

Selling Information Through Consulting*

Yiling Chen
Harvard University
yiling@seas.harvard.edu

Haifeng Xu
University of Virginia
hx4ad@virginia.edu

Shuran Zheng
Harvard University
shuran_zheng@seas.harvard.edu

Abstract

We consider a monopoly information holder selling information to a *budget-constrained* decision maker, who may benefit from the seller’s information. The decision maker has a utility function that depends on his action and an uncertain *state of the world*. The seller and the buyer each observe a private signal regarding the state of the world, which may be correlated with each other. The seller’s goal is to sell her private information to the buyer and extract maximum possible revenue, subject to the buyer’s budget constraints. We consider three different settings with increasing generality, i.e., the seller’s signal and the buyer’s signal can be independent, correlated, or follow a general distribution accessed through a black-box sampling oracle. For each setting, we design information selling mechanisms which are both optimal and simple in the sense that they can be naturally interpreted, have succinct representations, and can be efficiently computed. Notably, though the optimal mechanism exhibits slightly increasing complexity as the setting becomes more general, all our mechanisms share the same format of acting as a *consultant* who recommends the best action to the buyer but uses different and carefully designed payment rules for different settings. Each of our optimal mechanisms can be easily computed by solving a single polynomial-size linear program. This significantly simplifies exponential-size LPs solved by the Ellipsoid method in the previous work, which computes the optimal mechanisms in the same setting but *without* budget limit. Such simplification is enabled by our new characterizations of the optimal mechanism in the (more realistic) budget-constrained setting.

1 Introduction

Recent years have seen a growth of information markets thanks to the tremendous increase in the volume and variety of the online data sources. The information traded includes consumer information (e.g. Acxiom, Nielsen, Oracle), credit reports (e.g. Experian, TransUnion, FICO), recommendations (e.g. Waze, Tripadvisor), etc., which can be valuable to decision makers like advertisers, retailers and loan providers.

Motivated by these applications, this paper considers a very basic setting along this line and studies how a monopoly information provider can sell information optimally to a single decision maker. The decision maker’s utility depends on the action he takes (e.g. to lend the money or not) and a random *state of the world*, which is uncertain to the decision maker (e.g., whether the borrower will pay the debt). The information provider and the decision maker each observe a private signal regarding the state of the world. Their signals can be correlated to one another. The information provider’s goal is to sell her private signal to the decision maker in a way that maximizes the revenue, knowing that the buyer will choose what is best for himself.

This problem has been formulated and studied by [Babaioff et al. \(2012\)](#). They identified conditions under which there is an optimal one-round mechanism and provide characterizations of the optimal mechanism. In particular, they show that a one-round optimal mechanism exists in two special settings: (1) the seller’s signal and the buyer’s signal are independent, *or* (2) the buyer is *committed* to complete the entire protocol even if aborting the protocol will give him higher utility. The buyer commitment assumption is not as unrealistic as it might first sound. In particular, [Babaioff et al. \(2012\)](#) show that buyer commitment can be obtained using the following approach: first ask the buyer to deposit a large amount of money, then run the optimal mechanism for committed buyers, and at last refund the buyer his deposit less the payment in the mechanism. The deposit step enforces the buyer’s commitment to the payment of the mechanism.¹ However, the major drawback of this approach — as also highlighted by [Babaioff et al. \(2012\)](#) — is that both the amount of deposit and the money transfer required by the mechanism can be extremely large compared to players’ expected utilities. For in-

*Supported by the National Science Foundation under Grant No. CCF-1718549.

¹This is crucial in selling information since the payment amount itself can be correlated with the information for sale. Thus the buyer can learn information upon knowing the payment amount but before deciding to purchase. In such cases, an uncommitted buyer can simply learn the payment amount and then leave because the payment already reveals useful information to him.

stance, they give a simple example with buyer utility function taking values from $[0, 5]$. However, the optimal mechanism has to ask the buyer to deposit an amount of 25004 and then either returns 0 or returns 50000 to the buyer, yielding an optimal expected seller revenue no more than 2. Clearly, such a mechanism would *not* be practical in reality.

Given the above deficiencies, this work considers a more realistic setting with a *budget-constrained* buyer. In particular, the buyer in our model has a budget, indicating the maximum amount he can afford. We consider both *public* budget which is known to the seller a priori, and *private* budget which the seller also needs to elicit in the mechanism. In the latter case, the buyer may misreport his private budget if this is beneficial to him. We assume that the seller also has a budget as the maximum amount she can afford to pay the buyer (as seen from the previous example, the optimal mechanism sometimes requires the seller to pay the buyer). We consider three settings with increasing generality — i.e., independent signals, correlated signals, and signals drawn from a general distribution accessed through a black-box sampling oracle — and for each setting we are able to design information selling mechanisms which are both optimal (approximately optimal when signal distributions are accessed through samples) and simple in the sense that they can be naturally interpreted, have succinct representations, and can be efficiently computed. Notably, in contrast to the optimal mechanisms by Babaioff et al. (2012) which are formulated as exponential-size linear programs and solved by using the ellipsoid method on the dual LP, all our mechanisms can be directly computed by solving a single polynomial-size linear program thanks to our new and succinct characterization of the optimal mechanisms. In some sense, our results can be viewed as generalizations of the results by Babaioff et al. (2012) since when the payments in their mechanisms are smaller than the budget limits, our mechanisms (though simpler) also achieve optimality as their mechanisms whereas when the payments in their mechanisms are too large, our mechanisms are more realistic by imposing payment limits. Next, we elaborate our results in more details.

1.1 Our Results When the seller’s signal and the buyer’s signal are independent and the buyer’s budget is public, we prove that the optimal mechanism has the following simple format: (i) it asks the buyer to report his signal; (ii) based on the report, the seller charges the buyer a fixed amount of payment; (iii) the seller maps (possibly randomly) her signal to an action recommendation to the buyer. The mechanism has the following properties: (1) it is individually rational

(IR) and incentive compatible (IC); (2) each action recommended to the buyer is indeed his best response, which we call the *obedience* constraint; (3) the optimal mechanism can be computed directly from solving a polynomial-size linear program (LP). This mechanism is very much like the process of consulting, during which a client (the buyer) tells his type to the consultant (the seller) and pays the consultant for receiving the best action recommendation. As a result, we term this mechanism the *consulting mechanism with direct payment* (*CM-dirP*).

When the buyer budget is private, we show via an example that any mechanism with a single-round transfer *cannot* be optimal even when the buyer’s signal and the seller’s signal are independent. Nevertheless, for independent signals, we prove that the consulting mechanism with a slightly different payment method, coined the *consulting mechanism with deposit and return* (*CM-depR*), achieves the optimal revenue. In particular, instead of asking for a direct payment as in *CM-dirP*, the *CM-depR* mechanism first asks the buyer to deposit his (private) budget and then refund the buyer his budget less the payment (thus two-rounds of transfers). Besides satisfying IR, IC and obedience, the optimal *CM-depR* mechanism can also be computed by a polynomial-size LP.

For the general setting with correlated signals drawn from an explicitly given prior distribution, we prove that for both public and private budget cases, the consulting mechanism with a simple two-round transfer process, coined the *consulting mechanism with probabilistic return* (*CM-probR*), achieves the optimal revenue. The *CM-probR* has the following format: (i) it asks the buyer to report his signal and deposit his budget; (ii) based on the report, the mechanism maps the seller signal (randomly) to an action recommendation together with an amount of refund, which is either 0 (i.e., no refund) or the buyer’s full deposit plus the seller’s full budget. The main difference between *CM-probR* and the previous consulting mechanisms is that the payment amount is no longer pre-determined before any information disclosure, but is contingent upon the recommended action and thus contains information regarding the seller’s private signal. Similarly, besides satisfying IR, IC and obedience, the optimal *CM-probR* mechanism can also be computed by a polynomial-size LP. Our proof of the optimality and the computability of *CM-probR* mechanisms features a novel use of the duality theory, which may be of independent interest. Our starting point is an LP formulation \mathcal{P} for computing the optimal mechanism which however has exponentially many constraints *and* exponentially many variables. We simplify this LP as follows. We first turn

to the dual LP \mathcal{D} and leverage its special structure to eliminate a majority of the dual variables and convert \mathcal{D} into another linear program \mathcal{D}' with only polynomially many variables though still exponentially many constraints. At this point, we can employ the celebrated ellipsoid method to solve \mathcal{D}' , which however is neither practically efficient nor able to give us an interpretable mechanism. We instead turn to the dual of \mathcal{D}' , denoted as \mathcal{P}' , which is different from \mathcal{P} . Interestingly, despite that \mathcal{P}' has exponentially many variables, we manage to prove that its solutions correspond to a special class of information selling mechanisms and moreover, always admits an optimal solution which is a CM-probR. This yields the optimality of CM-probR mechanisms.

Finally, we extend our results to the setting where the seller does not know the prior distribution of the signals explicitly and can only access it through a black-box sampling oracle. This applies to the cases where distributions are only accessible through samples. Using a Monte-Carlo sampling approach, inspired by the idea of Dughmi and Xu (2016), we show that a mechanism with approximate optimality, IR, IC and obedience can be implemented in polynomial time with polynomially many samples from the prior distribution.

1.2 Related Work Most relevant to this paper is the work by Babaioff et al. (2012) who study the optimal mechanism design for selling information to an imperfectly informed decision maker (henceforth the *buyer*) without budget consideration. Our setting is essentially the same as that of Babaioff et al. (2012), but under *budget constraints*. This is motivated by the drawbacks of their optimal mechanisms which may require an extremely large amount of transfers even to extract a tiny amount of revenue. On the descriptive side, we show that budget constraints indeed affect the characterization of optimal mechanisms. For example, when the buyer has a private budget constraint, any mechanism with one-round transfer cannot be optimal even when signals are independent, which is opposed to the characterization in Babaioff et al. (2012) for the same setting but without budgets. On the prescriptive side, all the mechanisms in Babaioff et al. (2012) rely on solving exponentially large linear programs via the ellipsoid method whereas all our mechanisms can be computed by directly solving a polynomial-size linear program. This is enabled by our new characterizations of the optimal mechanism in much simpler and more interpretable forms.

The sale of information has attracted much attention in the economic literature. To our knowledge, Admati and Pfleiderer (1986, 1990) are among the first to explicitly consider the sale of information. They con-

sider a monopoly who sells information regarding a risky asset to a continuum of homogeneous traders in a financial market, which then affects the equilibrium price of the market. They show that the optimal selling method may involve revealing different and noisy signals to different traders. In contrast to the focus of Admati and Pfleiderer (1986, 1990) on selling information to multiple homogeneous traders with externalities, we study the sale of information to a single decision maker with heterogeneous types. Recently, Bergemann and Bonatti (2015) study the sale of cookies, a particular type of informational goods, motivated by the context of online advertising. Malenko and Malenko (2018) studies how to sell information to voters and characterize when the sale of information leads to a more informative voting outcome. Bergemann et al. (2018) considers the sale of information to a decision maker, which is similar to our setting but without budget constraints. Moreover, they restrict the design space to a particular format of mechanisms as a menu of statistical experiments. They establish various properties regarding the optimal mechanism and characterize the optimal mechanism in the cases of binary states and actions, or binary types. In particular, part of our proof for the characterization results of independent signals uses similar ideas by Bergemann et al. (2018) to reduce the number of signals required in the signaling scheme. Eso and Szentes (2007) studied the consultant's problem of optimally contracting with the client. They considered a different model in which a consultant incurs a cost to acquire a private observation, and once the observation is acquired, the consultant will fully reveal the information to the client.

Selling physical goods to buyers with budget constraints has been studied extensively in the literature of mechanism design (see, e.g., (Che and Gale, 2000; Chawla et al., 2011; Devanur and Weinberg, 2017) and references therein). Our work is conceptually related to, yet fundamentally different from, selling physical goods and meanwhile revealing information regarding the value of the item for sale. Emek et al. (2012); Bro Miltersen and Sheffet (2012); Badanidiyuru et al. (2018) study how an auctioneer can strategically reveal information to bidders in order to affect their valuation regarding the item with the ultimate goal of maximizing revenue in a fixed auction format (in particular, the second-price auction). Daskalakis et al. (2016) consider a seller of a physical item who also possesses private information regarding the item's value to buyers, and studies the optimal mechanism for jointly selling the item and the seller's private information. They show that this joint design reduces to optimal mechanism design for selling multiple items without information involved. Smolin (2019) considers a similar situation

where a seller sells an item, characterized by a vector of attributes, to a buyer. They also study the optimal joint design of the item pricing mechanism and information selling mechanism and provide a characterization of the optimal mechanism within a particular design space.

Finally, our work is also related to the rich literature of persuasion, a.k.a. *signaling* or *information design*. Our model is particularly relevant to persuasion of a privately informed receiver by Kolotilin et al. (2017) since the seller in our model reveals information by a signaling scheme to a buyer with a private signal. The techniques we use in Section 5 is inspired by algorithms for Bayesian persuasion designed by Dughmi and Xu (2016). Due to the space limit, we do not give a thorough review of the large literature in persuasion but refer the reader to the recent survey by Kamenica (2018) (from the economic perspective) and Dughmi (2017) (from the algorithmic perspective).

2 Preliminaries

Basic Setup We consider a monopoly information holder (call *her* the *seller*) selling information to a budget-constrained decision maker (call *him* the *buyer*), who needs to choose an action $a \in A$ and has a budget $b \in B$. The buyer's utility depends on both his action a and a random *state of the world*. The seller and the buyer each observe a private *signal* about the state of the world, denoted by $\omega \in \Omega$ and $\theta \in \Theta$, respectively. Sets A, B, Ω, Θ are all finite. Let $u(\omega, \theta, a)$ denote the buyer's utility of action a when the seller's signal and buyer's signal are realized to ω, θ , respectively.² For convenience, we sometimes also call ω and θ the *type* of the seller and buyer. Moreover, we assume that the seller also has a budget M to run the mechanism and M is *publicly known* (as we shall see, the optimal mechanism sometimes requires the seller to pay the buyer with some probability).

We assume that the signal pair ω, θ and the budget b are drawn from a publicly known prior distribution $\mu(\omega, \theta, b)$ and the utility function $u(\omega, \theta, a)$ is also public knowledge. The surplus for a buyer of type θ from fully observing the seller's signal ω , denoted as $\delta(\theta)$, is

$$\delta(\theta) = \mathbb{E}_\omega \left[\max_a u(\omega, \theta, a) \mid \theta \right] - \max_{a \in A} \mathbb{E}_\omega [u(\omega, \theta, a) \mid \theta] \geq 0.$$

The seller's goal is to sell her information regarding ω to the budget-constrained buyer to extract the largest possible revenue. As is typical in mechanism design, we assume that the seller is a trusted authority and will

²Alternatively and equivalently, one can think of the buyer's utility as a function u' of action a and the state of world denoted as by s . Then our definition of $u(\omega, \theta, a)$ simply corresponds to $\mathbb{E}_s [u'(s, a) \mid \omega, \theta]$.

not defect after posting a selling mechanism based on her knowledge of $\mu(\omega, \theta, b)$ and $u(\omega, \theta, a)$. However, the buyer will strategically respond to the mechanism in order to maximize his own utility and may defect if that is beneficial. We remark that this basic model has been studied in some previous works Babaioff et al. (2012); Bergemann et al. (2018). The key difference between our model and previous ones is that our seller and buyer are both *budget-constrained*.

The Design Space – Generic Interactive Protocols. We now describe the design space, i.e., the class of mechanisms within which our mechanism will be optimal. A key observation by Babaioff et al. (2012) is that a mechanism that interacts with the buyer for *multiple* rounds can possibly extract more revenue than any one-round mechanism. Therefore, we consider mechanisms that may interact with the buyer for multiple rounds via a *generic interactive protocol*, as defined by Babaioff et al. (2012). At a high level, the protocol/mechanism can be viewed as an extensive-form game represented as a tree with three types of nodes: *seller* nodes, *buyer* nodes and *transfer* node. Each move at a seller node can be seen as revealing some amount of information regarding ω to the buyer whereas each move at a buyer node reveals buyer information (e.g., the buyer's type) to the seller. All the payments — from the buyer to the seller or vice versa — happen at transfer nodes. More formally, the protocol is defined as follows.

DEFINITION 2.1. (BABAIOFF ET AL. (2012)) A generic interactive protocol is a finite-size protocol tree with the following three types of non-leaf nodes:

1. **Seller node:** at a seller node n , the seller randomly moves to one of the children of n according to her private signal ω . Let $C(n)$ be any child of n and $p_n(\omega, C(n))$ be the probability of moving to $C(n)$ when the seller's signal is ω . Note that $\sum_{C(n)} p_n(\omega, C(n)) = 1$ for all ω .
2. **Buyer node:** at a buyer node n , the buyer (with type θ and budget b) randomly moves to one of the children $C(n)$. A buyer's strategy at node n is a transition function $\phi_n((\theta, b), C(n))$, which indicates the probability of moving to the node $C(n)$. Naturally, it should satisfy $\sum_{C(n)} \phi_n((\theta, b), C(n)) = 1$.
3. **Transfer node:** any transfer node n only has one child and an associated payment to the seller of amount $t(n)$. Note that $t(n)$ can be negative, meaning the seller pays the buyer.

The whole protocol (including all the transition probabilities) is publicly known and both players know which node they are currently at during the execution.

However, each player's private type is only known to themselves. Nevertheless, the buyer at any buyer node can form a posterior belief on ω according to what has been observed so far along the tree path based on the transition function $p_n(\omega, C(n))$ at each seller node. During the game, the buyer can choose to leave the protocol if that is beneficial or the total spending exceeds his budget. Otherwise, the game terminates at some leaf node of the protocol. In either case, the buyer will then choose the best action based on his posterior belief of ω . Furthermore, the seller's budget constraint restricts the net transfer from seller to buyer to never exceeds M at any node of the protocol tree, and the buyer will never choose to visit a node of the protocol tree at which the net transfer from buyer to seller exceeds b .

We remark that all the optimality claims throughout the paper will be within the design space of all generic interactive protocols, though we tend to omit this emphasis for convenience. We first simplify the design space by invoking the revelation principle.

LEMMA 2.1. (REVELATION PRINCIPLE) *There exists an optimal generic interactive protocol in which the buyer truthfully reports type θ and budget b at the beginning and takes no other actions. Therefore, the protocol only has one buyer node, i.e., the root of the protocol tree. We call such a mechanism a revelation mechanism.*

The proof can be found in Appendix A. A revelation mechanism is *incentive compatible* (IC) if it is an optimal strategy for the buyer to report his true type θ and budget b at the beginning, and is *individually rational* (IR) if the expected utility of participating in the mechanism is non-negative for any buyer.

Information Revelation via Signaling Schemes. Though the generic interactive protocols reveal information sequentially in general and can be intricate, we prove that optimal mechanisms will only need to reveal information once (though transfers may happen at multiple rounds)³, and such information revelation can be described via a *signaling scheme*. Concretely, a signaling scheme is a *randomized* mapping from the seller's signal set Ω to a set of signals Σ , which can be fully described by the likelihood function $\{p(\omega, \sigma)\}_{\omega \in \Omega, \sigma \in \Sigma}$ where $p(\omega, \sigma)$ is the probability of sending signal σ given seller signal ω .

With slight abuse of notation, we use $\mu(\omega)$ to denote the buyer's belief regarding the seller's signal ω . After

receiving any signal σ from the signaling scheme, the buyer with prior belief μ will update his belief via a standard Bayesian update, and infer that the buyer signal is ω with probability

$$(2.1) \quad \Pr(\omega|\sigma) = \frac{\mu(\omega)p(\omega, \sigma)}{\sum_{\omega' \in \Omega} \mu(\omega')p(\omega', \sigma)}.$$

As an example, each buyer node n in the generic interactive protocol, together with its transitions described by $\{p_n(\omega, C(n))\}_{\omega, C(n)}$, is equivalent to a signaling scheme in which each $C(n)$ corresponds to a signal and $p_n(\omega, C(n))$ is the probability of sending signal $C(n)$ given the seller type ω . Under this interpretation, moving to node $C(n)$ is equivalent to that the seller sends a signal. The buyer then observes the signal and updates his belief regarding ω according to Equation (2.1). Though such information revelation and buyer's belief updates may happen multiple times in the generic interactive protocol, we will show that it suffices to reveal information only once in the optimal mechanism.

Player Utilities and the Revenue Maximization Problem. Let $Z_{\theta, b}$ denote the (random) node that the protocol ends at (so $Z_{\theta, b}$ could be a leaf, or a node where the buyer chooses to leave) when interacting with a buyer of type θ and budget b . Let $\mu(\omega|Z_{\theta, b})$ denote the buyer's posterior belief probability of ω at $Z_{\theta, b}$ and $t(Z_{\theta, b})$ be the summed transfer on the path from the root to $Z_{\theta, b}$. The expected utility of a buyer of type θ and budget b is then equal to

$$U(\theta, b) = \mathbb{E}_{Z_{\theta, b}} \left[\max_a \sum_{\omega \in \Omega} [u(\omega, \theta, a) \cdot \mu(\omega|Z_{\theta, b}) - t(Z_{\theta, b})] \right],$$

where the randomness of node $Z_{\theta, b}$ is due to the internal randomness of the mechanism. Note that the buyer can always leave at the beginning, with expected utility $\max_a \sum_{\omega \in \Omega} u(\omega, \theta, a) \cdot \mu(\omega|\theta, b)$.

On the other hand, the seller's revenue of a generic interactive protocol M is

$$\text{Rev}(M) = \sum_{\theta, b} \mu(\theta, b) \cdot \mathbb{E}_{Z_{\theta, b}} [t(Z_{\theta, b})],$$

where $\mu(\theta, b) = \sum_{\omega} \mu(\omega, \theta, b)$ is the prior probability of a buyer of type θ and budget b showing up. The seller's goal is to choose the optimal protocol M that maximizes the above expected revenue, assuming a rational buyer who will optimize his own utility in this mechanism.

3 Independent Signals

An interesting special case of the problem is that the buyer's signal and budget (θ, b) is independent from the seller signal ω , i.e., $\mu(\omega, \theta, b) = \mu(\omega) \cdot \mu(\theta, b)$ where $\mu(\cdot)$

³ If one imposes a no-positive-transfers constraint, that $t(n)$ cannot be negative at any transfer node, then Babaioff et al. (2012) present an example where the seller may need to reveal information more than once in the optimal mechanism.

denotes the prior distribution of the corresponding random variables. We study this special case separately due to two reasons. First, this setting has another natural interpretation — i.e., the θ can also be interpreted as the buyer’s private type that captures the buyer’s *information-irrelevant* characteristics (thus independent of the seller’s signal) such as, e.g., properties of his decision problem u . Second, the optimal mechanism of this case has a more concise representation. In particular, we prove that in this case, there exists an optimal budget-feasible mechanism with the following simple procedure: (1) it asks for the buyer’s type θ and budget b ; (2) charges the buyer a fixed amount (via direct payment or first ask for a deposit of b and then return a part of it); (3) recommend an action for the buyer to take. This mechanism is similar in spirit to a “consulting procedure” during which the buyer (a client) tells his type to the seller (the consultant). The seller then charges the buyer a fixed amount of money and recommends an action to the buyer to take based on his reported type. We thus call such mechanisms *consulting mechanisms*. It turns out that, with independent signals, the optimal mechanism needs to use slightly different payment methods for the public and private budget cases. We thus consider them separately in subsection 3.1 and 3.2.

3.1 Public Budget In this subsection, we prove that when the signals are independent and the buyer has a fixed public budget b , there always exists an optimal consulting mechanism with one-round direct payment from the buyer to the seller, formally defined as follows.

DEFINITION 3.1. *A consulting mechanism with direct payment (CM-dirP) for a buyer with publicly known budget b proceeds as follows:*

1. *The seller commits to a payment amount $t_\theta (\leq b)$ and an action recommendation policy described as a mapping $p_\theta : \Omega \rightarrow \Delta A$ for each buyer type θ , where $p_\theta(\omega) \in \Delta A$ with $p_\theta(\omega, a)$ denoting the probability of recommending action a given seller signal ω . p_θ is required to be obedient — i.e., conditioned on any recommended action, it must indeed be an optimal action for the buyer given his information.*
2. *The seller asks the buyer to report his type $\hat{\theta}$.*
3. *The seller charges the buyer an amount $t_{\hat{\theta}}$.*
4. *Based on her signal ω , the seller samples an action $a \sim p_{\hat{\theta}}(\omega)$ and recommends a to the buyer.*

We make a few remarks about CM-dirP. First, CM-dirP is indeed a generic interactive protocol, described as follows: (1) its root is the only buyer node at which he

reports his type $\hat{\theta}$; (2) each child of the root corresponds to a buyer type $\hat{\theta}$ and is a transfer node with transfer amount $t_{\hat{\theta}}$ (to the seller); (3) following each transfer node is a seller node whose children are leaves, each indexed by a buyer action $a \in A$, and the transition probability to leaf node $a \in A$ is $p_{\hat{\theta}}(\omega, a)$. Second, as described in Section 2, the transition from seller node to leaves is effectively a signaling scheme, in which the set of signals coincides with the set of actions. Upon receiving an action recommendation a , the buyer infers a posterior probability of ω by Bayes updates, as follows: $\Pr(\omega|a) = \frac{\mu(\omega)p_{\hat{\theta}}(\omega, a)}{\sum_{\omega' \in \Omega} \mu(\omega')p_{\hat{\theta}}(\omega', a)}$. Finally, an intrinsic constraint of CM-dirP is that the seller’s action recommendation must satisfy the *obedience constraint*. That is, the recommended action must indeed be an optimal action for the buyer’s reported type or, more formally,

$$a = \operatorname{argmax}_{a' \in A} \sum_{\omega} u(\omega, \hat{\theta}, a') \cdot \frac{\mu(\omega)p_{\hat{\theta}}(\omega, a)}{\sum_{\omega' \in \Omega} \mu(\omega')p_{\hat{\theta}}(\omega', a)}.$$

Our main result of this subsection is to prove the optimality of CM-dirP mechanisms, described as follows. The proof of Theorem 3.1 is deferred to Appendix B.1.

THEOREM 3.1. *When the buyer budget b is public and the buyer signal θ is independent from ω , there always exists an optimal mechanism that is an IC consulting mechanism with direct payment (CM-dirP).*

3.2 Private Budget When the budget is privy to the buyer, our mechanism will also need to ask the buyer to report his budget. It turns out that in this case, any mechanism with single-round transfer cannot achieve optimality, which is illustrated as follows.

EXAMPLE 3.1. *We consider a similar setup as the “treasure box” example used by Babaioff et al. (2012). In particular, imagine a box with a locker on it. There are two keys labeled with 0 and 1, exactly one of which can open the box. The buyer can choose one key and try it. If he opens the box, he gets the object inside. Let the type of the buyer $\theta \in \{0, 1\}$ encode his value z_θ for the object, where $z_0 = 120$ and $z_1 = 80$. The seller knows exactly what the correct key is, and let it be $\omega \in \{0, 1\}$. How should the seller sell her information to the buyer?*

In this problem, we have $\Omega = \Theta = A = \{0, 1\}$, $u(\omega, \theta, a) = \mathbf{1}\{\omega = a\} \cdot z_\theta$. A type-0 buyer has a budget of 50 whereas a type-1 buyer has a budget of 100. The signals θ and ω are independent and uniformly at random. Concretely, the prior distribution can be expressed as follows: $\mu(\omega, \theta, b) = 1/4$ when $(\omega, \theta, b) = (0, 0, 50), (1, 0, 50), (0, 1, 100), (1, 1, 100)$ and $\mu(\omega, \theta, b) =$

0 otherwise. Simple calculation shows that the buyer's surplus from fully knowing the seller's signal ω will be $\delta(0) = 60$ and $\delta(1) = 40$.

It can be shown that some **CM-dirP** mechanism achieves optimality among all mechanisms with a single-round of transfer. Moreover, the optimal **CM-dirP** mechanism is to charge a fixed price 40 and then reveal the exact value of ω , which achieves expected revenue 40.⁴ However, strictly more revenue can be extracted via the following mechanism with two rounds of transfers. Consider the mechanism that offers the buyer the following two options:

1. Pay 50 dollars to know the exact value of ω .
2. Pay 100 dollars first and then get a refund of 61 dollars together with the exact value of ω .

The type-0 buyer only has 50 dollars, so he can only choose Option 1. The type-1 buyer will choose Option 2. The expected revenue of this mechanism is thus $(50 + 39)/2 = 44.5$, which is higher than 40.

We show that consulting mechanisms can still achieve optimality but with a slightly different payment method: first ask the buyer to fully deposit his budget, and then return a part of the deposit back.

DEFINITION 3.2. A consulting mechanism with deposit and return (**CM-depR**) proceeds as follows: (1) The seller commits to a payment amount $t_{\theta,b} (\leq b)$ and an action recommend policy $p_{\theta,b} : \Omega \rightarrow \Delta A$ for each (θ, b) , under obedience constraints (similar to Step 1 of the **CM-dirP** mechanism); (2) Asks the buyer to report his type $\hat{\theta}$ and deposit his budget \hat{b} ; (3) Returns to the buyer an amount of $\hat{b} - t_{\hat{\theta},\hat{b}}$; (4) Samples an action $a \sim p_{\hat{\theta},\hat{b}}(\omega)$ and recommends a to the buyer.

As can be seen, a **CM-depR** mechanism is almost the same as **CM-dirP** except that it charges the buyer $t_{\hat{\theta},\hat{b}}$ by first asking him to deposit the budget \hat{b} and then immediately returning $\hat{b} - t_{\hat{\theta},\hat{b}}$ whereas **CM-dirP** just asks for a direct payment $t_{\hat{\theta}}$. This step is useful in the private budget case because it serves as a "verification" of the buyer's budget which helps the seller to learn more information about the buyer's type and thus increases her power of price discrimination. This intuition is also illustrated in Example 3.1.

THEOREM 3.2. When the buyer has a private budget and (θ, b) is independent from ω , there always exists an optimal mechanism that is an **IC** consulting mechanism with deposit and return (**CM-depR**).

The proof can be found in Appendix B.2.

⁴in Section 3.3, we will show how to compute the optimal **CM-dirP** mechanism via a polynomial-size linear program.

3.3 Computing Optimal Consulting Mechanisms

We now show how to efficiently compute the optimal consulting mechanisms for both public and private buyer budget. It turns out that in both cases, the optimal consulting mechanisms can be easily computed by simple linear programs with polynomial size. This much simplifies the optimal mechanism proposed by Babaioff et al. (2012) (for the setting without budget constraints), which requires solving exponentially-large linear programs by going through the dual program and the ellipsoid method. Our mechanism is simpler, more interpretable and can be more efficiently computed from a practical perspective.

Here, we only give the LP formulation for computing the optimal **CM-depR** mechanism since the formulation for the optimal **CM-dirP** is essentially the same. By definition, a **CM-depR** mechanism can be fully described by variables $p_{\theta,b}(\omega, a)$, which is the probability of recommending action a when the seller signal is realized to ω and the buyer reports (θ, b) , and the corresponding net transfer variable $t_{\theta,b}$. If a (θ, b) -buyer misreports type θ' and budget b' , he will receive a recommendation generated according to the random mapping $p_{\theta',b'}(\omega, a)$ from the seller. However, this recommended action does not have to be optimal for this buyer due to his misreport. Nevertheless, the buyer's optimal expected utility can still be computed as the sum of the optimal expected utilities from all action recommendations:

$$\sum_a \max_{a'} \sum_{\omega} \mu(\omega) p_{\theta',b'}(\omega, a) u(\omega, \theta, a') - t_{\theta',b'}.$$

As a result, the optimal consulting mechanism can be computed via a linear program (LP) as in Figure 1 with variables $p_{\theta,b}(\omega, a)$ and $t_{\theta,b}$.

The first constraint guarantees the *obedience* constraint in the consulting recommendation, i.e., each recommended action must indeed be an optimal action for the buyer's reported type. The second constraint guarantees incentive compatibility (**IC**) and the third constraint guarantees individual rationality (**IR**).

We remark that the **IC** constraints (above) are not linear (yet). However, one can easily transform them into linear constraints by introducing new variables $z_{\theta,\theta',b',a}$ to represent the max function, as follows

$$z_{\theta,\theta',b',a} = \max_{a'} \sum_{\omega} \mu(\omega) \cdot p_{\theta',b'}(\omega, a) \cdot u(\omega, \theta, a').$$

Then for any $\theta, \theta', b \geq b'$, the **IC** constraint can be reformulated as follows:

$$\begin{aligned} & \sum_{\omega, a} \mu(\omega) \cdot p_{\theta,b}(\omega, a) \cdot u(\omega, \theta, a) - t_{\theta,b} \\ & \geq \sum_a z_{\theta,\theta',b',a} - t_{\theta',b'} \end{aligned}$$

$$\begin{aligned}
\max \quad & \sum_{\theta,b} \mu(\theta,b) \cdot t_{\theta,b} \\
\text{s.t.} \quad & \sum_{\omega} \mu(\omega) p_{\theta,b}(\omega,a) u(\omega,\theta,a) \geq \sum_{\omega} \mu(\omega) p_{\theta,b}(\omega,a) u(\omega,\theta,a'), \quad \text{for } a,a',\theta,b \quad (\text{OB}) \\
& \sum_{\omega,a} \mu(\omega) p_{\theta,b}(\omega,a) u(\omega,\theta,a) - t_{\theta,b} \\
& \geq \sum_a \max_{a'} \sum_{\omega} \mu(\omega) p_{\theta',b'}(\omega,a) u(\omega,\theta',a') - t_{\theta',b'}, \quad \text{for } \theta,\theta',b,b' \quad (\text{IC}) \\
& \sum_{\omega,a} \mu(\omega) p_{\theta,b}(\omega,a) u(\omega,\theta,a) - t_{\theta,b} \geq \sum_{\omega} \mu(\omega) u(\omega,\theta,a'), \quad \text{for } (\theta,b),a' \quad (\text{IR}) \\
& t_{\theta,b} \leq b, \quad \text{for } (\theta,b) \quad (\text{budget}) \\
& \sum_a p_{\theta,b}(\omega,a) = 1, \quad \text{for } (\theta,b),\omega \\
& 0 \leq p_{\theta,b}(\omega,a) \leq 1
\end{aligned}$$

Figure 1: LP formulation that computes the optimal consulting mechanisms.

with additional constraints that define $z_{\theta,\theta',b',a}$:

$$z_{\theta,\theta',b',a} \geq \sum_{\omega} \mu(\omega) \cdot p_{\theta',b'}(\omega,a) \cdot u(\omega,\theta,a'), \quad \forall a' \in A.$$

These overall establish the following theorem.

THEOREM 3.3. *The optimal CM-dirP mechanism for public buyer budget and the optimal CM-depR mechanism for private buyer budget can each be computed by a single linear program of poly(| Θ |, | A |, | Ω |) size.*

4 Correlated Signals

In this section, we turn to the general setting with correlated signals. In this case, buyers of different signal θ 's will have different prior beliefs on ω due to the correlation between ω and θ . On one hand, this increases the seller's power to do price discrimination. On the other hand, it also complicates the design of the optimal mechanism. In particular, with independent signals, the payments of the optimal mechanisms only depend on buyer types but do not need to be contingent on the information revealed (recall that CM-dirP and CM-depR ask for payments *before* any information is revealed). However, Babaioff et al. (2012) show that this ceases to hold for correlated signals, even when there is no budget constraint. That is, when the seller's signal and the buyer's signal are correlated, the payment of the optimal mechanism has to depend on the information that is revealed to the seller.

In this section, we prove that the consulting mechanism coupled with a particular way of payments (which does depend on the information revealed) — together termed the *Consulting Mechanisms with Probabilistic Return (CM-probR)* — is optimal for both the public and private buyer budget. With this characterization result, we then show that the optimal mechanism can be computed efficiently, again by formulating a polynomial-size linear program to directly solve for the optimal CM-probR mechanism.

4.1 Optimality of CM-probR for Correlated Signals

We start by formally defining the Consulting

Mechanisms with Probabilistic Return (CM-probR), as follows

DEFINITION 4.1. *Let $[I] = \{+, -\}$ contain the indicators about whether the buyer will receive a return (“+”) or not (“−”). A consulting mechanism with probabilistic return (CM-probR) for a budget-constrained buyer proceeds as follows:*

1. *The seller commits to an action recommendation policy described as a mapping $p_{\theta,b} : \Omega \rightarrow \Delta(A \times I)$ for each (θ,b) , where $p_{\theta,b}(\omega) \in \Delta(A \times I)$ with $p_{\theta,b}(\omega, [a, i])$ denoting the probability of recommending action a with indicator $i \in \{+, -\}$ given seller signal ω . $p_{\theta,b}$ is required to be obedient — i.e., conditioned on any recommended action, it must indeed be the buyer's optimal action given his information.*
2. *The seller asks the buyer to report his type $\hat{\theta}$ and deposit his budget \hat{b} .*
3. *Based on her signal ω , the seller samples $[a, i] \sim p_{\hat{\theta}, \hat{b}}(\omega)$, and recommends a to the buyer.*
4. *If the sampled $i = -$, return $\hat{b} + M$ to the buyer; return 0 otherwise (recall M is the seller's budget).*

The key difference between CM-probR and CM-dirP is in their ways of payments. In CM-probR, the seller has some probability of receiving a return of $\hat{b} + M$ and this probability depends on the $[a, i]$ pair. One might wonder why we treat the two cases $+$ and $-$ (w.r.t. an action a) separately while not simply use the expected payment instead. We note that this would *not* work because the posterior probability of each payment is different for different buyer types due to their different prior beliefs on ω . The combined payment will *not* simply translate to the same expectation for different buyer types due to nonlinearity.

Our main result of this subsection is to prove the optimality of CM-probR for correlated signals.

THEOREM 4.1. *When θ, ω are correlated, there always exists an IC Consulting Mechanism with Probabilistic*

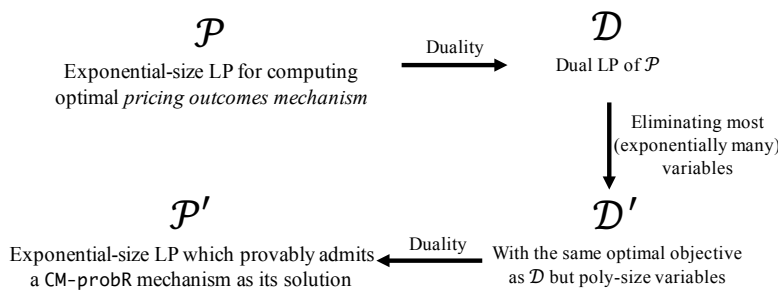


Figure 2: Illustration of the Main Idea for Proving Optimality of CM-probR.

Return (CM-probR) that maximizes the seller’s revenue in both public and private buyer budget settings.

The proof of Theorem 4.1 relies on a novel use of the duality theory. We defer the full proof to Appendix C and only provide a sketch here. Our starting point is a result by Babaioff et al. (2012) who proved that for the setting *without* budget constraints, there always exists an optimal mechanism of the following format: (1) ask the buyer for a (possibly extremely large) deposit; (2) reveal information through a *signaling scheme*; (3) return to the buyer some amount of the deposit which will depend on the signal sent. They term it a *pricing outcomes mechanism*. As a first step, we generalize this result and prove that it still holds even with budget constraints, but with Step (1) now being to ask the buyer to deposit his reported budget.

Observe that CM-probR mechanisms are a strict *subclass* of pricing outcomes mechanisms — in particular, CM-probR’s use at most $2n$ signals and only two possible payments, i.e., returning $b + M$ or 0. We want to prove that this much smaller class of mechanisms is still able to achieve the optimal revenue. Our idea goes as follows (also see illustrations in Figure 2). We first formulate the problem of computing the optimal pricing outcomes mechanism, which is a linear program \mathcal{P} with *exponentially* many variables and *exponentially* many constraints. Naturally, any naive approach — including the celebrated ellipsoid method — cannot be directly used to solve such an LP or its dual. We nevertheless work with the dual of \mathcal{P} , denoted as LP \mathcal{D} . Obviously, \mathcal{D} also has exponentially many variables and exponentially many constraints, which is still difficult to solve.

The crux of our approach is to prove that the dual LP \mathcal{D} can be reduced to another linear program \mathcal{D}' , which will achieve the same optimal objective value as of LP \mathcal{D} , still has exponentially many constraints but will have only polynomially many variables. At this point, one natural idea is to apply the ellipsoid method to solve \mathcal{D}' to obtain the optimal solution, which turns out indeed is doable in polynomial time by designing a separation oracle for \mathcal{D}' . However, this idea suffers

from several drawbacks. First, the ellipsoid method is widely known to be inefficient in practice. Second, the dual of \mathcal{D}' , denoted as \mathcal{P}' , is not our original LP \mathcal{P} any more (e.g., \mathcal{P}' now has exponentially many variables but polynomially many constraints) and it is not clear how it can help to recover the optimal solution to \mathcal{P} instead. We believe that by opening the black box of Ellipsoid method and after addressing the tedious arithmetic issues (since the separation oracle for \mathcal{D}' needs to solve a convex program, which can only be solved to be within ϵ precision in $\text{poly}(\log(1/\epsilon))$ time), one may indeed be able to recover the optimal solution to \mathcal{P} . However, this would take significantly more effort. Last but not least, the optimal solution obtained through such an optimization procedure is very unlikely to be a simple and interpretable mechanism.

As a result, we choose to work with the dual of LP \mathcal{D}' , denoted as \mathcal{P}' , which recall that now has polynomially many constraints but exponentially many variables. Surprisingly, we show that this transformation yields a simple characterization of the optimal mechanism as well as a much more efficient algorithm for computing it. This is enabled by two intriguing properties of \mathcal{P}' which we prove. First, any feasible solution to LP \mathcal{P}' corresponds to a particular format of pricing outcomes mechanisms, which only either asks the buyer for a net payment of b or pays the buyer the amount of M , and has no other format of transfers. Second, we prove that there always exists an optimal solution to LP \mathcal{P}' , which is a CM-probR mechanism, as described in Definition 4.1. As a result, we can conclude that CM-probR can achieve the optimal revenue since the optimal objective of \mathcal{P}' must equal that of \mathcal{P} due to strong duality.

4.2 Computing the Optimal CM-probR We now show how to efficiently compute the optimal CM-probR mechanism via a polynomial-size linear program. For notational convenience, instead of $p_{\theta,b}(\omega, [a, +])$ and $p_{\theta,b}(\omega, [a, -])$ as in Definition 4.1, we use variable $p_{\theta,b}^+(\omega, a)$ and $p_{\theta,b}^-(\omega, a)$ to denote the transition probabilities. Then the utility of the buyer with type θ and

$$\begin{aligned}
(4.2) \quad & \max \quad \sum_{\omega, \theta, b} \mu(\omega, \theta, b) \sum_a \left(b \cdot p_{\theta, b}^+(\omega, a) - M \cdot p_{\theta, b}^-(\omega, a) \right) \\
& \text{s.t.} \quad U_{\theta, b}(\theta, b) \geq U_{\theta, b}(\theta', b'), & \text{for } \theta, \theta', b \geq b' & \quad \text{(IC)} \\
& \quad U_{\theta, b}(\theta, b) \geq \sum_{\omega} \mu(\omega | \theta, b) u(\omega, \theta, a'), & \text{for } (\theta, b), a' \in A & \quad \text{(IR)} \\
& \quad U_{\theta, b}(\theta, b) \leq \sum_{a, \omega, \circ \in \{+, -\}} \mu(\omega | \theta, b) p_{\theta, b}^{\circ}(\omega, a) (u(\omega, \theta, a) - t^{\circ}), & \text{for } \theta, b & \quad \text{(OB)} \\
& \quad U_{\theta, b}(\theta', b') = \sum_a (z_{\theta, b, \theta', b', a}^+ + z_{\theta, b, \theta', b', a}^-), & \text{for } \theta, b, \theta', b' & \\
& \quad z_{\theta, b, \theta', b', a}^{\circ} \geq \sum_{\omega} \mu(\omega | \theta, b) p_{\theta', b'}^{\circ}(\omega, a) (u(\omega, \theta, a') - t^{\circ}), & \text{for } \theta, b, \theta', b', a, a', \circ \in \{+, -\} & \\
& \quad \sum_{a \in A} [p_{\theta, b}^+(\omega, a) + p_{\theta, b}^-(\omega, a)] = 1, & \text{for } (\theta, b), \omega & \\
& \quad 0 \leq p_{\theta, b}^+(\omega, a), p_{\theta, b}^-(\omega, a) \leq 1, & \text{for } \theta, b, \omega, a &
\end{aligned}$$

Figure 3: LP formulation for the optimal CM-probR mechanism.

budget b when he reports (θ', b') is equal to

$$\begin{aligned}
& U_{\theta, b}(\theta', b') \\
= & \sum_a \left(\max_{a'} \sum_{\omega} \mu(\omega | \theta, b) p_{\theta', b'}^+(\omega, a) (u(\omega, \theta, a') - b) \right. \\
& \quad \left. + \max_{a'} \sum_{\omega} \mu(\omega | \theta, b) p_{\theta', b'}^-(\omega, a) (u(\omega, \theta, a') + M) \right)
\end{aligned}$$

Introduce new variables

$$\begin{aligned}
z_{\theta, b, \theta', b', a}^+ &= \max_{a'} \sum_{\omega} \mu(\omega | \theta, b) p_{\theta', b'}^+(\omega, a) (u(\omega, \theta, a') - t^+), \\
z_{\theta, b, \theta', b', a}^- &= \max_{a'} \sum_{\omega} \mu(\omega | \theta, b) p_{\theta', b'}^-(\omega, a) (u(\omega, \theta, a') - t^-),
\end{aligned}$$

where $t^+ = b$ and $t^- = -M$. Then the LP for computing the optimal CM-probR mechanism can be formulated as in Figure 3.

These overall establish the following theorem.

THEOREM 4.2. *The optimal CM-probR mechanism can be computed by a single linear program of $\text{poly}(|\Theta|, |A|, |\Omega|)$ size for both public and private buyer budget setting.*

5 Extension to Black-box Prior Distributions

In this section, we extend our results to the setting where the seller does *not* know the prior distribution $\mu(\omega, \theta, b)$ and can only access the distribution by drawing i.i.d. samples from $\mu(\omega, \theta, b)$. We call this model the *black-box prior distribution*. This applies to the cases where distributions are only accessible through samples. We allow the signals to be correlated in this section, which includes independent signals as a special case. We provide a Monte-Carlo Sampling mechanism which draws polynomially many samples at the beginning and then computes an approximately optimal mechanism using the empirical distribution. This mechanism is nearly optimal and approximately achieves incentive compatibility and individual rationality when the buyer knows that the samples are drawn from the true distribution. This mechanism is similar to the one in [Dughmi and Xu \(2016\)](#) which computes nearly optimal persuasion under unknown prior distributions. Our problem is

Mechanism 1: CM-probR Mechanism for Black-box Prior Distribution

- (1) Observe private signal ω_1 . Ask the buyer to give his type θ_1 and deposit his budget b_1 .
 - (2) Draw $n - 1$ samples $T = \{(\theta_2, \omega_2, b_2), \dots, (\theta_n, \omega_n, b_n)\}$ from $\mu(\theta, \omega, b)$ to get $S = \{(\theta_1, \omega_1, b_1), \dots, (\theta_n, \omega_n, b_n)\} \sim \mu^n$.
 - (3) Solve an LP to get a CM-probR mechanism \mathcal{M}_S . The LP will be defined later.
 - (4) Use \mathcal{M}_S to sell the information to the buyer. Specifically, according to ω_1, θ_1, b_1 , the mechanism will randomly recommend an action $a_S(\theta_1, b_1)$, and charge a payment $t_S(\theta_1, b_1) \in \{b_1, -M\}$.
-

more intricate because in persuasion the receiver's type (corresponding to the seller in our case) is known to the sender, but in our problem the buyer's type is randomly drawn from a distribution and needs to be sampled. For completeness, we fully state our algorithm and results here and defer the full proof to Appendix D.

The seller will post the mechanism before observing ω_1 and commit to follow the mechanism. The buyer does not know $\mu(\omega, \theta, b)$ but knows how the mechanism works. Our mechanism will satisfy the following properties.

DEFINITION 5.1. (ε -IC) *Mechanism 1 is ε -IC if for any possible $\mu(\omega, \theta, b)$, it is ε -optimal for the buyer to truthfully report and follow the mechanism's recommendation in expectation. More specifically, for any $\mu(\omega, \theta, b)$, any $(\theta, b) \neq (\theta', b')$ with $b' \leq b$,*

$$\begin{aligned}
& \mathbb{E}_{\omega_1 \sim \mu(\omega | \theta, b)} \mathbb{E}_{T \sim \mu^{n-1}} \mathbb{E}_{\mathcal{M}_S} [u(\omega_1, \theta, a_S(\theta, b)) - t_S(\theta, b)] \\
& \geq \mathbb{E}_{\omega_1 \sim \mu(\omega | \theta, b)} \mathbb{E}_{T \sim \mu^{n-1}} \mathbb{E}_{\mathcal{M}_S} [u(\omega_1, \theta, a'(a_S(\theta', b')))) \\
& \quad - t_S(\theta', b')] - \varepsilon.
\end{aligned}$$

for any deviation from the recommendation $a' : A \rightarrow A$.

DEFINITION 5.2. (ε -IR) *Mechanism 1 is ε -IR if for any possible $\mu(\omega, \theta, b)$, the buyer's expected surplus is*

$$\begin{aligned}
(5.3) \quad & \max \quad \sum_{\omega, \theta, b} \widehat{\mu}(\omega, \theta, b) \sum_a \left(b \cdot p_{\theta, b}^+(\omega, a) - M \cdot p_{\theta, b}^-(\omega, a) \right) \\
& \text{s.t.} \quad U_{\theta, b}(\theta, b) \geq U_{\theta, b}(\theta', b') - \varepsilon, & \text{for } \theta, \theta', b \geq b' & \quad (\varepsilon\text{-IC}) \\
& \quad U_{\theta, b}(\theta, b) \geq \sum_{\omega} \widehat{\mu}(\omega | \theta, b) u(\omega, \theta, a') - \varepsilon, & \text{for } (\theta, b), a' \in A & \quad (\varepsilon\text{-IR}) \\
& \quad U_{\theta, b}(\theta, b) - \varepsilon \leq \sum_{a, \omega, \circ \in \{+, -\}} \widehat{\mu}(\omega | \theta, b) p_{\theta, b}^{\circ}(\omega, a) (u(\omega, \theta, a) - t^{\circ}), & \text{for } \theta, b & \quad (\varepsilon\text{-OB}) \\
& \quad U_{\theta, b}(\theta', b') = \sum_a (z_{\theta, b, \theta', b', a}^+ + z_{\theta, b, \theta', b', a}^-), & \text{for } \theta, b, \theta', b' & \\
& \quad z_{\theta, b, \theta', b', a}^{\circ} \geq \sum_{\omega} \widehat{\mu}(\omega | \theta, b) p_{\theta, b'}^{\circ}(\omega, a) (u(\omega, \theta, a') - t^{\circ}), & \text{for } \theta, b, \theta', b', a, a', \circ & \\
& \quad \sum_{a \in A} [p_{\theta, b}^+(\omega, a) + p_{\theta, b}^-(\omega, a)] = 1, & \text{for } (\theta, b), \omega & \\
& \quad 0 \leq p_{\theta, b}^+(\omega, a), p_{\theta, b}^-(\omega, a) \leq 1, & \text{for } \theta, b, \omega, a &
\end{aligned}$$

Figure 4: LP formulation for black-box prior distributions.

no less than $-\varepsilon$ when he truthfully reports (θ, b) and follows the recommendation. More specifically, for any $\mu(\omega, \theta, b)$ and any (θ, b) ,

$$\begin{aligned}
& \mathbb{E}_{\omega_1 \sim \mu(\omega | \theta, b)} \mathbb{E}_{T \sim \mu^{n-1}} \mathbb{E}_{\mathcal{M}_S} [u(\omega_1, \theta, a_S(\theta, b)) - t_S(\theta, b)] \\
& \geq \mathbb{E}_{\omega_1 \sim \mu(\omega | \theta, b)} [u(\omega_1, \theta, a')] - \varepsilon
\end{aligned}$$

for all $a' \in A$.

DEFINITION 5.3. (ε -OBEDIENCE) *Mechanism 1* is ε -obedient if for any possible $\mu(\omega, \theta, b)$, when the buyer truthfully reports, in expectation it is ε -optimal for the buyer to take action a when being recommended action a . More specifically, for any $\mu(\omega, \theta, b)$, any (θ, b) and mapping $a' : A \rightarrow A$,

$$\begin{aligned}
& \mathbb{E}_{\omega_1 \sim \mu(\omega | \theta, b)} \mathbb{E}_{T \sim \mu^{n-1}} \mathbb{E}_{\mathcal{M}_S} [u(\omega_1, \theta, a_S(\theta, b))] \\
& \geq \mathbb{E}_{\omega_1 \sim \mu(\omega | \theta, b)} \mathbb{E}_{T \sim \mu^{n-1}} \mathbb{E}_{\mathcal{M}_S} [u(\omega_1, \theta, a'(a_S(\theta, b)))] - \varepsilon.
\end{aligned}$$

We now define the LP that solves \mathcal{M}_S . The LP is basically the same as the one we use when μ is known but with estimated $\mu(\omega, \theta, b)$ and $\mu(\omega | \theta, b)$. For $\mu(\omega, \theta, b)$ in the objective function, we estimate it with the empirical distribution over set S , denoted by $\widehat{\mu}(\omega, \theta, b)$. For $\mu(\omega | \theta, b)$ in the constraints, we estimate it with the empirical distribution over the set $S_{\theta, b} = \{\omega_1\} \cup \{\omega_i : i \geq 2 \text{ and } (\theta_i, b_i) = (\theta, b)\}$, denoted by $\widehat{\mu}(\omega | \theta, b)$. Note that the definition of $\widehat{\mu}(\omega | \theta, b)$ does **not** use the buyer's report (θ_1, b_1) . The LP with variables p^+, p^- is shown in Figure 4.

LEMMA 5.1. *When \mathcal{M}_S is solved by the LP (5.3), Mechanism 1 is ε -IR, ε -IC and ε -obedient.*

The lemma is proved in Appendix D.1. We then show that the mechanism extracts nearly optimal revenue with a sufficient number of samples. WLOG we assume $u \in [0, 1]$, $M \in [0, 1]$ and $b \in [0, 1]$ for all $b \in B$. We assume that $\mu_{\min} = \min_{\theta, b} \mu(\theta, b)$ is an instance-dependent constant.

THEOREM 5.1. *When we use Mechanism 1 with LP (5.3) with*

$$\begin{aligned}
n & \geq \Theta \left(\ln(G/\delta) \cdot \max \left\{ \frac{|A|^2}{\varepsilon^2 \cdot \mu_{\min}}, \frac{1}{\mu_{\min}^2} \right\} \right) \\
& = \Theta \left(\frac{|A|^2 \cdot \ln(G/\delta)}{\varepsilon^2} \right),
\end{aligned}$$

where $G = \max\{|\Theta|, |B|, |A|\}$ and $\mu_{\min} = \min_{\theta, b} \mu(\theta, b)$, the mechanism will be ε -IR, ε -IC and ε -obedient, and extract revenue no less than $\text{Rev}_{\mu}(p^*) - \delta$. $\text{Rev}_{\mu}(p^*)$ is the expected revenue of the optimal solution of (4.2) p^* .

The theorem is proved in Appendix D.2.

References

- Anat R Admati and Paul Pfleiderer. 1986. A monopolistic market for information. *Journal of Economic Theory* 39, 2 (1986), 400–438.
- Anat R Admati and Paul Pfleiderer. 1990. Direct and indirect sale of information. *Econometrica: Journal of the Econometric Society* (1990), 901–928.
- Moshe Babaioff, Robert Kleinberg, and Renato Paes Leme. 2012. Optimal Mechanisms for Selling Information. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12)*. ACM, New York, NY, USA, 92–109. <https://doi.org/10.1145/2229012.2229024>
- Ashwinkumar Badanidiyuru, Kshipra Bhawalkar, and Haifeng Xu. 2018. Targeting and signaling in ad auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2545–2563.
- Dirk Bergemann and Alessandro Bonatti. 2015. Selling cookies. *American Economic Journal: Microeconomics* 7, 3 (2015), 259–94.
- Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. 2018. The design and price of information. *American economic review* 108, 1 (2018), 1–48.

- Peter Bro Miltersen and Or Sheffet. 2012. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 234–247.
- Shuchi Chawla, David L Malec, and Azarakhsh Malekian. 2011. Bayesian mechanism design for budget-constrained agents. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 253–262.
- Yeon-Koo Che and Ian Gale. 2000. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic theory* 92, 2 (2000), 198–233.
- Constantinos Daskalakis, Christos Papadimitriou, and Christos Tzamos. 2016. Does Information Revelation Improve Revenue?. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 233–250.
- Nikhil R Devanur and S Matthew Weinberg. 2017. The optimal mechanism for selling to a budget constrained buyer: The general case. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 39–40.
- S. Dughmi. 2017. Algorithmic Information Structure Design: A Survey. *ACM SIGecom Exchanges* 15 (2017), 2–24.
- Shaddin Dughmi and Haifeng Xu. 2016. Algorithmic bayesian persuasion. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM, 412–425.
- Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. 2012. Signaling Schemes for Revenue Maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12)*. ACM, New York, NY, USA, 514–531. <https://doi.org/10.1145/2229012.2229051>
- Peter Eso and Balazs Szentes. 2007. The price of advice. *The Rand Journal of Economics* 38, 4 (2007), 863–880.
- Emir Kamenica. 2018. Bayesian persuasion and information design. *Annual Review of Economics* 11 (2018).
- Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. 2017. Persuasion of a privately informed receiver. *Econometrica* 85, 6 (2017), 1949–1964.
- Andrey Malenko and Nadya Malenko. 2018. Proxy advisory firms: The economics of selling information to voters. *Journal of Finance*, forthcoming (2018).
- Alex Smolin. 2019. Disclosure and Pricing of Attributes. *Available at SSRN 3318957* (2019).

Appendix A Revelation Principle

Consider any generic interactive protocol M , we construct a mechanism M' that extracts no less revenue, and the buyer truthfully report type θ and budget b at the root. The root of M' is just the buyer node that asks the buyer to report type θ and b . Then we add a subtree $T_{\theta,b}$ below the root for each (θ, b) . The subtree $T_{\theta,b}$ is just a copy of M in which the seller simulates the optimal strategy of the buyer with type θ and budget b at each buyer node. So $T_{\theta,b}$ only has seller nodes and transfer nodes. Clearly the buyer will truthfully report θ and b and the revenue of M' is the same as M .

Appendix B Omitted Proofs in Section 3

B.1 Proof of Theorem 3.1 Our proof has two main steps. First, we show that it suffices to consider one-round mechanisms. We employ a characterization of [Babaioff et al. \(2012\)](#) who show that there always exists an optimal one-round mechanism which is a pricing mappings mechanism (defined formally next) for a buyer *without* budget constraints. We generalize this characterization to the public budget case. Our second step is then to show that such a pricing mappings mechanism can always be converted to a consulting mechanism without revenue loss.

DEFINITION B.1. A pricing mappings mechanism with direct payment *proceeds as follows*:

1. Ask the buyer to report his type $\hat{\theta}$. His budget is publicly known as b .
2. According to the reported type $\hat{\theta}$, charge the buyer an amount of $t_{\hat{\theta}} \leq b$.
3. Choose a signaling scheme $S(\hat{\theta})$ to send a signal to the buyer.

Formally, the protocol tree of a pricing mappings mechanism has three layers: a buyer node root, $|\Theta|$ transfer nodes below the root, and each transfer node connects to a signaling scheme $S(\theta)$.

We first show that there always exists an optimal mechanism which is a pricing mappings mechanism with direct payment.

LEMMA B.1. *When the buyer has a public fixed budget b , and the buyer signal θ is independent from the seller signal ω , there always exists an optimal mechanism which is an IC pricing mappings mechanism with direct payment.*

Proof. According to Lemma 2.1, there exists an optimal IC revelation mechanism M . We first reduce M so that the buyer will always follow the protocol until a leaf is reached. If a type- θ buyer abort the protocol at some node n after truthfully reporting his type, we can just cut away the subtree under n . We claim that this will not change the truthfulness guarantee and the revenue will also remain the same. First of all, for a buyer with type θ , the protocol is the same because we only remove the part he is going to abort. For buyers with other type $\theta' \neq \theta$, if it is better for them to misreport θ when the subtree under n is cut away, then it should also be better for them to misreport θ and abort at n in the original protocol, which is contradictory to that M is IC. Consequently, the mechanism is still truthful and everyone will pay the same amount of money.

Now let's just assume the buyer will always follow the protocol until reaching a leaf, after truthfully reporting his type. We construct a pricing mappings mechanism with direct payment that extracts the same revenue as M . Let l_1, \dots, l_m be all the leaves of M that can be reached with positive probability. Let $p_\theta(\omega, l_i)$ be the probability that the protocol ends at leaf l_i when the seller's private signal is realized to ω and the buyer reports type θ at the beginning. Define $\tau(l_i)$ to be the total transfer on the path from the root to l_i . Since the signals are independent, the buyer's prior for ω is always the marginal distribution $\mu(\omega)$. As a result, for a buyer with any type, if he reports type θ and follows the protocol until a leaf is reached, his expected total payment will always be equal to $\sum_{l_i} \tau(l_i) \sum_{\omega} \mu(\omega) p_\theta(\omega, l_i)$. Note that the budget b should always allow the buyer to follow any path from the root to a leaf, in other words, the buyer will never have to quit the protocol because he cannot afford the payment at some step. Because if this is not true, a buyer with a certain type will have to quit before reaching a leaf, which is contradictory to our assumption.

Therefore we construct an optimal pricing mappings mechanism with direct payment as follows:

- $t_\theta = \sum_{l_i} \tau(l_i) \sum_{\omega} \mu(\omega) p_\theta(\omega, l_i)$,
- the signaling scheme $S(\theta)$ has leaves l_1, \dots, l_m and the transition function is just $p_\theta(\omega, l_i)$, so that the buyer's posterior at each l_i remains the same.

This mechanism is still truthful because a buyer who gains from misreporting in this mechanism will prefer to

misreport in the same way and then follow the protocol until reaching a leaf in M . It should also hold that $t_\theta \leq b$, since $\tau(l_i) \leq b$ for all l_i . It is easy to see that the revenue remains the same.

We emphasize that the above proof relies crucially on the assumptions that (1) the signals are independent and (2) there is a public fixed budget. If θ and ω are correlated, the buyer's prior for ω will be dependent on θ , which is equal to $\mu(\omega|\theta)$. This means that buyers with different types will have different expectations of how much they're going to pay. As a result, we cannot define t_θ using the common prior. If the budget is not fixed, it may not be true that the buyer will always have enough budget to follow M until reaching a leaf when he *misreports*. So in the original mechanism M , the buyer may not be able to get the same utility as in the new mechanism when he misreports a type. Consequently, the new mechanism may not be IC. \square

Although we have now reduced the problem to designing simple one-round mechanisms, the optimal signal schemes $S(\theta)$ can still be very complicated. Next, we show that any pricing mappings mechanism can be converted to a consulting mechanism without revenue loss. In other words, the signal sent by the seller can be converted to a recommendation about the best action a . Since there are $|A|$ possible actions, the signal schemes $S(\theta)$ will have no more than $|A|$ leaves. Similar results appeared in Proposition 1 of Bergemann et al. (2018).

LEMMA B.2. *For any IC pricing mappings mechanism with direct payment, there exists an IC consulting mechanism with direct payment that extracts no less revenue.*

Proof. Let M be any IC pricing mappings mechanism with direct payment. Let $S(\theta)$ be the signaling scheme for the type- θ buyer. Let s_1, \dots, s_m be the leaves of $S(\theta)$, and let a_1, \dots, a_m be the optimal actions of a type- θ buyer when the signal sent by $S(\theta)$ is realized to s_1, \dots, s_m respectively.

Suppose a_j, a_k both equal $\bar{a} \in A$ for some j and k , we show that we can without loss merge the two leaves s_j, s_k . That is, whenever the seller moves to s_j , she moves to s_k instead. The transition function of $S(\theta)$ is updated as follows: $p'(\omega, s_k) \leftarrow p(\omega, s_j) + p(\omega, s_k)$ and $p'(\omega, s_j) \leftarrow 0$ for all ω . The payments keep unchanged. Therefore, the seller's revenue will not change if the buyer still truthfully reports.

We claim that this new mechanism will not change the utility of a type- θ buyer. First of all, the optimal action for the type- θ buyer at leaf s_k is still \bar{a} . This is because, by definition, we have

$$\bar{a} = \operatorname{argmax}_{a \in A} \sum_{\omega} \mu(\omega) p(\omega, s_j) u(\omega, \theta, a),$$

$$\bar{a} = \operatorname{argmax}_{a \in A} \sum_{\omega} \mu(\omega) p(\omega, s_k) u(\omega, \theta, a).$$

Therefore $\operatorname{argmax}_{a \in A} \sum_{\omega} \mu(\omega) p'(\omega, s_j) u(\omega, \theta, a)$ must also equal \bar{a} . As a result, the expected utility of type- θ buyer will not change as $\sum_{\omega} \mu(\omega) p(\omega, s_j) u(\omega, \theta, \bar{a}) + \sum_{\omega} \mu(\omega) p(\omega, s_k) u(\omega, \theta, \bar{a}) = \sum_{\omega} \mu(\omega) p'(\omega, s_j) u(\omega, \theta, \bar{a}) + \sum_{\omega} \mu(\omega) p'(\omega, s_k) u(\omega, \theta, \bar{a})$.

Next we argue that the utility of any other type- θ' buyer will not increase when he misreports θ . This is because his utility from the original two leaves s_j, s_k is

$$\max_{a \in A} \sum_{\omega} \mu(\omega) p(\omega, s_j) u(\omega, \theta', a)$$

and

$$\max_{a \in A} \sum_{\omega} \mu(\omega) p(\omega, s_k) u(\omega, \theta', a).$$

Their sum is at least

$$\max_{a \in A} \sum_{\omega} \mu(\omega) [p(\omega, s_j) + p(\omega, s_k)] \cdot u(\omega, \theta', a).$$

To sum up, the new mechanism will not change the utility of type- θ buyer and will not increase the other buyers' utility when they misreport θ , and thus will remain IC and IR. Moreover, the revenue does not change. We can perform such merging operation until each leaf corresponds to a different action $a \in A$. To that end, each signal can be viewed as an "honest" action recommendation which will indeed maximizes the buyer's expected utility. This is precisely a consulting mechanism, as desired.

□

Lemma B.1 and B.2 together yield a proof of Theorem 3.1.

B.2 Proof of Theorem 3.2 The proof consists of two parts as the public budget case. We will first prove the existence of a one-round optimal mechanism, a *pricing mappings mechanism with deposit and return*, and then convert this mechanism to a *consulting mechanism with deposit and return*.

DEFINITION B.2. A pricing mappings mechanism with deposit and return *proceeds as follows*:

1. Ask the buyer to report his type $\hat{\theta}$ and budget \hat{b} , and then deposit his budget \hat{b} .
2. According to the reported type $\hat{\theta}$ and the deposit amount \hat{b} , return the buyer an amount of $\hat{b} - t_{\hat{\theta}, \hat{b}} \geq 0$.
3. Choose a signaling scheme $S(\hat{\theta}, \hat{b})$ to send a signal to the buyer.

Formally, the protocol tree of a pricing mappings mechanism has three layers: a buyer node root, $|\Theta| \cdot |B|$ transfer nodes below the root, and each transfer node connects to a signaling scheme $S(\theta, b)$.

LEMMA B.3. When the buyer has a private budget, and (θ, b) is independent from ω , there always exists an optimal mechanism which is an IC pricing mappings mechanism with deposit and return.

Proof. According to Lemma 2.1, there exists an optimal IC revelation mechanism M . As in the proof of Lemma B.1, we can reduce M so that the buyer will always follow the protocol until reaching a leaf. Now let's just assume the buyer will always follow the protocol until reaching a leaf, after truthfully reporting his type and budget. We construct a pricing mappings mechanism with deposit and return that extracts the same revenue as M . Let l_1, \dots, l_m be all the leaves of M that can be reached with positive probability. Let $p_{\theta, b}(\omega, l_i)$ be the probability that the protocol ends at leaf l_i when the seller's private signal is realized to ω and the buyer reports θ and b at the beginning. Define $\tau(l_i)$ to be the total transfer on the path from the root to l_i . Since (θ, b) is independent from ω , the buyer's prior for ω is always the marginal distribution $\mu(\omega)$. As a result, for a buyer with any type and budget, if he reports type θ, b and follows the protocol until a leaf is reached, his expected total payment will always be equal to $\sum_{l_i} \tau(l_i) \sum_{\omega} \mu(\omega) p_{\theta, b}(\omega, l_i)$. Therefore we construct an optimal pricing mappings mechanism with deposit and return as follows:

- $t_{\theta, b} = \sum_{l_i} \tau(l_i) \sum_{\omega} \mu(\omega) p_{\theta, b}(\omega, l_i)$,
- the signaling scheme $S(\theta, b)$ has leaves l_1, \dots, l_m and the transition function is just $p_{\theta, b}(\omega, l_i)$, so that the buyer's posterior at each l_i remains the same.

Note that the buyer can never overreport his budget in this mechanism because he is asked to deposit the full budget at the beginning. This guarantees the truthfulness of the new mechanism because a (θ, b) -buyer who gains from misreporting (θ', b') with $b' \leq b$ can also get higher utility in M by misreporting (θ', b') and following the protocol until a leaf is reached. Since $b \geq b'$, he should always have enough budget to reach the leaves in M if he reports b' . It should also hold that $t_{\theta, b} \leq b$, since for any leaf l_i that can be reached by (θ, b) -buyer with positive probability must have $\tau(l_i) \leq b$. Finally it is easy to see that the revenue remains the same. □

Same as the public budget case, any pricing mappings mechanism with deposit and return can be without loss converted into a mechanism that recommends

actions to the buyer, with the same *deposit and return* payment method.

LEMMA B.4. *For any IC pricing mappings mechanism with deposit and return, there exists an IC consulting mechanism with deposit and return that extracts no less revenue.*

The proof follows by a simple adaptation of the proof of Lemma 3.1, with $S(\theta, b)$ in place of $S(\theta)$.

Appendix C Proof of Theorem 4.1

We start by describing another class of mechanisms defined as follows, which is a strictly broader class than CM-probR mechanisms and has been proved to contain an optimal mechanism for the setting with no budget.

DEFINITION C.1. *A pricing outcomes mechanism with deposit and return (POM-depR) proceeds as follows:*

1. *The seller commits to a signaling scheme for each (θ, b) , described by $\{p_{\theta, b}(\omega, \sigma)\}_{\omega \in \Omega, \sigma \in \Sigma}$ where $p_{\theta, b}(\omega, \sigma)$ denotes the probability of sending signal σ upon seller signal ω , and a payment amount $t_{\theta, b}(\sigma) (\leq b)$ that depends on the realized signal σ .*
2. *Ask the buyer to report his type $\hat{\theta}$ and deposit his budget \hat{b} .*
3. *Sample a signal $\sigma \sim p_{\hat{\theta}, \hat{b}}(\omega)$ and sends σ to the buyer*
4. *Return $\hat{b} - t_{\hat{\theta}, \hat{b}}(\sigma)$ back to the buyer.*

As our first step, we generalize the result of Babaioff et al. (2012) and show that for budget-constrained buyers, pricing outcomes mechanisms with deposit and return (POM-depR) can also achieve optimality.

LEMMA C.1. *For a budget-constrained buyer, there always exists an IC pricing outcomes mechanism with deposit and return (POM-depR) that maximizes the revenue.*

Proof. According to Lemma 2.1, there exists an IC optimal revelation mechanism M . We can reduce M so that the buyer will always follow the protocol until reaching a leaf. Now let's just assume the buyer will always follow the protocol until reaching a leaf, after truthfully reporting his type and budget. We construct a pricing outcomes mechanism with deposit and return that extracts the same revenue as M . Let l_1, \dots, l_m be all the leaves of M that can be reached with positive probability. Let $p_{\theta, b}(\omega, l_i)$ be the probability that the protocol ends at leaf l_i when the seller's private signal is realized to ω and the buyer reports θ and b at the

beginning. Define $\tau(l_i)$ to be the total transfer on the path from the root to l_i . Then we construct an optimal pricing outcomes mechanism as follows:

- the signaling scheme $S(\theta, b)$ has leaves l_1, \dots, l_m and the transition function is just $p_{\theta, b}(\omega, l_i)$,
- the refund below leaf l_i is set to $b - \tau(l_i)$ so that the net transfer equals $\tau(l_i)$.

Since the buyer will never defect in M , the refund amount $b - \tau(l_i)$ is always nonnegative, so the buyer will not have chance to defect in the last step.

This pricing outcomes mechanism is still truthful because a buyer who gains from misreporting in this mechanism will prefer to misreport in the same way in M . Then since the net transfer remains the same on each path, and the probability of ending at each leaf remains the same, the revenue does not change. \square

Observe that any CM-probR mechanism is also a POM-depR mechanism since we can view each $[a, i] \in A \times \{+, -\}$ as a signal. However, CM-probR is significantly simpler — it only has two payments, i.e., b and $-M$, and it uses at most $2n$ signals. Our goal is to show that this much smaller class of mechanisms can still achieve optimality. For notational convenience, we will prove our result for the case of public budget. It is straightforward to generalize the argument to private budget, simply by adding additional incentive compatibility constraints to elicit truthful budget report (or equivalently, by viewing the buyer type θ as including the budget information).

Our starting point is an LP formulation proposed by Babaioff et al. (2012) with exponentially many variables but polynomially many constraints for computing the optimal POM-depR mechanism. We now adapt that LP to incorporate buyer budget constraints, which leads to an LP with exponentially many variables and constraints.

To describe the LP, we will need a different but equivalent definition of a signaling scheme. Recall that we use $\{p_{\theta}(\omega, \sigma)\}_{\omega \in \Omega, \sigma \in \Sigma}$ to denote a signaling scheme for buyer type θ (omitting the subscript of budget) where $\sum_{\sigma \in \Sigma} p_{\theta}(\omega, \sigma) = 1$ for each ω , which is drawn from the prior distribution $\{\mu(\omega)\}_{\omega \in \Omega}$. Equivalently, this signaling scheme can be viewed as a *convex decomposition* of the prior $\mu \in \Delta\Omega$ into a set of posterior distributions. In particular, the posterior distribution $\{\frac{p_{\theta}(\omega, \sigma)}{\sum_{\omega' \in \Omega} p_{\theta}(\omega', \sigma)}\}_{\omega \in \Omega}$ has convex coefficient $\sum_{\omega' \in \Omega} p_{\theta}(\omega', \sigma)$ (i.e., the probability).

As a result, a signaling scheme p_{θ} can be equivalently described by a variable $x_{\theta}(q) \in [0, 1]$ for each possible posterior distribution $q \in \Delta\Omega$, satisfying $\sum_q x_{\theta}(q) \cdot$

$q = \mu$ so that $\{x_\theta(q)\}_{q \in \Delta\Omega}$ indeed represents a convex decomposition of μ . Crucially, given any posterior distribution q , each buyer will interpret it differently due to different information the buyer type θ possess regarding ω . Concretely, as observed by Babaioff et al. (2012), buyer type θ will interpret any posterior q as $D_\theta q$ (up to a normalization factor) where D_θ is a $\Omega \times \Omega$ diagonal matrix with $\mu(\theta|\omega)$ for all ω in the diagonal. A minor issue is that we now have infinitely many variables since each $q \in \Delta\Omega$ corresponds to a variable $x_\theta(q)$. This turns out to be fine for our approach. However, for convenience, we invoke a simplification by Babaioff et al. (2012) who show that one can without loss of generality focus on a set Q^* of finite but *exponentially* many possible posterior beliefs q 's since only these q 's can show up in the optimal solution.

Now we are ready to state the linear program for computing the optimal POM-depR mechanism, with variable $x_\theta(q)$ and $\tilde{t}_\theta(q) = x_\theta(q) \cdot t_\theta(q)$ where $t_\theta(q)$ is the corresponding buyer total payment for the realized signal inducing posterior q (negative $t_\theta(q)$ means that the seller will pay the buyer).

$$(3.5) \quad \begin{aligned} \max \quad & \sum_{\theta, q} (\mathbf{1}^t D_\theta q) \cdot \tilde{t}_\theta(q) \\ \text{s.t.} \quad & \sum_q [v_\theta(D_\theta q) x_\theta(q) - (\mathbf{1}^t D_\theta q) \tilde{t}_\theta(q)] \\ & \geq v_\theta(D_\theta \mu), \quad \text{for } \theta \in \Theta \\ & \sum_q [v_\theta(D_\theta q) x_\theta(q) - (\mathbf{1}^t D_\theta q) \tilde{t}_\theta(q)] \\ & \geq \sum_q [v_{\theta'}(D_\theta q) x_{\theta'}(q) - (\mathbf{1}^t D_\theta q) \tilde{t}_{\theta'}(q)], \quad \text{for } \theta' \neq \theta \\ & \sum_q x_\theta(q) \cdot q = \mu, \quad \text{for } \theta \in \Theta \\ & b \cdot x_\theta(q) \geq \tilde{t}_\theta(q) \geq -M \cdot x_\theta(q), \quad \text{for } \theta, q \in Q^* \\ & x_\theta(q) \geq 0, \quad \text{for } \theta, q \in Q^* \end{aligned}$$

Here $v_\theta(q) = \max_{a' \in A} \mathbf{E}_{\omega \sim q} u(\omega, \theta, a')$ denotes the type- θ buyer's optimal expected utility when he believes that ω follows distribution q . Note that q here does not have to be normalized to have l_1 norm 1 since $v_\theta(q)$ is linear in any re-scaling factor of q .

We mention a subtle issue here regarding the above LP formulation. Our LP (3.5) is similar to an LP formulation given in Babaioff et al. (2012) for computing the optimal pricing outcomes mechanism *without* budget constraints. More concretely, the LP of Babaioff et al. (2012) does not have the constraint $b \cdot x_\theta(q) \geq \tilde{t}_\theta(q) \geq -M \cdot x_\theta(q)$. However, that LP does *not* exactly compute the optimal pricing outcomes mechanism since it may have an optimal solution such that $\tilde{t}_\theta(q) \neq 0$ but $x_\theta(q) = 0$ for some θ, q , which does not correspond to any pricing outcomes mechanisms. This discrepancy is due to their variable exchange $\tilde{t}_\theta(q) = x_\theta(q) t_\theta(q)$ which does not result in an equivalent formulation since the constraint $x_\theta(q) = 0 \Rightarrow \tilde{t}_\theta(q) = 0$ cannot not be enforced after the variable change (and cannot be casted as any linear constraint). To deal with this issue, Babaioff et al. (2012) introduced a technique of

slightly perturbing the optimal solution and look for an ϵ -optimal solution which always satisfies $x_\theta(q) \neq 0$ whenever $\tilde{t}_\theta(q) \neq 0$. As a result, this technique can only give an approximately optimal pricing outcomes mechanism, and may result in infinitely large payment as the approximation error ϵ tends to 0.

Interestingly, it turns out that the budget constraints help us to get rid of the above discrepancy. In particular, our LP (3.5) exactly computes the optimal pricing outcomes mechanism. This is because the budget constraint $b \cdot x_\theta(q) \geq \tilde{t}_\theta(q) \geq -M \cdot x_\theta(q)$ automatically implies that $\tilde{t}_\theta(q) = 0$ when $x_\theta(q) = 0$, assuming all budgets are finite. We summarize this observation in the following lemma.

LEMMA C.2. *The optimal solution to LP (3.5) corresponds exactly to an optimal pricing outcomes mechanism.*

The main challenge for solving LP (3.5) is that it has both exponentially many variables and constraints (due to budget constraints) since the set Q^* is exponentially large. It is not even clear that it has an optimal solution with polynomial size, let alone solving it to output the optimal solution in polynomial time. We consider the dual program of LP (3.5), formulated as in (3.4) where $\{\alpha_\theta\}_{\theta \in \Theta}, \{\lambda_{\theta, \theta'}\}_{\theta \neq \theta'}, \{y_\theta \in \mathbb{R}^\Omega\}_{\theta \in \Theta}$ and $\{\gamma_{\theta, q}, \xi_{\theta, q}\}_{\theta, q}$ are the dual variables for the first, second, third and fourth sets of constraints in LP (3.5), respectively.

Obviously, LP (3.4) is not tractable neither as it still has exponentially many variables and constraints. Our next main step is to prove that the variables $\gamma_{\theta, q}, \xi_{\theta, q}$ in LP (3.4) can be eliminated to induce a different linear program which has the same optimal objective as LP (3.4) but only has polynomially many variables. For notational convenience, we define the following functions:

$$\begin{aligned} g_\theta(\lambda, \alpha; q) &= (\mathbf{1}^t D_\theta q)(1 - \alpha_\theta - \sum_{\theta' \neq \theta} \lambda_{\theta, \theta'}) + \sum_{\theta' \neq \theta} \lambda_{\theta', \theta} (\mathbf{1}^t D_{\theta'} q) \\ h_\theta(\lambda, \alpha; q) &= v_\theta(D_\theta q)(\alpha_\theta + \sum_{\theta' \neq \theta} \lambda_{\theta, \theta'}) - \sum_{\theta' \neq \theta} \lambda_{\theta', \theta} v_{\theta'}(D_{\theta'} q). \end{aligned}$$

Note that g_θ, h_θ are both *linear* functions of α, λ for any given $q \in Q^*$. Using the above simplification, we can re-write LP (3.4) as in (3.6).

LEMMA C.3. *LP (3.6) has the same optimal objective value as the following LP with only variables α, λ, y :*

$$(3.7) \quad \begin{aligned} \min \quad & \sum_\theta \mu \cdot y_\theta - v_\theta(D_\theta \mu) \cdot \alpha_\theta \\ \text{s.t.} \quad & y_\theta \cdot q - b \cdot g_\theta - h_\theta(\lambda, \alpha, q) \geq 0, \quad \text{for } \theta, q \\ & y_\theta \cdot q + M \cdot g_\theta - h_\theta(\lambda, \alpha, q) \geq 0, \quad \text{for } \theta, q \\ & \alpha, \lambda \geq 0 \end{aligned}$$

Proof. The main idea of the proof is to eliminate the variables $\gamma_{\theta, q}, \xi_{\theta, q}$.

$$(3.4) \quad \begin{aligned} \min \quad & \sum_{\theta} \mu \cdot y_{\theta} - v_{\theta}(D_{\theta}\mu) \cdot \alpha_{\theta} \\ \text{s.t.} \quad & \gamma_{\theta,q} - \xi_{\theta,q} - (\mathbf{1}^t D_{\theta}q)(1 - \alpha_{\theta} - \sum_{\theta' \neq \theta} \lambda_{\theta',\theta'}) - \sum_{\theta' \neq \theta} \lambda_{\theta',\theta}(\mathbf{1}^t D_{\theta'}q) = 0, \quad \text{for } \theta, q \\ & y_{\theta} \cdot q - b\gamma_{\theta,q} - M\xi_{\theta,q} - v_{\theta}(D_{\theta}q)(\alpha_{\theta} + \sum_{\theta' \neq \theta} \lambda_{\theta',\theta'}) + \sum_{\theta' \neq \theta} \lambda_{\theta',\theta} v_{\theta'}(D_{\theta'}q) \geq 0, \quad \text{for } \theta, q \\ & \alpha, \lambda, \gamma, \xi \geq 0 \end{aligned}$$

Figure 5: Dual LP of (3.5).

$$(3.6) \quad \begin{aligned} \min \quad & \sum_{\theta} \mu \cdot y_{\theta} - v_{\theta}(D_{\theta}\mu) \cdot \alpha_{\theta} \\ \text{s.t.} \quad & \gamma_{\theta,q} - \xi_{\theta,q} - g_{\theta}(\lambda, \alpha; q) = 0, \quad \text{for } \theta, q \\ & y_{\theta} \cdot q - b\gamma_{\theta,q} - M\xi_{\theta,q} - h_{\theta}(\lambda, \alpha; q) \geq 0, \quad \text{for } \theta, q \\ & \alpha, \lambda, \gamma, \xi \geq 0 \end{aligned}$$

Figure 6: Re-Writing of LP (3.4)

First, we claim that LP (3.6) always admits an optimal solution in which either $\gamma_{\theta,q} = 0$ or $\xi_{\theta,q} = 0$ for any θ, q . In particular, let $(\alpha^*, \lambda^*, y^*, \gamma^*, \xi^*)$ be any optimal solution to LP (3.6). We construct a new optimal solution which satisfies the desired property, as follows. For any index pair (θ, q) :

- If $\gamma_{\theta,q}^* \geq \xi_{\theta,q}^*$, define $\gamma'_{\theta,q} = \gamma_{\theta,q}^* - \xi_{\theta,q}^*$ and $\xi'_{\theta,q} = 0$;
- If $\gamma_{\theta,q}^* < \xi_{\theta,q}^*$, define $\gamma'_{\theta,q} = 0$ and $\xi'_{\theta,q} = \xi_{\theta,q}^* - \gamma_{\theta,q}^*$.

We claim that $(\alpha^*, \lambda^*, y^*, \gamma', \xi')$ is a feasible and optimal solution to LP (3.4). The new variable values achieve the same (optimal) objective value as $(\alpha^*, \lambda^*, y^*, \gamma^*, \xi^*)$ since the objective function only depends on α^*, y^* . We only need to argue that $(\alpha^*, \lambda^*, y^*, \gamma', \xi')$ is feasible. The non-negativity of $\gamma'_{\theta,q}, \xi'_{\theta,q}$ is obvious from construction. The first set of constraint is satisfied because $\gamma_{\theta,q}^* - \xi_{\theta,q}^* = \gamma'_{\theta,q} - \xi'_{\theta,q}$ in our construction and g_{θ} does not depend on $\gamma_{\theta,q}, \xi_{\theta,q}$. The second set of constraints are also feasible because $y_{\theta}^* \cdot q - b\gamma_{\theta,q}^* - M\xi_{\theta,q}^* - h_{\theta}(\lambda^*, \alpha^*; q) \geq 0$, and the constructed variable value $\gamma'_{\theta,q}, \xi'_{\theta,q}$ satisfy $\gamma'_{\theta,q} \leq \gamma_{\theta,q}^*, \xi'_{\theta,q} \leq \xi_{\theta,q}^*$ and thus will only increase the left-hand side of the above inequality. Therefore the “ \geq ” inequality still holds.

Thanks to the above argument, we can impose the complementarity constraints $\gamma_{\theta,q} \cdot \xi_{\theta,q} = 0$ into LP (3.6), without changing its optimal objective value. As a result, the equality constraint $\gamma_{\theta,q} - \xi_{\theta,q} - g_{\theta}(\lambda, \alpha; q) = 0$ in LP (3.6), together with $\gamma_{\theta,q} \cdot \xi_{\theta,q} = 0$, can be used to uniquely determine the value of $\gamma_{\theta,q}$ and $\xi_{\theta,q}$. In particular, if $g_{\theta}(\lambda, \alpha; q) \geq 0$, we must have $\gamma_{\theta,q} = g_{\theta}(\lambda, \alpha; q)$ and $\xi_{\theta,q} = 0$. Otherwise, we must have $\gamma_{\theta,q} = 0$ and $\xi_{\theta,q} = -g_{\theta}(\lambda, \alpha; q)$. To sum up, $\gamma_{\theta,q}$ and $\xi_{\theta,q}$ can be uniquely determined as follows:

$$\begin{aligned} \gamma_{\theta,q} &= \mathbf{1} \{g_{\theta}(\lambda, \alpha, q) \geq 0\} g_{\theta}(\lambda, \alpha, q), \\ \xi_{\theta,q} &= -\mathbf{1} \{g_{\theta}(\lambda, \alpha, q) \leq 0\} g_{\theta}(\lambda, \alpha, q), \end{aligned}$$

where $\mathbf{1}(X)$ is the indicator function which equals 1 when condition X holds.

We can now plug in the above value of $\gamma_{\theta,q}, \xi_{\theta,q}$ into the left-hand side of the second constraint of LP (3.6), and obtain the following inequality constraint for any θ, q :

$$(3.8) \quad \begin{aligned} & y_{\theta} \cdot q - b \cdot \mathbf{1} \{g_{\theta}(\lambda, \alpha, q) \geq 0\} g_{\theta}(\lambda, \alpha, q) \\ & + M \cdot \mathbf{1} \{g_{\theta}(\lambda, \alpha, q) \leq 0\} g_{\theta}(\lambda, \alpha, q) - h_{\theta}(\lambda, \alpha; q) \geq 0 \end{aligned}$$

Observe that

$$\begin{aligned} & -b \cdot \mathbf{1} \{g_{\theta}(\lambda, \alpha, q) \geq 0\} g_{\theta}(\lambda, \alpha, q) \\ & + M \cdot \mathbf{1} \{g_{\theta}(\lambda, \alpha, q) \leq 0\} g_{\theta}(\lambda, \alpha, q) \\ & = \min\{-b \cdot g_{\theta}(\lambda, \alpha, q), M \cdot g_{\theta}(\lambda, \alpha, q)\} \end{aligned}$$

which is the minimum of two linear functions and thus is concave. As a result, Inequality constraint (3.8) can be re-written as the following

$$\begin{aligned} & y_{\theta} \cdot q - b \cdot \mathbf{1} \{g_{\theta}(\lambda, \alpha, q) \geq 0\} g_{\theta}(\lambda, \alpha, q) + M \cdot \\ & \quad \mathbf{1} \{g_{\theta}(\lambda, \alpha, q) \leq 0\} g_{\theta}(\lambda, \alpha, q) - h_{\theta}(\lambda, \alpha; q) \geq 0 \\ \Leftrightarrow & y_{\theta} \cdot q + \min\{-b \cdot g_{\theta}(\lambda, \alpha, q), M \cdot g_{\theta}(\lambda, \alpha, q)\} \\ & \quad - h_{\theta}(\lambda, \alpha; q) \geq 0 \\ \Leftrightarrow & \min\{-b \cdot g_{\theta}(\lambda, \alpha, q), M \cdot g_{\theta}(\lambda, \alpha, q)\} \\ & \quad \geq h_{\theta}(\lambda, \alpha; q) - y_{\theta} \cdot q \\ \Leftrightarrow & -b \cdot g_{\theta}(\lambda, \alpha, q) \geq h_{\theta}(\lambda, \alpha; q) - y_{\theta} \cdot q \\ & \quad \& M \cdot g_{\theta}(\lambda, \alpha, q) \geq h_{\theta}(\lambda, \alpha; q) - y_{\theta} \cdot q \end{aligned}$$

Substituting Inequality constraint (3.8) by the above to inequalities, we obtain LP (3.7) which has the same optimal objective as LP (3.6). \square

We now plugging the concrete form of g and h into LP (3.7) and obtain (3.9), which has the same objective as LP (3.4) but with only polynomially many constraints.

At this point, a natural approach is to design an efficient separation oracle for LP (3.9) and then employ the celebrated Ellipsoid method to solve it with a hope to also obtain an optimal solution to LP (3.5). Due to the reasons we mentioned in the main body, here we adopt a different approach and instead turn to the dual program of LP (3.9). However, the

$$\begin{aligned}
(3.9) \quad & \min \quad \sum_{\theta} \mu \cdot y_{\theta} - v_{\theta}(D_{\theta}\mu) \cdot \alpha_{\theta} \\
& \text{s.t.} \quad -\alpha_{\theta} \left[v_{\theta}(D_{\theta}q) - b(\mathbf{1}^t D_{\theta}q) \right] - \sum_{\theta'} \lambda_{\theta, \theta'} \left[v_{\theta}(D_{\theta}q) - b(\mathbf{1}^t D_{\theta}q) \right] + \\
& \quad \sum_{\theta'} \lambda_{\theta', \theta} \left[v_{\theta'}(D_{\theta'}q) - b(\mathbf{1}^t D_{\theta'}q) \right] + y_{\theta}q \geq b(\mathbf{1}^t D_{\theta}q), \quad \text{for } \theta, q \\
& \quad -\alpha_{\theta} \left[v_{\theta}(D_{\theta}q) + M(\mathbf{1}^t D_{\theta}q) \right] - \sum_{\theta'} \lambda_{\theta, \theta'} \left[v_{\theta}(D_{\theta}q) + M(\mathbf{1}^t D_{\theta}q) \right] + \\
& \quad \sum_{\theta'} \lambda_{\theta', \theta} \left[v_{\theta'}(D_{\theta'}q) + M(\mathbf{1}^t D_{\theta'}q) \right] + y_{\theta}q \geq -M(\mathbf{1}^t D_{\theta}q), \quad \text{for } \theta, q \\
& \quad \alpha, \lambda \geq 0
\end{aligned}$$

Figure 7: Re-Writing of LP (3.7)

benefit of considering the dual program is not clear at this point yet since it will have exponentially many variables. Interestingly, we show that by exploring some “local convexity” structure of the formulation the dual program actually can be solved much more efficiently.

We start by deriving the dual of LP (3.9) as in (3.10), with variable $x_{\theta}^{+}(q), x_{\theta}^{-}(q)$ for the two different sets of constraints. Interestingly, it turns out that solutions to LP (3.10) have a natural interpretation as a *special type* of pricing outcomes mechanism, as formalized in the following lemma.

LEMMA C.4. *Any feasible solution to LP (3.10) is an incentive compatible pricing outcomes mechanism which uses only two different payment methods: either buyer pays b to the seller or the seller pays M to the buyer.*

Proof. The proof is almost evident from the formulation of LP (3.10). Any feasible solution to LP (3.10) corresponds to, for any buyer type θ , a convex decomposition of the prior μ where the posterior is q with probability $x_{\theta}^{+}(q) + x_{\theta}^{-}(q)$. Moreover, $\frac{x_{\theta}^{+}(q)}{x_{\theta}^{+}(q) + x_{\theta}^{-}(q)}$ of the time the buyer pays b to the seller in which case the buyer’s remaining utility is $[v_{\theta}(D_{\theta}q) - b(\mathbf{1}^t D_{\theta}q)]$ (up to a normalization factor). Otherwise, the seller pays M to the buyer in which case the buyer’s remaining utility is $[v_{\theta}(D_{\theta}q) + M(\mathbf{1}^t D_{\theta}q)]$. As a result, it is easy to see that the first and second sets of constraints in LP (3.10) are precisely the individual rationality and incentive compatibility constraints. Moreover, it is also easy to see that the objective function is precisely the seller’s revenue under this mechanism. This concludes the proof. \square

Next, we show that LP (3.10) always admits an optimal solution which is a **CM-probR**. More precisely, $x_{\theta}^{+}, x_{\theta}^{-}$ support on at most $|A|$ different posteriors, each resulting in a different best response action for buyer type θ .

LEMMA C.5. *There always exists an optimal solution to LP (3.10) where, for any $\theta \in \Theta$, both x_{θ}^{+} and x_{θ}^{-} are*

*supported on at most $|A|$ different posterior q ’s, each inducing a different best response action for buyer type θ . In other words, LP (3.10) always admits an optimal solution which is a **CM-probR** as described in Definition 4.1 by viewing each posterior distribution as an action recommendation.*

Proof. We prove the lemma by converting any optimal solution to LP (3.10) to a feasible and optimal solution that satisfies the desired property. Let $\mathbf{x} = \{x_{\theta}^{+}(q), x_{\theta}^{-}(q)\}_{\theta \in \Theta, q \in Q^*}$ be an arbitrary optimal solution to LP (3.10). Let $q_1, q_2 \in Q^*$ be two posteriors such that $x_{\theta}^{+}(q_1), x_{\theta}^{+}(q_2) > 0$ and buyer utility $v_{\theta}(D_{\theta}q_1), v_{\theta}(D_{\theta}q_2)$ are both achieved by the same best action \hat{a} . We now adjust \mathbf{x} as follows. Denote

$$q = \frac{1}{x_{\theta}^{+}(q_1) + x_{\theta}^{+}(q_2)} \left[x_{\theta}^{+}(q_1) \cdot q_1 + x_{\theta}^{+}(q_2) \cdot q_2 \right] \in \Delta_{\Omega}$$

and let $\hat{x}_{\theta}^{+}(q) = x_{\theta}^{+}(q_1) + x_{\theta}^{+}(q_2)$ whereas $\hat{x}_{\theta}^{+}(q_1) = 0, \hat{x}_{\theta}^{+}(q_2) = 0$. Observe that the buyer θ ’s optimal action on posterior q is still \hat{a} because $\hat{a} = \arg \max_{a \in A} \sum_{\omega \in \Omega} (q_1)_{\omega} u(w, \theta, a)$ and $\hat{a} = \arg \max_{a \in A} \sum_{\omega \in \Omega} (q_2)_{\omega} u(w, \theta, a)$ imply that $\hat{a} = \arg \max_{a \in A} \sum_{\omega \in \Omega} [c_1(q_1)_{\omega} + c_2(q_2)_{\omega}] u(w, \theta, a)$ for any constants $c_1, c_2 > 0$. As a result, we have

$$(3.11) \quad c_1 \cdot v_{\theta}(D_{\theta}q_1) + c_2 \cdot v_{\theta}(D_{\theta}q_2) = v_{\theta}(D_{\theta}[c_1 \cdot q_1 + c_2 \cdot q_2])$$

We now argue that our adjustment to \mathbf{x} maintains feasibility and optimality. The adjustment does not change the objective because

$$\begin{aligned}
& b(\mathbf{1}^t D_{\theta}q) \cdot \hat{x}_{\theta}^{+}(q) \\
&= b \left[\mathbf{1}^t D_{\theta} \cdot (\hat{x}_{\theta}^{+}(q) \cdot q) \right] \\
&= b \left[\mathbf{1}^t D_{\theta} \cdot (x_{\theta}^{+}(q_1) \cdot q_1 + x_{\theta}^{+}(q_2) \cdot q_2) \right] \\
&= b(\mathbf{1}^t D_{\theta}q_1) \cdot x_{\theta}^{+}(q_1) + b(\mathbf{1}^t D_{\theta}q_2) \cdot x_{\theta}^{+}(q_2)
\end{aligned}$$

We thus only need to prove that it remains feasible. The adjustment does not change the first constraint,

$$\begin{aligned}
(3.10) \quad & \max \quad \sum_{\theta, q} \left[b(\mathbf{1}^t D_{\theta q}) \cdot x_{\theta}^+(q) - M(\mathbf{1}^t D_{\theta q}) \cdot x_{\theta}^-(q) \right] \\
& \text{s.t.} \quad \sum_q \left[v_{\theta}(D_{\theta q}) - b(\mathbf{1}^t D_{\theta q}) \right] x_{\theta}^+(q) + [v_{\theta}(D_{\theta q}) + M(\mathbf{1}^t D_{\theta q})] x_{\theta}^-(q) \geq v_{\theta}(D_{\theta} \mu), \quad \text{for } \theta \in \Theta \\
& \sum_q \left[v_{\theta'}(D_{\theta' q}) - b(\mathbf{1}^t D_{\theta' q}) \right] x_{\theta'}^+(q) + [v_{\theta'}(D_{\theta' q}) + M(\mathbf{1}^t D_{\theta' q})] x_{\theta'}^-(q) \\
& \quad \geq \sum_q \left[v_{\theta'}(D_{\theta' q}) - b(\mathbf{1}^t D_{\theta' q}) \right] x_{\theta'}^+(q) + [v_{\theta'}(D_{\theta' q}) + M(\mathbf{1}^t D_{\theta' q})] x_{\theta'}^-(q), \quad \text{for } \theta' \neq \theta \\
& \sum_q [x_{\theta}^+(q) + x_{\theta}^-(q)] \cdot q = \mu, \quad \text{for } \theta \in \Theta \\
& x_{\theta}(q) \geq 0, \quad \text{for } \theta, q
\end{aligned}$$

Figure 8: Dual LP of (3.9).

because

$$\begin{aligned}
& \left[v_{\theta}(D_{\theta q}) - b(\mathbf{1}^t D_{\theta q}) \right] \widehat{x}_{\theta}^+(q) \\
= & u \left[D_{\theta} \cdot [\widehat{x}_{\theta}^+(q) \cdot q] \right] - b \left[\mathbf{1}^t D_{\theta} \cdot (\widehat{x}_{\theta}^+(q) \cdot q) \right] \\
= & u \left[D_{\theta} \cdot [x_{\theta}^+(q_1) \cdot q_1 + x_{\theta}^+(q_2) \cdot q_2] \right] \\
& - b \left[\mathbf{1}^t D_{\theta} \cdot (x_{\theta}^+(q_1) \cdot q_1 + x_{\theta}^+(q_2) \cdot q_2) \right] \\
= & \left[v_{\theta}(D_{\theta q_1}) - b(\mathbf{1}^t D_{\theta q_1}) \right] x_{\theta}^+(q_1) \\
& + \left[v_{\theta}(D_{\theta q_2}) - b(\mathbf{1}^t D_{\theta q_2}) \right] x_{\theta}^+(q_2)
\end{aligned}$$

where the last equality is due to Inequality (3.11).

Next, we show that the second set of constraints also remain feasible under our adjustment. The above argument shows that the value of its left-hand side remains the same for any θ, q . We now show that the adjustment will *not increase* the value of the right-hand side. The only possibility that the value of the right-hand side could change is when θ' in $x_{\theta'}^+(q)$ correspond to our focus of buyer type θ since all other $x_{\theta'}^+(q)$'s are not changed. In this case, we have for any $\theta' \neq \theta$

$$\begin{aligned}
& [v_{\theta'}(D_{\theta' q}) - b(\mathbf{1}^t D_{\theta' q})] \widehat{x}_{\theta'}^+(q) \\
\leq & \left[v_{\theta'}(D_{\theta' q_1}) \cdot \frac{x_{\theta'}^+(q_1)}{x_{\theta'}^+(q_1) + x_{\theta'}^+(q_2)} \right. \\
& \left. + v_{\theta'}(D_{\theta' q_2}) \cdot \frac{x_{\theta'}^+(q_2)}{x_{\theta'}^+(q_1) + x_{\theta'}^+(q_2)} \right] \cdot \widehat{x}_{\theta'}^+(q) \\
& - b(\mathbf{1}^t D_{\theta' q}) \cdot \widehat{x}_{\theta'}^+(q) \\
= & \left[v_{\theta'}(D_{\theta' q_1}) \cdot x_{\theta'}^+(q_1) + v_{\theta'}(D_{\theta' q_2}) \cdot x_{\theta'}^+(q_2) \right] \\
& - b(\mathbf{1}^t D_{\theta' q_1}) \cdot x_{\theta'}^+(q_1) - b(\mathbf{1}^t D_{\theta' q_2}) \cdot x_{\theta'}^+(q_2)
\end{aligned}$$

where the inequality is due to the convexity of $v_{\theta'}$ function and our definition of $q = \frac{1}{x_{\theta'}^+(q_1) + x_{\theta'}^+(q_2)} \left[x_{\theta'}^+(q_1) \cdot q_1 + x_{\theta'}^+(q_2) \cdot q_2 \right]$ as a convex combination of q_1, q_2 . This

shows that our adjustment will not increase the value of the right-hand side of the second constraint for any θ, θ' . It is easy to see that the third set of constraints also hold. As a result, the adjusted solution remains feasible.

The above adjustment can be done for any θ and for both x_{θ}^+ and x_{θ}^- . Therefore, ultimately, we can obtain an optimal solution with the desired property as described in the lemma. \square

The following lemma concludes our proof of the characterization part in Theorem 4.1.

LEMMA C.6. *The optimal revenue can be achieved by an IC Consulting Mechanism with Probabilistic Return (CM-probR).*

Proof. This is because the optimal objective of LP (3.10) — i.e., the maximum revenue achieved by the optimal CM-probR — equals the optimal objective of LP (3.5), i.e., the maximum revenue achieved by the optimal pricing outcomes mechanisms, which is proved to be an optimal mechanism by Lemma C.1. \square

Appendix D Omitted Proofs in Section 5

D.1 Proof of Lemma 5.1 We first prove that ε -IC can be guaranteed by the IC constraint in the LP. Consider any $\mu(\omega, \theta, b)$, any (θ, b) . The expected payoff of a type- (θ, b) buyer when reporting θ, b is equal to

$$\begin{aligned}
U_{\theta, b}(\theta, b) &= \\
& \mathbb{E}_{\omega_1 \sim \mu(\omega|\theta, b)} \mathbb{E}_{T \sim \mu^{n-1}} \mathbb{E}_{\mathcal{M}_S} [u(\omega_1, \theta, a_S(\theta, b)) - t_S(\theta, b)].
\end{aligned}$$

Consider this expected payoff conditioning on $S_{\theta, b}$. Let $\mathcal{U}(S_{\theta, b})$ be the uniform distribution over $S_{\theta, b}$. Then

$$\begin{aligned}
& U_{\theta, b}(\theta, b | S_{\theta, b}) \\
&= \mathbb{E}_{\omega_1 \sim \mu(\omega|\theta, b)} \mathbb{E} [u(\omega_1, \theta, a_S(\theta, b)) - t_S(\theta, b) | S_{\theta, b}] \\
&= \mathbb{E}_{\omega_1 \sim \mathcal{U}(S_{\theta, b})} \mathbb{E} [u(\omega_1, \theta, a_S(\theta, b)) - t_S(\theta, b) | S_{\theta, b}]
\end{aligned}$$

Conditioning on $S_{\theta,b}$, we should have $\omega_1 \sim \mathcal{U}(S_{\theta,b})$ because $\omega_1 \in S_{\theta,b}$ and the elements in $S_{\theta,b}$ are independently drawn from $\mu(\omega|\theta, b)$. Recall that $\hat{\mu}(\omega|\theta, b)$ is defined to be the uniform distribution over $S_{\theta,b}$. So our IC constraints for type- (θ, b) buyer guarantees that when ω_1 is drawn from the uniform distribution over $S_{\theta,b}$, it is ε -optimal to truthfully report (θ, b) , i.e.,

$$U_{\theta,b}(\theta, b|S_{\theta,b}) \geq U_{\theta,b}(\theta', b'|S_{\theta,b}) - \varepsilon$$

for all $(\theta', b') \neq (\theta, b)$. Here

$$\begin{aligned} & U_{\theta,b}(\theta', b'|S_{\theta,b}) \\ &= \mathbb{E}_{\omega_1 \sim \mathcal{U}(S_{\theta,b})} \mathbb{E} [u(\omega_1, \theta, a'(a_S(\theta', b')))) - t_S(\theta', b') | S_{\theta,b}] \end{aligned}$$

represents the conditional expected utility when reporting (θ', b') and then taking arbitrary action based on the recommendation. Take expectation over $S_{\theta,b}$, we prove ε -IC.

The proofs for ε -IR and ε -obedience are basically the same. ε -IR is guaranteed by the IR constraints for each $S_{\theta,b}$. ε -obedience is guaranteed by the OB constraints for each $S_{\theta,b}$.

D.2 Proof of Theorem 5.1 We first show that the optimal solution of (4.2) is a feasible solution of (5.3) with high probability.

LEMMA D.1. *With probability at least $1 - \delta/2$, the optimal solution of (4.2) p^* is a feasible solution of the LP (5.3), as long as*

$$\begin{aligned} n &\geq \Theta \left(\ln(G/\delta) \cdot \max \left\{ \frac{|A|^2}{\varepsilon^2 \cdot \mu_{\min}}, \frac{1}{\mu_{\min}^2} \right\} \right) \\ &= \Theta \left(\frac{|A|^2 \cdot \ln(G/\delta)}{\varepsilon^2} \right). \end{aligned}$$

where $G = \max\{|\Theta|, |B|, |A|\}$ and $\mu_{\min} = \min_{\theta,b} \mu(\theta, b)$.

Proof. Define $\mu_{\min} = \min_{\theta,b} \mu(\theta, b)$. We first show that with probability at least $1 - \delta/4$, we collect at least $\frac{n \cdot \mu_{\min}}{2}$ samples for each buyer type θ, b . By the Chernoff bound, for each θ, b , with probability at least $1 - \delta/(4|\Theta| \cdot |B|)$,

$$\begin{aligned} \mu(\theta, b) - \frac{\sum_{i=1}^n \mathbf{1}\{(\theta_i, b_i) = (\theta, b)\}}{n} &\leq \sqrt{\frac{\ln(4|\Theta| \cdot |B|/\delta)}{2n}} \\ \implies \frac{\sum_{i=1}^n \mathbf{1}\{(\theta_i, b_i) = (\theta, b)\}}{n} &\geq \mu(\theta, b) - \sqrt{\frac{\ln(4|\Theta| \cdot |B|/\delta)}{2n}} \\ \implies \frac{\sum_{i=1}^n \mathbf{1}\{(\theta_i, b_i) = (\theta, b)\}}{n} &\geq \mu_{\min} - \sqrt{\frac{\ln(4|\Theta| \cdot |B|/\delta)}{2n}} \end{aligned}$$

So by the union bound when $n \geq \frac{2 \ln(4|\Theta| \cdot |B|/\delta)}{\mu_{\min}^2}$ we have $\frac{\sum_{i=1}^n \mathbf{1}\{(\theta_i, b_i) = (\theta, b)\}}{n} \geq \mu_{\min}/2$ for all θ, b with probability $1 - \delta/4$.

Now assume that we collect at least $\frac{n \cdot \mu_{\min}}{2}$ samples for each buyer type θ, b . Let $p_{\theta,b}^{*,\circ}(\omega, a)$ be the variables of the optimal solution of (4.2) p^* , where $\circ \in \{+, -\}$. To prove that p^* is a feasible solution of (5.3) with high probability, it suffices to show that with high probability, we have

$$\left| \sum_{\omega} (\hat{\mu}(\omega|\theta, b) - \mu(\omega|\theta, b)) p_{\theta,b}^{*,\circ}(\omega, a) (u(\omega, \theta, a) - t^\circ) \right| \leq \frac{1}{4|A|}$$

for all $\theta, \theta', b, b', a, a'$. Now consider the buyer with type θ and budget b . Let K be the number of samples with $(\theta_i, b_i) = (\theta, b)$. Consider fixed θ', b', a, a' and define $h_{\omega}^{\circ} = p_{\theta',b'}^{*,\circ}(\omega, a) (u(\omega, \theta, a') - t^\circ)$. According to the Chernoff bound, with probability at least $1 - \delta/(8|\Theta|^2|B|^2|A|^2)$,

$$\begin{aligned} & \left| \frac{\sum_{i=1}^n \mathbf{1}\{(\theta_i, b_i) = (\theta, b)\} h_{\omega_i}^{\circ}}{K} - \sum_{\omega} \mu(\omega|\theta, b) h_{\omega}^{\circ} \right| \\ & \leq \sqrt{\frac{2 \ln(8|\Theta|^2|B|^2|A|^2/\delta)}{K}} \leq \sqrt{\frac{4 \ln(8|\Theta|^2|B|^2|A|^2/\delta)}{n \cdot \mu_{\min}}}. \end{aligned}$$

Setting the right hand side = $\frac{1}{4|A|}$, we get $n \geq \frac{64|A|^2 \cdot \ln(8|\Theta|^2|B|^2|A|^2/\delta)}{\varepsilon^2 \cdot \mu_{\min}}$. And by the union bound, with probability $1 - \delta/4$ we have the left-hand side $\leq \frac{1}{4|A|}$ for all $\theta, \theta', b, b', a, a', \circ \in \{+, -\}$.

Therefore when $n \geq$

$$\begin{aligned} & \max \left\{ \frac{64|A|^2 \cdot \ln(8|\Theta|^2|B|^2|A|^2/\delta)}{\varepsilon^2 \cdot \mu_{\min}}, \frac{2 \ln(2|\Theta| \cdot |B|/\delta)}{\mu_{\min}^2} \right\} \\ &= \Theta \left(\ln(G/\delta) \cdot \max \left\{ \frac{|A|^2}{\varepsilon^2 \cdot \mu_{\min}}, \frac{1}{\mu_{\min}^2} \right\} \right), \end{aligned}$$

with probability $1 - \delta/2$, p^* is a feasible solution of (5.3), where $G = \max\{|\Theta|, |B|, |A|\}$. \square

Define $\text{Rev}_S(\mathcal{M})$ to be the expected revenue of \mathcal{M} when the underlying distribution is the uniform distribution over S . Let \mathcal{E} be the event that p^* is a feasible solution of the LP (5.3). Then the expected revenue of Mechanism 1 is equal to

$$\begin{aligned} & \mathbb{E}_S \text{Rev}_S(\mathcal{M}_S) \\ &= \mathbb{E}_S \mathbf{1}(\mathcal{E}) \text{Rev}_S(\mathcal{M}_S) + \mathbb{E}_S \mathbf{1}(\bar{\mathcal{E}}) \text{Rev}_S(\mathcal{M}_S) \\ &\geq \mathbb{E}_S \mathbf{1}(\mathcal{E}) \text{Rev}_S(p^*) \\ &\quad + \mathbb{E}_S \mathbf{1}(\bar{\mathcal{E}}) (\text{Rev}_S(p^*) + \text{Rev}_S(\mathcal{M}_S) - \text{Rev}_S(p^*)) \\ &= \mathbb{E}_S \text{Rev}_S(p^*) + \mathbb{E}_S \mathbf{1}(\bar{\mathcal{E}}) (\text{Rev}_S(\mathcal{M}_S) - \text{Rev}_S(p^*)) \\ &\geq \mathbb{E}_S \text{Rev}_S(p^*) - 2 \cdot \mathbb{E}_S \mathbf{1}(\bar{\mathcal{E}}) \\ &\geq \text{Rev}_{\mu}(p^*) - \delta. \end{aligned}$$