{z}_j^i(\theta^i, \tilde{\theta}^i) \geq \sum_{\omega \in \Omega} \theta_{\omega}^i \hat{\pi}_{\omega,j}^i(\tilde{\theta}^i) u_{\omega, a_k}^i && \forall i, \forall \theta^i, \tilde{\theta}^i, \forall j, k \\
& && \sum_{j \in [m]} \sum_{\omega \in \Omega} \theta_{\omega}^i \hat{\pi}_{\omega,j}^i(\theta^i) u_{\omega, a_j}^i + (1 - \hat{p}^i(\theta^i)) u^i(\theta^i) - \hat{t}^i(\theta^i) \geq u^i(\theta^i), && \forall i, \forall \theta^i \text{ (IIR)} \\
& && \hat{p}^i(\theta^i) = \sum_{j \in [m]} \hat{\pi}_{\omega,j}^i(\theta^i) && \forall i, \forall \omega, \forall \theta^i \\
& && \{\hat{\pi}_{\omega,j}^i(\theta^i)\}_{\omega \in \Omega, i \in [n], \theta^i \in \Theta^i, j \in [m]} \in \mathcal{P}(F) && \text{(Feasibility)}
\end{aligned}$$

■ **Figure 3** A linear program to find the reduced form of the revenue-optimal mechanism in the multi-agent setting.

4.1 Feasibility of Reduced Forms

To design a separation oracle for the set $\mathcal{P}(F)$, we invoke the equivalence between Optimization and Separation in Linear Programming [29, 33], which states that being able to optimize any linear function over a convex set P is equivalent to having a separation oracle for P . It is a well-known fact that given a separation oracle for P one can optimize any linear function using the ellipsoid method. Interestingly, the reverse is also true. If there is an algorithm to optimize any linear function over P , one can construct a separation oracle for P using the ellipsoid method. We state a strengthened version of the equivalence due to Cai et al. [7].

81:16 How to Sell Information Optimally: An Algorithmic Study

The reason that we need to be able to decompose a feasible point into corners of the polytope is that we eventually need to be able to implement the reduced forms as a feasible mechanism. We elaborate more on this later.

► **Theorem 16** (Adapted from Theorem H.1 of [7]). *Let P be a d -dimensional closed convex region, and let \mathcal{A} be any polynomial-time algorithm that takes any direction $w \in \mathbb{R}^d$ as input and outputs the extreme point $\mathcal{A}(w) \in P$ in direction w such that $\mathcal{A}(w) \cdot w \geq \max_{x \in P} x \cdot w$. Then we can design a polynomial time separation oracle SO for P such that, whenever $SO(x) = \text{“yes”}$, the execution of SO explicitly finds directions w_1, \dots, w_k such that x lies in the convex hull of $\{\mathcal{A}(w_1), \dots, \mathcal{A}(w_k)\}$.*

To apply this equivalence, we need to show how we can optimize a linear function over the set of feasible reduced-form variables. Recall that we use $\Pi^i(\theta) = (\pi_{\omega, j}^i(\theta))_{\omega \in \Omega, j \in [m]}$ and $\hat{\Pi}^i(\theta^i) = (\hat{\pi}_{\omega, j}^i(\theta^i))_{\omega \in \Omega, j \in [m]}$ to denote the ex-post signaling scheme and its reduced form. We will treat $\Pi^i(\theta)$ and $\hat{\Pi}^i(\theta^i)$ as $m|\Omega|$ -dimensional vectors. The following maximization problem plays a crucial role in our approach.

► **Definition 17.** *Consider any type profile θ . Let $\{X^i(\theta^i)\}_{i \in [n], \theta^i \in \Theta_i}$ be a collection of $m|\Omega|$ -dimensional vectors. We define a Virtual Payoff Maximizer (VPM) w.r.t. these weight vectors $VPM(\{X^i(\theta^i)\}_{i \in [n], \theta^i \in \Theta_i})$ to be the ex-post signaling scheme $\Pi(\theta)$ that maximizes the following quantity $\sum_i \Pi^i(\theta^i) \cdot X^i(\theta^i)$ for every type profile θ . The corresponding reduced form $\hat{\Pi}(\theta)$ is called $rVPM(\{X^i(\theta^i)\}_{i \in [n], \theta^i \in \Theta_i})$. In order to ensure that the maximizer is unique, we break ties lexicographically.*

When there is no confusion, we also write $VPM(w), rVPM(w)$ as the maximizers for the weight vector w .

► **Lemma 18.** *Given an arbitrary collection of weights $\{X^i(\theta^i)\}_{i \in [n], \theta^i \in \Theta_i}$ we can find the exact optimal solution of*

$$\max_{\hat{\Pi} \in \mathcal{P}(F)} \sum_{i \in [n]} \sum_{\theta^i \in \Theta_i} \hat{\Pi}^i(\theta^i) \cdot X^i(\theta^i) \text{ in time } O\left(m|\Omega| \left(\sum_{i \in [n]} |\Theta_i|\right) + \left(\sum_{i \in [n]} |\Theta_i|\right)^2\right).$$

Proof of Lemma 18. We first rewrite the maximization problem

$$\begin{aligned} \max_{\hat{\Pi} \in \mathcal{P}(F)} \sum_{i \in [n]} \sum_{\theta^i \in \Theta_i} \hat{\Pi}^i(\theta^i) \cdot X^i(\theta^i) &= \max_{\Pi} \sum_{i, \theta^i, \theta^{-i}} F^{-i}(\theta^{-i}) \Pi^i(\theta^i, \theta^{-i}) \cdot X^i(\theta^i) = \\ &= \max_{\Pi} \sum_{i, \theta} F(\theta) \Pi^i(\theta) \cdot \frac{X^i(\theta^i)}{F^i(\theta^i)} = \max_{\Pi} \sum_{\theta} F(\theta) \sum_i \Pi^i(\theta) \cdot \tilde{X}^i(\theta^i), \end{aligned}$$

where $\tilde{X}^i(\theta^i) = \frac{X^i(\theta^i)}{F^i(\theta^i)}$.

Let θ be a type profile. We now characterize the solution of $\max_{\Pi(\theta)} \sum_i \Pi^i(\theta) \cdot \tilde{X}^i(\theta^i)$. If we allocate the experiment to buyer i , the maximum value we can derive is $v^i(\theta^i) = \sum_{\omega} \max_a \tilde{X}_{\omega, a}^i(\theta^i)$. Clearly, the optimal solution of the linear function above is to always allocate the experiment to the buyer with the largest $v^i(\theta^i)$.

The ex-post signaling scheme that maximizes $\sum_{\theta} F(\theta) \sum_i \Pi^i(\theta) \cdot \tilde{X}^i(\theta^i)$ is the one that always allocates the experiment to the buyer with the largest $v^i(\theta^i)$ for every type profile θ . To solve $\max_{\hat{\Pi} \in \mathcal{P}(F)} \sum_{i \in [n]} \sum_{\theta^i \in \Theta_i} \hat{\Pi}^i(\theta^i) \cdot X^i(\theta^i)$, we only need to calculate the reduced form of this ex-post signaling scheme and we denote it using $\hat{\Pi}_*$.

We first compute $v^i(\theta^i) = \sum_{\omega} \max_a x_{\omega,a}^i(\theta^i)$ for every buyer i and every type θ^i . This step takes time $O\left(m|\Omega| \left(\sum_{i \in [n]} |\Theta^i|\right)\right)$, and there are $\sum_{i \in [n]} |\Theta^i|$ different such values. Next, for each buyer, we sort $v^i(\theta^i)$. This step takes time $O\left(\sum_{i \in [n]} |\Theta^i| \log |\Theta^i|\right)$. To compute $\hat{\Pi}_*^i(\theta^i)$, we only need to calculate the probability of the event that over the random draws of θ_{-i} , there exists another buyer $\ell \neq i$ either $v_{\ell}(\theta^{\ell}) > v^i(\theta^i)$ or $\ell < i$ and $v_{\ell}(\theta^{\ell}) = v^i(\theta^i)$. This probability can be computed in time $O\left(\sum_{i \in [n]} |\Theta^i|\right)$ for each buyer i and type θ^i . Hence, in total we can optimize the linear function in time $O\left(m|\Omega| \left(\sum_{i \in [n]} |\Theta^i|\right) + \left(\sum_{i \in [n]} |\Theta_i|\right)^2\right)$. ◀

Combining Theorem 16 and Lemma 18, we have a polynomial time algorithm to solve the LP in Figure 3, but we still need to turn the reduced form into an ex-post signaling scheme. We again use an idea from computing the optimal multi-item auctions, that is, first decomposing the optimal reduced form into a distribution over extreme points of $\mathcal{P}(F)$, then implementing all the extreme points that appear in the distribution using a VPM ex-post signaling scheme.

► **Theorem 19.** *We design an algorithm to compute the revenue-optimal mechanism in time $\text{poly}\left(n, m, |\Omega|, \sum_{i \in [n]} |\Theta^i|\right)$. Moreover, the mechanism can be implemented as a distribution over $m|\Omega| \left(\sum_{i \in [n]} |\Theta^i|\right) + 1$ VPM ex-post signaling schemes.*

5 Further Extensions and Future Directions

In this section, we discuss further generalizations of the model by Bergemann et al. [4] and future research directions that we believe are interesting to pursue.

5.1 Extensions of the Original Model

Enlarged Buyer Type

Recall that in the original model the only private information of the buyer is his private belief of the underlying state, which is realized at the beginning of the interaction with the seller. Importantly, the payoffs are public knowledge and remain the same across different buyers. A natural generalization one can consider is to allow the buyer to draw not only his prior belief θ , but also his payoff function $u : \Omega \times A \rightarrow [0, 1]$ from some distribution.

To be more specific, we consider the setting where a buyer's type $\rho = (\theta, u)$ is drawn from some distribution F at the beginning of the interaction between the buyer and the seller, and ρ is private to the buyer. As in the original model, we assume that the seller has access to this distribution. We remark that all of our positive results from Section 3 and Section 4, except for Theorem 11, hold in this extended model as well. The only difference in our constructions is that instead of indexing the variables by θ we now index them by ρ .

Misspecified Model

Another generalization we consider in the single-agent setting is the *misspecified* model. In this model the seller has access to some type distribution \tilde{F} which is within ε in TV-distance with the real type distribution F . Moreover, the seller has access to a type space $\tilde{\Theta}$ with the following two properties: $|\tilde{\Theta}| = |\Theta|$, for all $\tilde{\theta} \in \tilde{\Theta}$ there is some $\theta \in \Theta$ for which $d_{TV}(\tilde{\theta}, \theta) \leq \varepsilon$. Then, the menu that the seller designs for the misspecified distributions can be modified so that it guarantees only a negligible revenue loss when it is evaluated in the true setting. Lemma 20 formalizes this claim.

► **Lemma 20.** *Let \tilde{F}, F be the distributions of the types that the seller has access to and the true distribution of the types, respectively. Let also $\{\tilde{\theta}_i\}_{i \in [k]}, \{\theta_i\}_{i \in [k]}$ be the types that the seller has access to and the true types, respectively. Assume that $d_{TV}(\tilde{F}, F) \leq \varepsilon_1, d_{TV}(\tilde{\theta}_i, \theta_i) \leq \varepsilon_2, \forall i \in [k]$. We also let $\tilde{\mathcal{M}} = \{\tilde{E}(\tilde{\theta}_i), \tilde{t}(\tilde{\theta}_i)\}_{i \in [k]}$ be an IC, IR menu that has revenue $REV(\tilde{\mathcal{M}})$ under the misspecified distributions and uses at most $|S|$ signals. Then, we can compute a set of prices $\{t(\tilde{\theta}_i)\}_{i \in [k]}$ so that $\mathcal{M} = \{E(\theta_i), t(\theta_i)\}_{i \in [k]}$ is IR, IC and has $REV(\mathcal{M}) \geq REV(\tilde{\mathcal{M}}) - O\left(\varepsilon_1 + \sqrt{|\Omega|\varepsilon_2}\right)$ under the true distributions.*

Note that Lemma 20 allows us to generalize our results and obtain approximately-optimal menus when we only have black-box access to the distribution of the types. That is, we take enough samples to learn the distribution within Total Variation distance ε and then apply Lemma 20.

5.2 Future Directions

We believe that the design of Information Markets is a very important problem that has not received sufficient attention by the Theory of Computation community. There are many interesting questions waiting to be addressed.

1. In the single-buyer setting where we only have access to the action space via a BR oracle the running time of our algorithms is exponential in the number of states. An immediate question to ask is whether we can get an FPTAS or even a PTAS that has a better dependence on the number of states.
2. In the multi-agent setting, we consider the case in which the seller does not reveal the signaling scheme that she uses to send a signal to the winner. An interesting question is whether we can have efficient algorithms in the setting where the seller reveals the signaling scheme to the buyer.
3. Currently, in the multi-agent setting we assume that the ex-post utility of each buyer depends only on the state of the world and the action he takes. Is the problem of designing the optimal mechanism when the ex-post utilities also depend on the actions of the other buyers tractable?

References

- 1 Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. Bayesian Optimal Auctions via Multi- to Single-agent Reduction. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012.
- 2 Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014.
- 3 Moshe Babaioff, Robert Kleinberg, and Renato Paes Leme. Optimal mechanisms for selling information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 92–109, 2012.
- 4 Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American economic review*, 108(1):1–48, 2018.
- 5 Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. An Algorithmic Characterization of Multi-Dimensional Mechanisms. In *the 44th Annual ACM Symposium on Theory of Computing (STOC)*, 2012.
- 6 Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal Multi-Dimensional Mechanism Design: Reducing Revenue to Welfare Maximization. In *the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2012.

- 7 Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Reducing Revenue to Welfare Maximization : Approximation Algorithms and other Generalizations. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013.
- 8 Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013.
- 9 Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *the 48th Annual ACM Symposium on Theory of Computing (STOC)*, 2016.
- 10 Yang Cai, Federico Echenique, Hu Fu, Katrina Ligett, Adam Wierman, and Juba Ziani. Third-party data providers ruin simple mechanisms. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(1):1–31, 2020.
- 11 Yang Cai, Argyris Oikonomou, Grigoris Velezkas, and Mingfei Zhao. An efficient epsilon-bic to bic transformation and its application to black-box reduction in revenue maximization. In *the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2021. [arXiv:1911.10172](#).
- 12 Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 170–183, 2017. [doi:10.1145/3055399.3055465](#).
- 13 Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*, 2007.
- 14 Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010.
- 15 Yiling Chen, Haifeng Xu, and Shuran Zheng. Selling information through consulting. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2412–2431. SIAM, 2020.
- 16 Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. Mixture selection, mechanism design, and signaling. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1426–1445. IEEE, 2015.
- 17 Jacques Cremer and Richard P. McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.
- 18 Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. *Econometrica*, 85(3):735–767, 2017.
- 19 Constantinos Daskalakis, Christos Papadimitriou, and Christos Tzamos. Does information revelation improve revenue? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 233–250. ACM, 2016.
- 20 Constantinos Daskalakis and S. Matthew Weinberg. Symmetries and Optimal Multi-Dimensional Mechanism Design. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012.
- 21 Nikhil Devanur, Kira Goldner, Raghuvansh Saxena, Ariel Schwartzman, and S Matthew Weinberg. Optimal mechanism design for single-minded agents. *arXiv preprint*, 2020. [arXiv:2002.06329](#).
- 22 Nikhil R Devanur and S Matthew Weinberg. The optimal mechanism for selling to a budget constrained buyer: The general case. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 39–40, 2017.
- 23 Shaddin Dughmi. On the hardness of signaling. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 354–363. IEEE, 2014.
- 24 Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. Constrained signaling in auction design. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1341–1357. Society for Industrial and Applied Mathematics, 2014.

- 25 Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 412–425. ACM, 2016.
- 26 Amos Fiat, Kira Goldner, Anna R Karlin, and Elias Koutsoupias. The fedex problem. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 21–22, 2016.
- 27 Yiannis Giannakopoulos and Elias Koutsoupias. Duality and optimality of auctions for uniform distributions. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 259–276, 2014. doi:10.1145/2600057.2602883.
- 28 Yiannis Giannakopoulos and Elias Koutsoupias. Selling two goods optimally. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, pages 650–662, 2015. doi:10.1007/978-3-662-47666-6_52.
- 29 Martin Grötschel, László Lovász, and Alexander Schrijver. The Ellipsoid Method and its Consequences in Combinatorial Optimization. *Combinatorica*, 1(2):169–197, 1981.
- 30 Nima Haghpanah and Jason D. Hartline. Reverse mechanism design. *CoRR*, abs/1404.1341, 2014. arXiv:1404.1341.
- 31 Sergiu Hart and Noam Nisan. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012.
- 32 Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- 33 Richard M. Karp and Christos H. Papadimitriou. On linear characterizations of combinatorial optimization problems. *SIAM J. Comput.*, 11(4):620–632, 1982. doi:10.1137/0211053.
- 34 Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.
- 35 Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- 36 John Riley and Richard Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, 98(2):267–289, 1983.
- 37 Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *SODA*, 2015. arXiv:1406.3278.