THÈSE de DOCTORAT

de l'UNIVERSITÉ TOULOUSE CAPITOLE

Présentée et soutenue par

Madame Thi Hien PHAM

Le 18 décembre 2023

Essai sur la théorie microéconomique

École doctorale : **TSE Toulouse Sciences Economiques**
Spécialité : **Sciences Economiques - Toulouse**
Unité de recherche : **TSE-R - Toulouse School of Economics - Recherche**

Thèse dirigée par *Monsieur Johannes HORNER et Monsieur Takuro YAMASHITA*

Composition du jury
Rapporteur : M. Peter ESö
Rapporteure : Mme Sarah AUSTER
Directeur de thèse : M. Johannes HORNER
Co-directeur de thèse : M. Takuro YAMASHITA
Président : M. Jacques CRéMER

Essays on Microeconomic Theory

Hien Pham

Abstract

This thesis comprises four essays on mechanism and information design. The first two chapters focus on the joint mechanism and information design by a monopolistic seller. In Chapter 1, a monopolistic seller jointly designs allocation rules and (new) information about a pay-off relevant state to a buyer with private types. When the new information flips the ranking of willingness to pay across types, a *screening* menu of prices and threshold disclosures is optimal. Conversely, when its impact is marginal, *bunching* via a single posted price and threshold disclosure is (approximately) optimal. While information design expands the scope for random mechanisms to outperform their deterministic counterparts, its presence leads to an equivalence result regarding sequential versus. static screening. These findings explain distinct strategies of pricing and information provision adopted across industries and stages of a product's lifecycle, rationalizing the prevalence of free information in many markets.

Chapter 2 investigates the interaction of buyer's optimism, information design, and price discrimination. The model features a buyer who holds a *biased* and *private* prior belief about the product's match quality. The seller can provide additional information about the product to refine the buyer's belief. We fully characterize the revenue-maximizing menu of price-information bundles that follows a simple cutoff structure. While neither the diversity in the (biased) priors nor information design suffices to trigger price discrimination, their combination induces the optimal mechanism featuring both information and price discrimination. Moreover, we show that it is optimal to offer information free of charge.

Chapter 3 (joint with Takuro Yamashita) considers an auction design problem with private values, where the seller and bidders may enjoy heterogeneous priors about their (possibly correlated) valuations. Each bidder forms an (interim) belief about the others based on his own prior, updated by observing his own value. If the seller faces uncertainty about the bidders' priors, even if he knows that the bidders' priors are within any given distance from his, he may

find it worst-case optimal to propose a dominant-strategy auction mechanism. This provides a foundation for dominant-strategy mechanisms in auctions with heterogeneous priors.

Chapter 4 (joint with Daniil Larionov, Takuro Yamashita, and Shuguang Zhu) studies mechanism design with flexible but costly information acquisition. There is a principal and four or more agents, sharing a common prior over the set of payoff-relevant states. The principal proposes a mechanism to the agents who can then acquire information about the state of the world by privately designing a signal device. As long as it is costless for each agent to acquire a signal that is independent of the state, there exists a mechanism which allows the principal to implement any social choice rule at zero information acquisition cost to the agents. Two applications are considered, including auctions with common value and collective decision-making.

Acknowledgements

I am greatly indebted my advisors, Johannes Hörner and Takuro Yamashita. They have taught me much more than this thesis shows. Long discussions in countless meetings with Takuro over the last four years have turned me from a student who knew nothing about mechanism and information design to being able to explore, develop, and answer my own questions. While Johannes only became my advisor in the final year of my Ph.D., I have learnt a lot from him. His sharp mind guided me to think logically, write precisely, and present effectively. It has been a privilege being their student. I am also grateful to Jacques Crémer, who has always been willing to read my drafts, even when they were unpolished; lift my spirits when I felt lost; and give huge encouragement for my small achievements.

I also thank faculty members and fellow graduate students at TSE, and in particular to Nour Meddahi and Christian Hellwig for their invaluable support; to Daniel Garrett, Niccolò Lomys, Thomas Mariotti, David Martimort, and Alex Smolin for helpful discussions; and to Mudit Dhakar, Dakang Huang, Chanwoo Kim, Cyril Thomson, and Philipp Wangner for making graduate school much more enjoyable. A special thanks to Andrew Rhodes, who offered valuable insights for my research and tough decisions, offering his ERC funding without my having to ask for it.

Finally, I am thankful to my family, especially my sister and my younger brother who have been counting the days until I come back home. I would also like to thank Nuong Nguyen for always being by my side.

This work receives financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreements No. 759733, No. 714693, and No. 101088307).

Contents

Chapter 1

How Information Design Shapes Optimal Selling Mechanisms

1 INTRODUCTION

The evolution of informational technology has significantly broadened sellers' ways of selling their products. They can design not only *allocation rules* which specify how to allocate products and charge payments to buyers, but also *information policies* which control how much buyers learn about the products, thereby refining their willingness to pay. For instance, they may offer a posted price, associated with full information, to everyone. Alternatively, they could propose a rich menu of allocation rules and information policies.

As an example, many software such as McAfee and various (mobile) apps like Spotify provide users with a *single* free trial version, followed by a *single* subscription fee schedule. The trial version is, therefore, *merely* a learning opportunity for potential buyers to make well-informed purchasing decisions. An opposite example is travel agency platforms such as Priceline and Hotwire practice so-called "opaque pricing" by which, buyers either book hotels with detailed information at standard prices or opt for limited details at discounted prices. Thus, these travel agencies *screen* their buyers via a menu of prices and information policies.

Price and information discrimination is also in the form of pre-order offers for buyers of not-yet-released products, as exemplified by Google's recent pre-order bonus for the Pixel 8. By contrast, well-known products are typically sold via a *single* posted price, coupled with a *single* timeframe for free return to all buyers.

What leads to these diverse selling strategies? In particular, when is a single posted price and disclosure policy optimal and conversely, when is it necessary to provide a screening menu of prices and information? In addition, is there any benefit from offering random mechanisms? Given that classical mechanism design results (Myerson (1981)) predict that a posted price is optimal when the informational environment is *fixed*, answering these questions explains how information design shapes optimal selling mechanisms. Regarding the timing, can the seller's revenue be improved by contracting with the buyer at the "*interim*" stage where he knows his type but before the seller's information disclosure? Or equivalently, should she allow the buyer to walk away at the "*posterior*" stage where he observes both his type and the information provided? Answering this question helps understand the impact of consumer protection regulations that grant the consumer a withdrawal right such as the European directive 2011/83/EU.[1] Finally, if the buyer privately observes the information disclosed by the seller, can the buyer enjoy any rent induced from such an *endogenously private* information?

This paper aims to answer these questions. The model, as formally described in Section 2, features a seller (she) who sells an object to a buyer (he) with a privately known initial valuation (*initial type*). The seller controls how much the buyer learns about an *additional component* in his valuation. For example, this additional component represents what the buyer learns via product trials. The seller designs a menu of *information policies* for different types of the buyer, and *allocation rules* for different types and signals. Therefore, she solves a joint mechanism and information design problem in which information plays a dual role. First, it allows the seller to screen the buyer's type through discriminatory disclosure policies. Second, disclosed information serves as input for designing allocation rules. We focus on the case where the buyer privately observes the new information (private signals) and investigate the case with public signals as a benchmark.

1.1 Summary of results

First, we establish a revenue-equivalence result regarding sequential vs. static screening. Specifically, we show that for any *feasible* and *deterministic* mechanism, there exists a mechanism that generates the same revenue for the seller and non-negative payoff for the buyer at any type and signal realization. As a consequence, there is no revenue loss if contracting at the posterior stage when the buyer knows both his type and signal. This result counters the well-established idea in sequential screening suggesting that the seller's revenue is strictly higher if contracting with the

[1] For a detailed discussion on such policies, see Krähmer and Strausz (2015*b*).

buyer before, rather than after, he learns additional information.[2] The basic intuition is that the seller's ability to flexibly design information can crowd out the advantages of sequential over static screening. A practical implication is that afore-mentioned consumer protections do not necessarily harm the seller, rationalizing the prevalence of free information in many markets.

Second, we investigate the (ir)relevance of signal privacy. In the benchmark problem with public signals, only *expected* allocations and payments (over signals) matter. Hence, this benchmark admits multiple solutions, including $\mathbf{M}^\star$, a *screening* menu of threshold disclosures $\pi^\star$ and prices paid conditional on trade.[3] We provide a simple way to verify the (ir)relevance of signal privacy, which is to check if, under $\mathbf{M}^\star$, the highest type pays the lowest price. If this is true, privacy of signals is irrelevant and $\mathbf{M}^\star$ solves the seller's original problem. We find that this is not always the case and consequently, not observing signals generally hurts the seller. Moreover, *per-signal* allocations and payments matter, which significantly complicates the characterization of optimal mechanisms In particular, it is not *a priori* clear how many signals are needed and which incentive compatibility (IC) constraints are relevant. The seller must also handle double deviations when the buyer lies about both his type and observed signal. Leveraging techniques for mechanisms with non-convex type spaces, we make it *always* possible for the buyer to "correct his lie," facilitating the characterization of optimal double deviations and thereby, optimal mechanisms.

Our main result characterizes optimal mechanisms, starting with binary types. The seller faces a trade-off between maximizing virtual surplus and minimizing the *posterior* rent. A threshold disclosure rule, under which signal realization is either "*good news*" if the state is above some cutoff or "*bad news*" otherwise, is optimal in both targets.[4] Under the optimal mechanism, the seller either *screens* the buyer's types (via a menu of threshold disclosures and posted prices) or *bunches* them (via a single posted price and threshold disclosure), depending on whether the threshold disclosure $\pi^\star$ induces a *threshold flip* of type order: the high type's value after "*bad news*" is *lower* than the low type's after "*good news*." Specifically, screening is optimal when this flip of type order occurs, and bunching otherwise.

To grasp the intuition, note that such a flip of type order occurs when the variation of valuations is mainly driven by the unknown component, leaving some room for the threshold disclosure $\pi^\star$ to reverse the ranking of valuation. Information (about the unknown component) matters,

[2] See Courty and Li (2000) and Krähmer and Strausz (2015*b*).

[3] See Definition 4 for a formal description of $\mathbf{M}^\star$.

[4] See Definition 1 for our formal definition of a threshold disclosure.

serving as a screening tool. Conversely, if the buyer's type is the main driver, which prevents $\pi^\star$ from flipping the type order, information is not crucial and screening disappears. The optimal mechanism echoes its counterpart in standard mechanism design where the buyer's valuation is his type: a posted price (but associated with threshold disclosure) is optimal.

The significance of this bunching vs. screening result is two-fold. First, it implies that in the above-mentioned scenarios, eliciting signals and random mechanisms are worthless. Second, it rationalizes observed mechanisms in practice. For coming-soon items, the unknown component's impact on the variation of valuations is large and a screening menu is employed. By contrast, its impact is marginal for well-known products where bunching comes into play. The significance of the unknown component also varies across different industries. In the realm of hotels, it matters much more than in software or mobile apps, leading to screening for the former and bunching for the latter.

Having characterized the optimal mechanism for the binary-type setting, we consider larger type spaces. With more than two types, there are also cases where an information policy reverses the ranking of valuations within a group of types but fails to do so for another. Consequently, not only information but also trading probabilities are needed to screen the buyer, leading to a *random* solution. However, the two scenarios of bunching/screening extend to the case with finitely many types, under stronger notions of flip (no flip) of type order. Specifically, a screening menu is optimal under a *partition* flip by $\pi^\star$ of type order - which generalizes the threshold flip of type order by $\pi^\star$, taking into account medium types and their associated cut-off states. Instead, bunching via a fixed price and threshold disclosure maximizes the seller's revenue when there is *uniformly* no threshold flip of type order under which, the type order is to be preserved between *any* pair of types and after *any* threshold disclosure. This strong requirement of type order preservation helps deal with the challenge of determining the lowest type being served in a rich type space.

As binding (IC) constraints can involve local, global, and upward ones, characterizing optimal random mechanisms becomes difficult. We thus focus on shedding light on how random mechanisms outperform their deterministic counterparts.[5] We first establish the "*no randomization at the top*" result, extending the well-known "no distortion at the top" to a setting with information design: the highest type receives an efficient (and hence, deterministic) allocation. In turn, this implies an optimal contract for this type, featuring a posted price and no disclosure. While

[5] In the Online Appendix, we solve for the optimal random mechanism in several examples.

randomization is not needed for the highest type, it can be helpful for the lower types, leading to a better balance of the efficiency vs. rent trade-off.[6] We analyze, by examples, how random mechanisms facilitate screening distant types as well as screening signals.

Finally, we consider a setting with a continuum of types. In this case, the optimality of a screening menu of posted prices and threshold disclosures under a partition flip of type order extends readily. Particularly, in a "continuous" model when valuation shifts smoothly across types and states, this notion corresponds to the ranking of valuations at the zero-virtual-value states by types being reversed. On the other hand, the fact that there are always types whose valuations are close to others' makes it impossible to flip the ranking of willingness to pay across *all* types. We show that when the type order is *almost* preserved, bunching via a fixed price-threshold disclosure bundle is *approximately* optimal. If there is only two states, we establish the "exact" optimality of bunching within the class of deterministic mechanisms..

1.2 Related literature

We contribute to the literature on joint mechanism and information design, comprising two main strands. The first, more related, strand endows the buyer with a private type, initiated by Eső and Szentes (2007) who focus on full disclosure. Most other papers focus on *posted-price* mechanisms,[7] which in turn, makes it without loss of generality to focus on *binary-signal* information structures (Li and Shi (2017), Guo et al. (2022), Wei and Green (2023), Smolin (2023)).[8] Our findings imply that these restrictions are not innocuous in general.

Our model builds on Eső and Szentes (2007) who focus on full disclosure and an environment with (i) the above-mentioned "continuous" model and (ii) certain assumptions on the valuation function. Under such an environment, they show that the upper bound of revenue with public signals can be achieved via full disclosure, associated with a screening menu of prices (for the good) and information fees. However, their optimal mechanism is not incentive compatible

[6]While it is natural to expect the two-dimensionality feature of the buyer's valuation to lead to random mechanisms, the seller has another tool for randomization: the distribution of signals, which potentially makes random mechanisms redundant. However, signal misreporting off-path shuts off this additional instrument. Thus, random mechanisms arise to deter double deviations, minimizing the *posterior* rent.

[7]In posted-price mechanisms, each type receives a posted price for the good and in some cases, a posted fee for information.

[8]Exceptions include Zhu (2023) and Krähmer (2020) who establish full surplus extraction results when the seller can correlate information disclosed to multiple buyers, and when randomizing over information structures is allowed and the buyer's type correlates with the unknown component, respectively.

and moreover, privacy of signals generally matters outside their environment.[9] Not only do we allow for general information structures, we also characterize a *joint* design of information and allocation rules in a more general environment of type space and valuation functions. This allows us to uncover how information design reshapes the optimal selling mechanism which features not just screening, but also bunching and a random mechanism. At the same time, we strengthen Eső and Szentes (2007)'s finding by showing that the irrelevance of signals extends to other (but not all) environments, with appropriate information design.

Bergemann and Wambach (2015) and Wei and Green (2023) revisit Eső and Szentes (2007)'s continuous model, showing that the latter's *optimal* allocation can be implemented under stronger participation constraints. We show that with deterministic allocations (including Eső and Szentes (2007)'s), this is true for any *feasible* allocations, not just optimal. In turn, this provides an alternative proof for Wei and Green (2023).[10]

In the second, less related, strand of this literature, the buyer's valuation is the unknown component itself. See, for example, Lewis and Sappington (1994), Bergemann and Pesendorfer (2007), Bergemann et al. (2022). Without the buyer's private types, information cannot serve as a screening tool. Moreover, the buyer's private information (about his valuation) arrives only once, making the seller's problem static.[11]

We also contribute to the literature on dynamic mechanism design in which handling off-path misreporting is a notable issue. Eső and Szentes (2007) explicitly characterize an agent's optimal double deviation, which is to "correct the lie". However, such a lie correction is feasible only if the agent's payoff shares a *common support* across types, which is rather restrictive. We show that by leveraging mechanism design techniques for a non-convex type space, lie correction is feasible even with non-common supports. Moreover, the existing literature (for instance, Battaglini (2005), Eső and Szentes (2007), Pavan et al. (2014)) extensively relies on the first-order approach considering only local incentive compatibility constraints.[12] Instead, we characterize different scenarios of binding constraints, showing that global deviations (associated with

[9] See Krähmer and Strausz (2015*a*) for a detailed discussion

[10] Wei and Green (2023) also shows that information disclosure triggers reverse price discrimination. We show that this can also be derived from the properties of Eső and Szentes (2007)'s optimal mechanism.

[11] If the buyer in our model has no private type, the seller fully extracts the surplus by offering no disclosure and a posted price for the good, which is equal to the expected valuation.

[12] The validity of this approach usually requires certain regularity conditions, which are not easy to satisfy, see Battaglini and Lamba (2019).

double deviation off-path) lead to bunching and random solutions.[13]

Finally, we contribute to the recent literature on Bayesian persuasion following Kamenica and Gentzkow (2011), where a sender designs *only* information disclosure to affect a receiver's action. When the latter has a private type, Kolotilin et al. (2017) show that with binary actions and linear valuation functions, non-discriminatory disclosure is optimal. In our *joint* design problem, the buyer's action space (which is the menu of allocations and payments) is endogenous and can consist of more than two options. We show that the optimality of non-discriminatory disclosure, while not being true in general, holds in some environments even if the seller also designs allocation rules and the valuation function is non-linear.

2 MODEL

2.1 Environment

The principal, a seller (she) sells an object to an agent, the buyer (he). The buyer's valuation for the object, $v(\theta, x) \in \mathbb{R}_+$, depends on two components: (i) the buyer's type $\theta \in \Theta \subset \mathbb{R}$ and (ii) an unknown state $x \in X \subset \mathbb{R}$. There are a finite number of possible types and states, *i.e.*, $|\Theta| < \infty$ and $|X| < \infty$.[14] Random variables θ and x are independent. Let $f(\theta)$ be the probability of each type θ and $\mu(x)$ of each state x. Without loss of generality, assume $f(\theta) > 0$ and $\mu(x) > 0$ for all θ and x.

The realization of $\theta \in \Theta$ is privately known by the buyer. Neither the seller nor the buyer knows the state $x \in X$. The seller commits to a policy of information disclosure about the state, formally defined in Section 2.2.

To define payoffs, let $q \in [0,1]$ be the trading probability and $p \in \mathbb{R}$ the expected transfer from the buyer to the seller. The seller's *ex post* payoff is then p and the buyer's is $v(\theta, x)q - p$.

For expositional clarity, we use the following notations $\overline{\theta} \equiv \max \Theta$, $\underline{\theta} \equiv \min \Theta$, $\theta^+ \equiv \min_{\theta' \in \Theta}\{\theta' \mid \theta' > \theta\}$ for $\theta < \overline{\theta}$, $\theta^- \equiv \max_{\theta' \in \Theta}\{\theta' \mid \theta' < \theta\}$ for $\theta > \underline{\theta}$, and $\overline{\theta}^+ = \overline{\theta}$, $\underline{\theta}^- = \underline{\theta}$. We define $\overline{x}$ and $\underline{x}$ similarly. Let

$$\phi(\theta, x) \equiv v(\theta, x) - [v(\theta^+, x) - v(\theta, x)] \frac{\sum_{\theta' > \theta} f(\theta)}{f(\theta)}$$

[13] Even with full disclosure, which makes our problem become a standard dynamic screening problem, random mechanisms can outperform their deterministic counterparts. See Example 3(b).

[14] We study the infinite type and state spaces in Section 7.

denote the buyer's virtual value. Throughout, assume that both the valuation and virtual valuation increase in the buyer's type and the state.

Assumption 1 (Monotone value). *$v(\theta, x)$ increases in θ and x.*

Assumption 2 (Monotone virtual value). *$\phi(\theta, x)$ increases in θ and x.*

2.2 Selling mechanism

The seller designs, and *ex ante* commits to a *grand* mechanism or a menu of (i) information policies for different types of the buyer and (ii) allocation rules for different types and information received by the buyer.

Information policies: We model information policies as information structures (experiments) $\Pi \equiv (S, \pi)$, which consists of a countable set of signals $S \subset \mathbb{R}$,[15] and a mapping π, which associates to each state θ a distribution over signals $\pi(\cdot \mid x) \in \Delta(S)$. Given a mapping π and a signal realization $s \in S$, the corresponding posterior belief $\Psi(\cdot|s) \in \Delta(X)$ is obtained by Bayes' rule whenever possible, and is given by

$$\mu_{s,\pi}(x) = \frac{\mu(x)\pi(s \mid x)}{\sum_{x' \in x} g\left(x'\right) \pi\left(s \mid x'\right)}$$

An example of information structures is the threshold rule, defined as follows.

Definition 1 (Threshold diclosure). *If the information policy follows a threshold rule, each signal realization is classified as either "good news" or "bad news". Moreover,*

$$\pi(\text{“}good\ news\text{"}, x) = \begin{cases} 1 & if \quad x > \hat{x}, \\ \lambda & if \quad x < \hat{x}, \\ \lambda & if \quad x = \hat{x}, \end{cases} \quad \text{for some } \hat{x} \in X \text{ and } \lambda \in [0,1].$$

Thus, a threshold disclosure is represented by a pair $(\hat{x}, \lambda)$ where $\hat{x}$ is the cut-off state and λ the probability with which "*good news*" is sent at the cut-off state. It informs the buyer whether the state is (weakly) higher or lower than $\hat{x}$. To simplify notations, throughout the paper, we use "s^g" to represent "*good news*" and "s^b" for "*bad news*".

A menu of experiments is a set $\{\pi_\theta\}_{\theta \in \Theta}$. The paper focuses on the case in which the buyer privately observes the signal. The benchmark case with public signals is examined in Section 3.3.

[15] Assuming S is a countable set of $\mathbb{R}$ is without loss.

Without loss of generality, assume that signals are ordered such that upon observing a higher signal, the buyer's posterior valuation is higher, as follows.

Assumption 3 (Ranking of signals)**.**

$$s > s' \Leftrightarrow \sum_x v(\theta, x)\mu_{s,\pi_\theta}(x) \geq \sum_x v(\theta, x)\mu_{s',\pi_\theta}(x)$$

Allocation rules: An allocation rule specifies the trading probability, q, and the expected transfer from the buyer to the seller, p. Given the information structure, by the revelation principle (see, for example, Myerson (1986*b*)), we focus on direct allocation rules $\{q(\theta, s), p(\theta, s)\}_{\theta,s}$.

Thus, a selling mechanism is a tuple $\mathbf{M} \equiv \left\{\pi_\theta, \left(q(\theta, s), p(\theta, s)\right)\right\}_{\theta,s}$. The formal definitions of a deterministic mechanism and its random counterpart are as follows.

Definition 2. *An mechanism* $\mathbf{M}$ *is deterministic if under* $\mathbf{M}$*,* $q(\theta, s) \in \{0, 1\}$ *for all* $\theta \in \Theta$ *and* $s \in S$*.* $\mathbf{M}$ *is random otherwise.*

Timing: The timing of interactions is as follows:

1. The seller offers a selling mechanism **M**.
2. The buyer learns his type θ and decides to accept or reject the offer. In case of acceptance, he reports a type $\hat{\theta}$ to receive information generated from $\pi_{\hat{\theta}}$.
3. The buyer *privately* observes a signal s and reports a signal $\hat{s}$.
4. The allocation $(q(\hat{\theta}, \hat{s}), p(\hat{\theta}, \hat{s}))$ is implemented.

According to this timing, the buyer's participation is decided at the *interim* state, as commonly assumed in the mechanism design literature. See our discussion on the timing structure in Section 3.2.

2.3 Seller's problem

An optimal mechanism refers to a revenue-maximizing mechanism. By the revelation principle, it is without loss of generality to focus on direct mechanisms such that the buyer finds it optimal to (i) participate in the mechanism, (ii) truthfully report his type, and (iii) truthfully report his signal conditional on being truthful about his type. Let

$$u(\theta, \theta', s, s') \equiv \sum_x [v(\theta, x)q(\theta', s') - p(\theta', s')]\Psi_\theta(x|s)$$

denote the *ex post* payoff for type-θ buyer, who reports θ', observes s, and reports s'. Note that if the buyer lies about his type, he may want to lie again about the signal. In other words, double deviations from truth-telling may be attractive. Let

$$s^{\star}(\theta,\theta',s) \in \underset{s'}{\operatorname{argmax}}\, u(\theta,\theta',s,s')$$

be the optimal signal reporting of type-θ buyer who reports θ' and observes signal s.[16] The *ex ante* payoff for type-θ buyer, who reports θ' and then $s^{\star}(\theta,\theta',s)$, is then given by

$$U(\theta,\theta') \equiv \sum_x \sum_s u(\theta,\theta',s,s^{\star}(\theta,\theta',s))\pi(s|x).$$

With abuse of notation, let $u(\theta,s) \equiv u(\theta,\theta,s,s)$, $u(\theta,s,s') \equiv u(\theta,\theta,s,s')$, and $U(\theta) \equiv U(\theta,\theta)$. For the buyer to truthfully report his signal on the equilibrium path (conditional on reporting his type truthfully), it must be that for all θ and s,

$$u(\theta,s) \geq u(\theta,s,s'). \tag{IC-signal}$$

For the buyer to truthfully report his type, it must be that for all θ and θ',

$$U(\theta) \geq U(\theta,\theta'). \tag{IC-type}$$

Finally, the buyer participates in the mechanism if and only if

$$U(\theta) \geq 0. \tag{IR}$$

Definition 3. *A mechanism is feasible if it satisfies all constraints (IR), (IC-type), and (IC-signal),*

Formally, the seller's maximization problem is given by

$$\sup_{\left\{\pi_\theta,q(\theta,s),p(\theta,s)\right\}_{s,\theta}} \sum_\theta \sum_x \sum_s p(\theta,s)\pi(s|x)\mu(x)f(\theta)$$
$$s.t. \quad \text{(IR), (IC-type), (IC-signal).}$$

[16] In case the buyer is indifferent between signals off the equilibrium path, fix arbitrarily one of the seller-preferred signals.

3.1 No distortion at the top and no rent at the bottom

First, we establish that the solution to the seller's joint design problem bears commonly known features: the highest type receives an efficient allocation while the lowest is fully extracted.

Lemma 1. *Under any optimal mechanism,*

(a) the lowest type gets a zero payoff: $U(\underline{\theta}) = 0$*, and*

(b) the highest type an efficient allocation: $q(\overline{\theta}, x) = \begin{cases} 1 & \text{if} \quad v(\overline{\theta}, x) > 0, \\ \in [0,1] & \text{if} \quad v(\overline{\theta}, x) = 0. \end{cases}$.

To prove Part (a) or the "no rent at the bottom" feature, we first show that the buyer's rent $U(\theta)$ increases in θ under any incentive-compatible mechanism. For this purpose, consider the buyer of type $\theta > \underline{\theta}$. He reveals his type only if his payoff from truth-telling is at least that obtained by mimicking some type $\theta' < \theta$ and reporting signals truthfully. Formally,

$$U(\theta) \geq \sum_x \left[v(\theta, x) q(\theta', s) - p(\theta', s) \right] \pi_{\theta'}(s|x) g(x) \tag{1.1}$$

As $v(\theta, x) \geq v(\theta', x)$ for all x, (1.1) implies $U(\theta) \geq \sum_x [v(\theta', x) q(\theta', s) - p(\theta', s)] \pi_{\theta'}(s|x) g(x) = U(\theta')$. Therefore, $U(\cdot)$ is an increasing function. Armed with this result, we now show that $U(\underline{\theta}) = 0$ under optimal mechanisms. By contradiction, suppose $U(\underline{\theta}) = \varepsilon > 0$ under an optimal mechanism. Then, if increasing $p(\theta, s)$ by ε for all θ and s, the seller strictly increases her revenue while not violating any IC and IR conditions. A contradiction. Thus, $U(\underline{\theta}) = 0$ at optimum.

We leave the proof of Part (b) or the "no distortion at the bottom" result in Appendix A.1. The idea is that whenever this type does not trade with probability 1 (at some state), it is possible to improve the seller's revenue by letting him always trade under no disclosure and a posted price being equal to his original expected payment, adding the new surplus.

It is worth noting that by Lemma 1(b), random allocations are not needed for the highest type. This is not necessarily true for the lower types to which, offering efficient allocations is generally sub-optimal. See Section 6.1 for a detailed discussion.

3.2 Sequential vs. static screening

This section establishes an irrelevance result regarding the timing structure of interactions. We show that contracting at the *posterior* stage (after the buyer observes both his type and signal) does not necessarily hurt the seller. Specifically, within the class of deterministic mechanisms, there is no revenue loss if the buyer can walk away after information disclosure.

Proposition 1. *For any deterministic and feasible mechanism, there exists a mechanism which generates the same revenue for the seller and a non-negative ex post pay-off for the buyer.*

Let $\mathbf{M}^d \equiv \{q(\theta,s), p(\theta,s), \pi_\theta\}$ be an arbitrary deterministic and feasible mechanism. Hence, under $\mathbf{M}^d$, $q(\theta,s) \in \{0,1\}$ for any θ and s. Fix $\theta \in \Theta$. Let $S_\theta^g \equiv \{s \mid q(\theta,s)=1\}$ and $S_\theta^b \equiv \{s \mid q(\theta,s)=0\}$. To induce signal truth-telling by θ, $p(\theta,s) = p(\theta,s') = \underline{p}(\theta)$ if $s \in S_\theta^g$; and $p(\theta,s) = p(\theta,s') \equiv \overline{p}(\theta)$ and if $s \in S_\theta^b$. Let

$$\mathcal{Q}(\theta) \equiv \sum_x \sum_{s \in S_\theta^g} \pi_\theta(s|x) g(x)$$

represent type θ's trade probability. Consider the following two cases:

Case 1: $\underline{p}(\theta) < 0$. If $\exists s$ such that $\pi(\theta,s) < 0$, then type θ who observes s (strictly) prefers to misreport signal s' with $q(\theta,s')=0$ to receive a negative transfer without buying the good. This contradicts with $\mathbf{M}^d$ being feasible. Hence, $\pi(\theta,s) \geq 0$ for all s in this case.

Case 2: $\overline{p}(\theta) \geq 0$. Revise $\mathbf{M}^d$ as follows. For any $s \in S_\theta^g$, replace s with signal "s^g" ; and for any $s \in S_\theta^b$, replace it with signal "s^b". In addition, each type θ now receives a posted price

$$\tilde{p}(\theta) = \overline{p}(\theta) + \underline{p}(\theta)\frac{1-\mathcal{Q}(\theta)}{\mathcal{Q}(\theta)}, \tag{1.2}$$

As θ pays only if he decides to buy the good, his payoff is non-negative at any signal realization. In Appendix A.2 we show that the seller's revenue remains unchanged under this revision, which completes the proof for Proposition 1.

Proposition 1 has two implications. First, if only deterministic mechanisms are allowed, there is no loss for the seller to contract after the buyer observes both type and signal realizations. Therefore, despite the sequential arrival of his private information, sequential screening the buyer is not beneficial, unless random mechanisms are necessary.

Second, by the "no rent at the bottom", the lowest type earns a zero payoff under optimal mechanisms. Therefore, if there exists θ such that $\underline{p}(\theta) < 0$, the lowest type mimics θ and not buying the good to enjoy a positive payoff. Consequently, if $\mathbf{M}^d$ is optimal, $\underline{p}(\theta) \geq 0 \ \forall\theta$. Then, as argued above, $\mathbf{M}^d$ do no better than "posted-price" mechanisms, which are *signal-independent.* As we will show, such posted-price mechanisms are optimal in several, but not all, environments.

3.3 (Ir)relevance of signal privacy

As the buyer privately observes signals only after the contract is signed, one might expect that the privacy of signals does not hurt the seller's revenue. To investigate this conjecture, we first consider the benchmark problem with public signals. There, the buyer's payoff and the seller's revenue depends only on *expected* payments and *expected* allocations over signals, defined as

$$\mathbb{Q}(\theta,x) \equiv \sum_x \sum_s q(\theta,s)\pi_\theta(s|x)\mu(x), \quad \mathbb{P}(\theta) \equiv \sum_x \sum_s p(\theta,s)\pi_\theta(s|x)\mu(x).$$

As a result, the seller's problem reduces to

$$(\overline{\mathscr{P}}) \quad \sup_{\mathbb{Q},\mathbb{P}} \sum_\theta \mathbb{P}(\theta) f(\theta)$$

$$s.t. \quad \sum_x v(\theta,x)\mathbb{Q}(\theta,x)\mu(x) - \mathbb{P}(\theta) \geq \sum_x v(\theta,x)\mathbb{Q}(\theta',x)\mu(x) - \mathbb{P}(\theta'), \qquad (\overline{IC}\text{-type})$$

$$\sum_x v(\theta,x)\mathbb{Q}(\theta,x)\mu(x) - \mathbb{P}(\theta) \geq 0. \qquad (\overline{IR})$$

Under Assumption 1 and 2, only local IC constraints bind under $(\overline{\mathscr{P}})$. By standard arguments (omitted), this problem reduces to point-wise maximization w.r.t $\mathbb{Q}$ only:

$$\sup_{\mathbb{Q}} \sum_\theta \sum_x \phi(\theta,x)\mathbb{Q}(\theta,x)\mu(x) f(\theta). \qquad (\star)$$

A solution to $(\star)$ exists and is generically unique:[17] $\mathbb{Q}(\theta,x) = \mathbb{1}_{\phi(\theta,x)\geq 0}$. *Expected* payment (over signals) is pinned down by $(\overline{IC}_{\theta^+\to\theta})$ and $(\overline{IR}_{\underline{\theta}})$. For any $\theta \geq \underline{\theta}$, let

$$x_\theta \equiv \begin{cases} \min\{x \mid \phi(\theta,x) \geq 0\} & \text{if} \quad \phi(\theta,\overline{x}) \geq 0, \\ +\infty & \text{if} \quad \phi(\theta,\overline{x}) < 0 \end{cases} \qquad (1.3)$$

denote the lowest state at which type θ's virtual value is non-negative. Note that $x^\star_\theta$ decreases in θ by Assumption 2.

Lemma 2 (Benchmark problem). *With public signals, the optimal mechanism is generically unique, given by*

$$\mathbb{Q}(\theta,x) = \mathbb{1}_{\phi(\theta,x)\geq 0}, \qquad (1.4)$$

$$\mathbb{P}(\theta^+) = \mathbb{P}(\theta) + \sum_{x_{\theta^+} \leq x < x_\theta} v(\theta^+,x)\mu(x) \quad \forall \theta \geq \underline{\theta}. \qquad (1.5)$$

[17] When $\phi(\theta,x) = 0$, any $\mathbb{Q}(\theta,x) \in [0,1]$ is optimal.

The seller retains a certain level of freedom in designing disclosure and *per-signal* allocation rules as long as (i) upon observing any signal, one knows whether the state is above or below the cut-off x_θ and (ii) expected terms are given by equations in Lemma 2. This leads to a multiplicity of solutions to $(\overline{\mathscr{P}})$, including the following menu of threshold disclosures and prices (paid conditional on trade), under which each type of the buyer (i) knows whether his virtual value is positive or not and (ii) pays only if trade happens. Formally:

Definition 4. *Under* $\boldsymbol{M}^\star \equiv \{p^\star(\theta,s), q^\star(\theta,s), \pi^\star_\theta\}_{\theta\in\Theta, s\in\{s^g,s^b\}}$ *is a menu of threshold disclosures and prices, in which*

1. $\pi^\star_\theta(s^g|x) = \mathbb{1}_{x\geq x_\theta}$, *where* x_θ *is given by equation (1.3).*

2. $\big(q^\star,(\theta,s), p^\star(\theta)\big) = \begin{cases} (1, \dfrac{\mathbb{P}(\theta)}{\sum_{x\geq x_\theta}\mu(x)}) & \text{if} \quad s = s^g, \\ (0,0) & \text{if} \quad s = s^b, \end{cases}$ *where* $\mathbb{P}(\theta)$ *is given by equation (1.5).*

Let $V(\overline{\mathscr{P}})$ represent the value of problem $(\overline{\mathscr{P}})$. Then, $V(\overline{\mathscr{P}})$ is an upper bound on the seller's revenue with private signals. Under a mild condition, Proposition 2(a) below shows that if this upper bound is achieved via some mechanism, it is via $\mathbf{M}^\star$. The basic intuition is that relative to other solutions to $(\overline{\mathscr{P}})$, $\mathbf{M}^\star$ provides less information (just enough to know the sign of virtual values) and a higher price for the good (payments are paid only when trade happens). Hence, if there exists a solution that induces truth-telling with private signals, so does $\mathbf{M}^\star$. This is the case, by Proposition 2(b), if and only if the highest type pays the lowest price under $\mathbf{M}^\star$.

To formally state Proposition 2, let $R_\mathbf{M}$ represent the revenue level obtained with private signals from an arbitrary mechanism $\mathbf{M}$.

Proposition 2.

a) Suppose $\phi(\theta, x_\theta) > 0 \ \forall\theta$. *If there exists* $\boldsymbol{M}$ *such that* $R_{\boldsymbol{M}} = V(\overline{\mathscr{P}})$, *then* $R_{\boldsymbol{M}^\star} = V(\overline{\mathscr{P}})$.

b) $R_{\boldsymbol{M}^\star} = V(\overline{\mathscr{P}})$ *if and only if* $p^\star(\overline{\theta}, s^g) = \min_\theta\{p^\star(\theta,s^g)\}$.

It seems counter-intuitive that the highest type pays the lowest price (conditional on buying the good). However, it is worth noting that information disclosure can flip the ranking of (*posterior*) willingness to pay across types, leading to non-monotone price discrimination.[18] As will be

[18]That information disclosure can lead to non-monotone price discrimination has been observed in Bang and Kim (2013) and Wei and Green (2023) where prices decrease in types. Throughout our paper, several examples are presented where under $\mathbf{M}^\star$, prices can be decreasing, increasing and even concave in types (see Example 6).

shown formally in later sections, this occurs in some, but not all environments.

4 A RESTATEMENT OF THE SELLER'S PROBLEM

Without loss of generality, assume that each signal induces a single (on-path) posterior valuation. Therefore, each signal s observed by type-θ buyer corresponds to his on-path posterior value after observing such a signal, given by

$$\omega^{\pi_\theta}(\theta, s) \equiv \sum_x v(\theta, x)\mu_{s,\pi_\theta}(x)$$

Moreover, that the buyer reveals the realized signal is equivalent to him reporting his posterior valuation. For any type θ, let

$$\Omega_\theta \equiv \{\omega \mid \omega = \omega^{\pi_\theta}(\theta, s) \text{ for some } s \in S\}$$

be the set of all possible on-path posterior values for type θ. Then, requiring signal truth-telling on-path is equivalent to ensuring truth-telling about on-path posterior values, or

$$\omega q(\theta, \omega) - p(\theta, \omega) \geq \omega q(\theta, \omega') - p(\theta, \omega') \quad \forall \theta, \forall \omega, \omega' \in \Omega_\theta$$

As mentioned, the buyer may want to coordinate lies about the realized type and signal. Given that the signal space is endogenous, this significantly complicates the characterization of truth-telling conditions. To facilitate characterizing the buyer's optimal double deviation, we extend the allocation rule to be defined on the set of all possible on-path and off-path posterior valuations, denoted by

$$\Omega \equiv [v(\underline{\theta}, \underline{x}), v(\overline{\theta}, \overline{x})].$$

Moreover, it is without loss of generality to require truthful signal reporting on this set Ω, rather than in only $\{\Omega_\theta\}_\theta$,[19] *i.e.*,

$$\omega q(\theta, \omega) - p(\theta, \omega) \geq \omega q(\theta, \omega') - p(\theta, \omega') \quad \forall \theta, \forall \omega, \omega' \in \Omega \qquad \text{(IC-value)}$$

The characterization of (IC-value) is standard.

Lemma 3 (Myerson, 1981)**.** *An allocation rule* $(q, p) : \Theta \times \Omega \to [0, 1] \times \mathbb{R}$ *satisfies (IC-value) if and only if*

[19] See, for example, Skreta (2006), for mechanism design with non-convex type spaces.

1. $\omega q(\theta,\omega) - p(\theta,\omega) = \hat{\omega} q(\theta,\hat{\omega}) - p(\theta,\hat{\omega}) + \int_{\hat{\omega}}^{\omega} q(\theta,z)dz,$
2. $q(\theta,\omega)$ *increases in* ω.

It then follows from Lemma 3 that the buyer, after having lied about his type, reveals his true (off-path) posterior valuation.

Lemma 4 (Optimal double deviations)**.** *Under any allocation rule* $(q,p):\Theta\times\Omega\to[0,1]\times\mathbb{R}$ *that satisfies (IC-value), it is optimal for type* θ *who mimics* θ' *and observe signal* s *to report his off-path posterior valuation, given by*

$$\omega^{\pi_\theta}(\theta',s) \equiv \sum_x v(\theta,x)\mu_{s,\pi_{\theta'}}(x)$$

The proof (omitted) is similar to what is called "correcting the lie" in the dynamic mechanism design literature. Often, this lie correction is made feasible by assuming that the agent's (new) private information shares a common support across types.[20] This is not applicable in our model as the buyer's new private information, which is his posterior valuation, is endogenous. By extending the allocation rule to be defined in the extended signal space Ω, we make it possible for the buyer to "correct his lie."[21]

Consider $\theta,\theta'\in\Theta$ with $\theta>\theta'$. Then,

$$\begin{aligned}
U(\theta,\theta') \equiv& \sum_x\sum_s [\omega^{\pi_{\theta'}}(\theta,s)q(\theta',\omega^{\pi_{\theta'}}(\theta,s)) - p(\theta',\omega^{\pi_{\theta'}}(\theta,s))]\pi_{\theta'}(s|x)\mu(x)\\
=&\sum_x\sum_s\left[[\omega^{\pi_{\theta'}}(\theta',s)q(\theta',\omega^{\pi_{\theta'}}(\theta',s)) - p(\theta',\omega^{\pi_{\theta'}}(\theta',s))] + \sum_s\int_{\omega^{\pi_{\theta'}}(\theta',s)}^{\omega^{\pi_{\theta'}}(\theta,s)} q(\theta',z)dz\right]\pi_{\theta'}(s|x)\mu(x)\\
=&U(\theta') + \sum_x\sum_s\int_{\omega^{\pi_{\theta'}}(\theta',s)}^{\omega^{\pi_{\theta'}}(\theta,s)} q(\theta',z)dz\pi_{\theta'}(s|x)\mu(x).
\end{aligned}$$

Thus, θ does not benefit from misreporting θ' if and only if

$$U(\theta)-U(\theta') \geq \sum_x\sum_s\int_{\omega^{\pi_{\theta'}}(\theta',s)}^{\omega^{\pi_{\theta'}}(\theta,s)} q(\theta',z)dz\pi_{\theta'}(s|x)\mu(x).$$

By similar arguments, θ' does not benefit from misreporting θ if and only if

$$U(\theta)-U(\theta') \leq \sum_x\sum_s\int_{\omega^{\pi_{\theta}}(\theta',s)}^{\omega^{\pi_{\theta}}(\theta,s)} q(\theta,z)dz\pi_{\theta}(s|x)\mu(x).$$

[20] See Eső and Szentes (2007) and Krähmer and Strausz (2015*b*) for example.

[21] This trick can also be helpful in other dynamic mechanism design problems where the agent(s)' private information does not share common support across types.

To sum up, the seller's problem can be expressed as follows.

$$(\mathscr{P}) \quad \max_{(\pi,q,U)} \quad \sum_{\theta} f(\theta)\Big[\sum_{x}\sum_{s} v(\theta,x)q(\theta,\omega^{\pi_\theta}(\theta,s))\pi_\theta(s|x)\mu(x) - U(\theta)\Big]$$

$$s.t: \quad \forall\theta, \quad U(\theta)-U(\theta') \geq \sum_{x}\sum_{s}\int_{\omega^{\pi_{\theta'}}(\theta',s)}^{\omega^{\pi_{\theta'}}(\theta,s)} q(\theta',z)dz\pi_{\theta'}(s|x)\mu(x) \quad \forall\theta'<\theta \qquad \text{(dwIC-type)}$$

$$U(\theta)-U(\theta') \leq \sum_{x}\sum_{s}\int_{\omega^{\pi_{\theta}}(\theta,s)}^{\omega^{\pi_{\theta}}(\theta',s)} q(\theta,z)dz\pi_{\theta}(s|x)\mu(x) \quad \forall\theta'>\theta \qquad \text{(uwIC-type)}$$

$$U(\theta) \geq 0 \qquad (IR)$$

$$q(\theta,\omega) \text{ increases in } \omega. \qquad (MON)$$

5 OPTIMAL MECHANISM FOR $|\Theta| = 2$

In this section, we characterize the optimal mechanism for binary types. We derive two findings. First, screening is optimal if and only if the ranking of willingness to pay is flipped under a certain threshold disclosure and bunching is optimal otherwise. Second, eliciting signals and random mechanisms are worthless. Formally, $\Theta = \{h, l\}$ and hence, the seller's problem reduces to $(\mathscr{P}_b)$, given by

$$(\mathscr{P}_b) \quad \max_{(\pi,q,U)} \quad \sum_{\theta} f(\theta)\Big[\sum_{x}\sum_{s} v(\theta,x)q(\theta,\omega^{\pi_\theta}(\theta,s))\pi_\theta(s|x)\mu(x) - U(\theta)\Big]$$

$$s.t: \quad U(h)-U(l) \geq \sum_{x}\sum_{s}\int_{\omega^{\pi_l}(l,s)}^{\omega^{\pi_l}(h,s)} q(l,z)dz\pi_l(s|x)\mu(x) \qquad (IC_{hl})$$

$$U(h)-U(l) \leq \sum_{x}\sum_{s}\int_{\omega^{\pi_h}(l,s)}^{\omega^{\pi_h}(h,s)} q(h,z)dz\pi_h(s|x)\mu(x) \qquad (IC_{lh})$$

$$U(h) \geq 0 \qquad (IR_h)$$

$$U(h) \geq 0 \qquad (IR_l)$$

$$q(\theta,\omega) \text{ increases in } \omega.$$

To state the main result of this section, we introduce the following notion of type order flip, which shapes the optimal mechanism. Recall that $\pi^\star$ is the threshold disclosure associated with $\mathbf{M}^\star$ formally defined in Definition 4, with $\pi_l^\star(s^g|x) = \mathbb{1}_{x\geq x_l}$.

Definition 5 (Threshold flip of type order by $\pi_l^\star$)**.**

If $\pi_l^\star$ induces the threshold flip of type order, $\mathbb{E}[v(h,x) \mid x < x_l] \leq \mathbb{E}[v(l,x) \mid x \geq x_l]$.

By Definition 5, $\pi_l^\star$ induces the threshold flip of type order when $\omega^{\pi_l^\star}(h, s^b) \leq \omega^{\pi_l^\star}(h, s^g)$. In words, this threshold disclosure overturns the ranking of willingness to pay with h's valuation after "*bad news*" being *lower* than l 's after "*good news*". Intuitively, this is the case when the unknown component x causes significant variations of valuations, creating room for $\pi_l^\star$ to flip the type order. By contrast, it does not happen in, for example, an extreme case in which valuation is constant with respect to this component (i.e., $g(\cdot)$ is a degenerate distribution with $\overline{x} = \underline{x}$), as in standard mechanism design problems.

We are now ready to state the main result of this section, assuming that type l 's virtual value is either strictly positive or negative, *i.e.*, $\phi(l, x_l) > 0$. Accordingly, the benchmark allocation is *unique*, given by $\mathbb{Q}(l, x) = \mathbb{1}_{x \geq x_l}$.

Theorem 1 (Binary types). *Fix $\Theta = \{h, l\}$. There exists some $\lambda \in [0,1]$ and $\hat{x}_l \in$, such that in the unique optimal mechanism, the allocation is given by*

$$q(h,x) = 1 \quad \forall x, \quad q(l,x) = \begin{cases} 1 & \text{if} \quad x > \hat{x}_l, \\ 0 & \text{if} \quad x < \hat{x}_l, \\ \lambda & \text{if} \quad x = \hat{x}_l. \end{cases}$$

Moreover,

(a) If $\pi_l^\star$ induces the threshold flip of type order, $(\hat{x}_l, \lambda) = (x_l, 1)$. A menu of posted prices and threshold disclosures is optimal.

(b) If $\pi_l^\star$ does not induce the threshold flip of type order, $(\hat{x}_l, \lambda) \neq (x_l, 1)$. A posted price, associated with a uniform threshold disclosure, is optimal.

In short, Theorem 1 states that the optimal mechanism features *screening* whenever $\pi_l^\star$ leads to the threshold flip of type order and *bunching* otherwise. Intuitively, when the unknown component x dominates the buyer's private type θ in triggering the variation of the buyer's valuation (to induce the threshold flip of type order), the new information matters and helps screen the buyer. Conversely, when the ranking of willingness to pay mainly depends on the buyer's type, screening disappears. Then, the optimal mechanism closely resembles its counterpart in standard mechanism design where the state is known: a posted price (but associated with threshold disclosure) is optimal. The optimal mechanism in each case is explicitly characterized in the remainder of this section. To illustrate Theorem 1, consider the following examples.

Example 1 (Binary types and states). $\Theta = \{l, h\}$ *and* $X = \{b, g\}$. *Types and states are equally likely.*

Assume that $\phi(\theta_1, x_1) < 0 < \phi(\theta_1, x_2)$ *to make the problem non-trivial.*

In this simple binary-type, binary-state setting, there are two scenarios of optimal mechanisms. If $v(\theta_3, x_1) \geq v(\theta_1, x_2)$, then $\pi^\star$ does not induce the threshold flip of type order. By Theorem 1(a), a fixed price and threshold disclosure is optimal. On the other hand, if $v(\theta_3, x_1) < v(\theta_1, x_2)$, then $\pi^\star$ leads to the threshold flip of type order. By Theorem 1(b), a menu of prices and threshold disclosures is optimal.

Example 2. $\Theta = \{l, h\}$ *and* X *is a finite subset of* $\mathbb{N}$*. Types and states are equally likely. Valuations are given by:* $v(\theta, x) = \theta + x$.

Let

$$\Delta_\theta \equiv v(h, x) - v(l, x) = h - l \quad \forall x,$$
$$\Delta_x \equiv v(\theta, \overline{x}) - v(\theta, \underline{x}) = \overline{x} - \underline{x} \quad \forall \theta.$$

Then, Δ_θ represents the variation of valuation due to the buyer's type, whereas Δ_x due to the state x. For any state $\hat{x} \in \Omega$,

$$\mathbb{E}[v(h, x) \mid x < \hat{x}] - \mathbb{E}[v(l, x) \mid x \geq \hat{x}] = \left(h + \frac{\hat{x} - 1 + \underline{x}}{2}\right) - \left(l + \frac{\hat{x} + \overline{x}}{2}\right) = \Delta_\theta - \frac{\Delta_x + 1}{2},$$

Thus, the threshold flip of type order happens if and only if

$$\Delta_\theta \leq \frac{\Delta_x + 1}{2}, \tag{1.6}$$

which is the case when the impact of the buyer's type is relatively small, relative to that of the unknown component. By Theorem 1, when (1.6) holds, it is optimal to offer a menu of threshold disclosures and posted prices. Otherwise, a posted price, coupled with uniform threshold disclosure, maximizes the seller's revenue.[22]

Remark 1. *Theorem* 1 *and its proof extends readily to the case with a continuum of states. As an example, fix* $\Theta = \{l, h\}$ *and* $X = [0, 10]$*, and both* θ *and* x *are uniformly distributed. Then, for any state* $\hat{x} \in \Omega$, $\mathbb{E}[v(h, x) \mid x < \hat{x}] - \mathbb{E}[v(l, x) \mid x \geq \hat{x}] = \Delta_\theta - 5$*. Thus, a menu of prices and information is optimal if* $\Delta \geq 5$ *and a fixed price coupling with a threshold disclosure (for all types) is optimal if* $\Delta < 5$.

Theorem 1 has two important implications:

[22] In particular, when Δ_θ is too high, the seller does not benefit from information disclosure. In this case, the optimal threshold for type l is the highest state ($\hat{x}_l = \overline{x}$), which means no disclosure is provided.

Corollary 1. *With* $|\Theta| = 2$*, privacy of signals does not matter when the threshold flip of type order happens under* $\gamma_l^\star$*. It matters otherwise.*

Corollary 2. *With* $|\Theta| = 2$*, the seller does not strictly benefit from using random mechanisms, nor from eliciting signals.*

What leads to the (ir)relevance of signal privacy and the optimality of deterministic mechanisms, signal-independent allocations will be explained when we present the key steps of the proof of Theorem 1, to which we turn next.

5.1 Proof of Theorem 1

To prove Theorem 1, we solve a relaxed problem, denoted by $(\mathscr{RP}_b)$, ignoring (IC_{lh}) and (IR_h) and provide an implementation. Formally, this relaxed problem is as follows.

$$(\mathscr{RP}_b) \quad \max_{(\pi,q,U)} \quad \sum_\theta f(\theta)\Big[\sum_x\sum_s v(l,x)q(\theta,\omega^{\pi_\theta}(\theta,s))\pi_\theta(s|x)\mu(x) - U(\theta)\Big]$$

$$s.t: \quad U(h) - U(l) \geq \sum_x\sum_s \int_{\omega^{\pi_l}(l,s)}^{\omega^{\pi_l}(h,s)} q(l,z)dz\pi_l(s|x)\mu(x) \qquad (IC_{h\to l})$$

$$U(l) \geq 0 \qquad (IR_l)$$

$$q(\theta,\omega) \text{ increases in } \omega. \qquad (MON)$$

The characterization of the solution to $(\mathscr{RP}_b)$ is done via the following steps. First, we prove the optimality of deterministic allocation rules. This step, while standard, is helpful in decomposing the buyer's rent into two components: the *ex ante* rent (due to privacy of types) and the *posterior* rent (due to privacy of signals). Using this rent decomposition, we establish the optimality of binary-signal experiments and furthermore, of threshold disclosures. Finally, we characterize the optimal allocation and implement it.

First, to obtain the optimality of deterministic allocations, note that $(IC_{h\to l})$ and (IR_l) must bind in $(\mathscr{RP}_b)$, *i.e.*,

$$U(l) = 0, \quad U(h) = \sum_x\sum_s \int_{\omega^{\pi_l}(l,s)}^{\omega^{\pi_l}(h,s)} q(l,z)dz\pi_l(s|x)\mu(x),$$

Thus, transfers have been eliminated, reducing the seller's relaxed problem to

$$\begin{aligned}\max_{q,\pi}\quad & f(h)\sum_x\sum_s v(h,x)q(h,\omega^{\pi_\theta}(\theta,s))\pi_h(s|x)\mu(x)\\ & +f(l)\sum_x\sum_s\left[v(l,x)q(l,\omega^{\pi_\theta}(\theta,s))-\int_{\omega^{\pi_l}(l,s)}^{\omega^{\pi_l}(h,s)}q(l,z)dz\right]\pi_l(s|x)\mu(x)\\ s.t\quad & q(\theta,\omega)\text{ increases in }\omega. \qquad (MON)\end{aligned}$$

Fix π. Given that the objective function is linear and the only constraint is (MON), there exists an optimal allocation rule that is deterministic and exhibits a cut-off structure. Moreover, as $v(h,x)$ is always non-negative, h receives an efficient allocation.

Lemma 5 (Deterministic allocations). *In ($\mathscr{RP}_b$), there exists an optimal allocation rule, given by $q(\theta,\omega)=\mathbb{1}_{\omega\geq\hat{\omega}_\theta}$, where $\hat{\omega}_h=v(l,\underline{x})$.*

Second, we derive the sufficiency of binary-signal experiments. As $q(h,s)=1$ for all s, any π_h is optimal. The relaxed problem reduces to finding the optimal π_l. Let

$$R_l \equiv f(l)\sum_x\sum_s\left[v(l,x)q(l,\omega^{\pi_l}(l,s))-\int_{\omega^{\pi_l}(l,s)}^{\omega^{\pi_l}(h,s)}q(l,z)dz\right]\pi_l(s|x)\mu(x)$$

denote the term involving π_l in the seller's objective function (revenue) in ($\mathscr{RP}_b$). Using $q(l,\omega)=\mathbb{1}_{\omega\geq\hat{\omega}_l}$ by Lemma 5, we obtain

$$\begin{aligned}R_l=&f(l)\sum_x\left[\sum_{\hat{s}_l}^{\overline{s}}v(l,x)\pi_l(s|x)-\sum_{\hat{s}_l}^{\overline{s}}\int_{\omega^{\pi_l}(l,s)}^{\omega^{\pi_l}(h,s)}dz\frac{f(h)}{f(l)}\pi_l(s|x)-\sum_{\underline{s}}^{\hat{s}_l}\int_{\omega^{\pi_l}(l,\hat{s}_l)}^{\omega^{\pi_l}(h,s)}dz\frac{f(h)}{f(l)}\pi_l(s|x)\right]\mu(x)\\ =&f(l)\sum_x\Bigg[\underbrace{\sum_{\hat{s}_l}^{\overline{s}}v(l,x)}_{l\text{'s surplus}}-\underbrace{\sum_{\hat{s}_l}^{\overline{s}}[\omega^{\pi_l}(h,s)-\omega^{\pi_l}(l,s)]\frac{f(h)}{f(l)}}_{h\text{'s } ex\ ante \text{ rent}}-\underbrace{\sum_{\underline{s}}^{\hat{s}_l}[\omega^{\pi_l}(h,s)-\omega^{\pi_l}(l,\hat{s}_l)]\frac{f(h)}{f(l)}}_{h\text{'s } posterior \text{ rent}}\Bigg]\pi_l(s|x)\mu(x).\end{aligned}$$

R_l depends on (i) the buyer's expected value (on path for l and off path for h), conditional on whether $s\geq\hat{s}_l$ or $s<\hat{s}_l$, and (ii) the cut-off signal, $\hat{s}_l$. By (i), there is no revenue loss in replacing all signals $s\geq\hat{s}_l$ with "*good news*" (s^g) and all $s<\hat{s}_l$ with "*bad news*" (s^b). At the same time, such a binary-signal experiment for type l increases the cut-off signal because $\omega^{\pi_l}(l,s^g)=\mathbb{E}[v(l,x)\mid s\geq\hat{s}_l]\geq\hat{s}_l$. In turn, this improves R_l, which increases in the cut-off signal. We thus obtain the optimality of binary-signal experiments.

Lemma 6 (Binary signals). *In ($\mathscr{RP}_b$), there exists an optimal experiment for type l where the signal realization can be either "good news" (s^g) or "bad news" (s^b).*

Third, we prove the optimality of threshold disclosures. By replacing all signals $s \geq \hat{s}_l$ (resp., $s < \hat{s}_l$) with "*good news*" (resp., "*bad news*"), R_l becomes

$$f(l)\sum_x \Big[\underbrace{v(l,x)\pi_l(s^g|x)}_{l\text{ 's surplus}} - \underbrace{[v(h,x)-v(l,x)]\frac{f(h)}{f(l)}\pi_l(s^g|x)}_{h\text{'s } \textit{ex ante} \text{ rent}} - \underbrace{[\omega^{\pi_l}(h,s^b)-\omega^{\pi_l}(l,s^g)]\frac{f(h)}{f(l)}\pi_l(s^b|x)}_{h\text{'s } \textit{posterior} \text{ rent}} \Big]\mu(x)$$

$$= \underbrace{f(l)\sum_x \phi(l,x)\pi_l(s^g|x)\mu(x)}_{l\text{ 's virtual value}} - f(l)\sum_x \underbrace{\max\Big\{[\omega^{\pi_l}(h,s^b)-\omega^{\pi_l}(l,s^g)]\frac{f(h)}{f(l)},0\Big\}\pi_l(s^b|x)\mu(x)}_{h\text{'s } \textit{posterior} \text{ rent}}.$$

Fix $\pi_l(s^b)$. Then, a threshold disclosure minimizes h's posterior rent by simultanenously maximizing $\omega^{\pi_l}(h,s^b)$ and minimizing $\omega^{\pi_l}(h,s^b)$. Moreover, as $\phi(l,x)$ increases in x, a threshold disclosure maximizes l 's expected virtual value. Therefore:

Lemma 7 (Threshold structure)**.** *In ($\mathscr{RP}_b$), a threshold disclosure for l is optimal.*

Last, we characterize the optimal allocation and provide an implementation. Let $\hat{x}_l \in X$ be the cut-off state associated with the optimal threshold disclosure for l and $\lambda \in [0,1]$ be the probability with which "*good news*" is sent at $\hat{x}_l$. Then, by Lemmas 5, 6, and 7, the optimal allocation is given by

$$q(h,x) = 1 \quad \forall x, \quad q(l,x) = \begin{cases} 1 & \text{if} \quad x > \hat{x}_l, \\ 0 & \text{if} \quad x < \hat{x}_l, \\ \lambda & \text{if} \quad x = \hat{x}_l. \end{cases}$$

Solving for the optimal allocation reduces to solving for the optimal $(\hat{x}_l, \lambda)$. As will be shown, there are two cases, depending on whether $\pi_l^\star$ triggers the threshold flip of type order. In the first case, when this flip happens, offering $\pi_l^\star$ with $(\hat{x}_l,\lambda) = (x_l, 1)$ is optimal. Not only does it induce zero *posterior* rent for h, given that

$$\omega^{\pi_l^\star}(h,s^b) - \omega^{\pi_l^\star}(l,s^g) \leq 0$$

when the threshold flip occurs under $\pi_l^\star$, but it also creates the highest expected virtual value for l 's, given by $f(l)\sum_{x\geq x_l}\phi(l,x)\pi_l(s^g|x)\mu(x)$.

With $(\hat{x}_l,\lambda) = (x_l, 1)$, l 's allocation coincides with the benchmark $\mathbb{Q}(l,x) = \mathbb{1}_{x\geq x_l}$. To find out payments, without loss of generality, assume the buyer pays only if "*good news*" is realized (or trade happens). Thus, $p(h,s^b) = p(l,s^b) = 0$. Then, $p(l,s^g) = \omega^{\pi_l^\star}(l,s^g)$ by (IR_l), and $p(h,s^g)$ is such that $(IC_{h\to l})$ holds, or $U(h) = U(h,l)$, which implies

$$p(h,s^g) = \mathbb{E}[v(h,x)] - \left[\omega^{\pi_l^\star}(h,s^g) - p(l,s^g)\right]\pi_l^\star(s^g).$$

Now, verify that ignored constraints are satisfied. First,(IR_h) hold because

$$U(h) = \left[\omega^{\pi_l^\star}(h,s^g) - \omega^{\pi_l^\star}(l,s^g)\right]\pi_l^\star(s^g) \geq 0$$

Second, $IC_{l\to h}$ is satisfied given that

$$\begin{aligned}
U(l,h) &= \mathbb{E}[v(l,x)] - p(h,s^g) \\
&= \mathbb{E}[v(l,x)] - \mathbb{E}[v(h,x)] + \left[\omega^{\pi_l^\star}(h,s^g) - p(l,s^g)\right]\pi_l^\star(s^g) \\
&= \mathbb{E}[v(l,x)] - \mathbb{E}[v(h,x)] + \left[\omega^{\pi_l^\star}(h,s^g) - \omega^{\pi_l^\star}(l,s^g)\right]\pi_l^\star(s^g) \\
&= \left[\omega^{\pi_l^\star}(l,s^b) - \omega^{\pi_l^\star}(h,s^b)\right]\pi_l^\star(s^b) < 0 = U(l)
\end{aligned}$$

Moreover, under no threshold flip of type order by $\pi^\star$, $\omega^{\pi_l^\star}(h,s^b) \leq \omega^{\pi_l^\star}(l,s^g) = p(l,s^g)$. Therefore, if h mimics l, it is optimal for him to report signals truthfully. This deviating behavior is not beneficial for h by the construction of $p(h,s^g)$. We thus obtain Theorem 1(a):

Lemma 8 (With threshold flip by $\pi_l^\star$)**.** *If $\pi_l^\star$ induces the threshold flip of type order, $q(l,x) = \mathbb{Q}(l,x) = \mathbb{1}_{x\geq\hat{x}_l}$, and $\boldsymbol{M}^\star \equiv \{p^\star(\theta),\pi_\theta^\star\}_\theta$ is optimal.*

By contrast, when $\pi_\theta^\star$ preserves the type order, or $\omega^{\pi_l^\star}(h,s^b) > \omega^{\pi_l^\star}(l,s^g)$, offering $\pi_l^\star$ to l induces a strictly positive *posterior* rent for h. Consequently, the seller trades off between l 's expected virtual value and h's *posterior* rent. On the one hand, she wants the threshold to be close to the cut-off x_l, maximizing l 's expected value. On the other hand, she desires to induce a small *posterior* rent for h.

Let $\pi_l^{\star\star}$ be an optimal experiment for l , associated with $\left(x^{\star\star}(l),\lambda^{\star\star}\right)$. Suppose, $\pi_l^{\star\star}$ can flip the type order, *i.e.*, $v^{\pi_l^{\star\star}}(h,s^b) < \omega^{\pi_l^{\star\star}}(l,s^g)$. Then, given that $\omega^{\pi_l^\star}(h,s^b) \leq \omega^{\pi_l^\star}(l,s^g)$, we can construct $\tilde{\pi}_l$ associated with $(\tilde{\omega},\tilde{\lambda})$ such that (i) $\left(x^{\star\star}(l),\lambda^{\star\star}\right)$ is closer to $\left(x_l^\star,1\right)$ and (ii) $v^{\tilde{\pi}_l}(h,s^b) \leq v^{\tilde{\pi}_l}(l,s^g)$. By (i), l 's expected virtual value under $\tilde{\pi}_l$ is higher than that under $\pi_l^\star$, whereas by (ii), h's *poterior* rent is zero under $\tilde{\pi}_l$. This contradicts with $\pi_l^{\star\star}$ being optimal. Therefore, $\pi_l^{\star\star}$ must not affect the type order. Formally:

Claim 1. $\omega^{\pi_l^{\star\star}}(h,s^b) \geq \omega^{\pi_l^{\star\star}}(l,s^g)$.

The detailed proof is in Appendix B.1. By Claim 1, R_l reduces to

$$\begin{aligned} & f(l)\sum_x \phi(l,x)\pi_l^{\star\star}(s^g|x)\mu(x) - f(h)\left[\omega^{\pi_l^{\star\star}}(h,s^b) - \omega^{\pi_l^{\star\star}}(l,s^g)\right]\pi_l(s^b) \\ =& \omega^{\pi_l^{\star\star}}(l,s^g)\left[f(l)\pi_l^{\star\star}(s^g|x)\mu(x) + f(h)\right] - \mathbb{E}[v(h,x)]. \end{aligned}$$

Therefore,

$$\pi_l^{\star\star} \in \underset{\pi_l}{\text{argmax}}\, \omega^{\pi_l}(l,s^g)\left[f(l)\pi_l^{\star\star}(s^g,x) + f(h)\right]. \tag{1.7}$$

To find optimal transfers, note that by Claim 1, h's value after *"bad news"* is higher than l 's after "*good news.*" Hence, if h mimics l, he always reports "*good news,*" and always buys the good. Consequently, l 's allocation is the same as h 's from the latter's perspective. This leads to a bunching solution. As information is of no value for h, the seller can offer $\pi_l^{\star\star}$ to both types. Moreover, as h always gets the good either on or off-path, by $(IC_{h\to l})$, both types receive the same posted price.[23] Then, by (IR_l),

$$p^{\star\star}(h) = p^{\star\star}(l) = \omega^{\pi_l^{\star\star}}(l,s^g). \tag{1.8}$$

This bunching mechanism satisfies ignored constraints, and hence, is optimal.

Lemma 9 (No threshold flip by $\pi_l^\star$)**.** *If $\pi_l^\star$ does not induce the threshold flip of type order, $(\hat{x}_l,\lambda) \neq (x_l,1)$. A single-option menu, $\{\pi_l^{\star\star}, p^{\star\star}(l)\}$ given by (1.7) and (1.8), is optimal.*

6 OPTIMAL MECHANISM FOR $|\Theta| \geq 3$

With binary types, there are two scenarios of the optimal mechanism (screening/bunching), depending on whether after information disclosure, the threshold flip of type order occurs or not. With richer type sets, it can be the case that information disclosure flips the order of a group of types but fails to do so for another group. Consequently, the characterization of optimal mechanisms cannot be obtained as a simple extension of that in the binary-type case. Moreover, as we will show, random mechanisms could be used to effectively screen signals and distant types. Despite these complications, we show that the optimality of a rich (respectively, single-option) menu of prices and threshold disclosure extends beyond the binary-type setting to a general model under stronger notions of type order flip (respectively, preservation). This result is presented in Section 5, followed by an analysis on the role of random mechanisms in 6.1.

[23] With deterministic allocations, it is without loss to offer a menu of posted prices. See Proposition 1.

6.1 Revenue improvement via random mechanisms

Using Example 3 below, we illustrate how random mechanisms outperform their deterministic counterparts in two aspects (i) screening distant types and (ii) screening signals to improve the seller's revenue and efficiency.[24]

Example 3. $\Theta = \{\theta_1, \theta_2, \theta_3\}$ *and* $X = \{x_1, x_2\}$. *Types and states are equally likely. Valuations are as follows.*

$v(\theta, x)$	x_1	x_2
θ_3	6.5	10
θ_2	0	7
θ_1	0	4

Table 1.1: Example 3(a)

$v(\theta, x)$	x_1	x_2
θ_3	5	5
θ_2	2	5
θ_1	0	4

Table 1.2: Example 3(b)

Example 3(a) - Screening distant types: In this example, type θ_3's value is always higher than type θ_1's. This leaves room for random mechanisms to "separate" these two types. To see this, note the following. If the seller employs deterministic mechanisms, type θ_1 either trades or not at any signal realization. Therefore, if type θ_1 trades (with probability 1) for some signal, it is optimal for θ_3 whose posterior value is always higher than type θ_1's, having mimicked θ_1, to (mis)report the realized signal such that he always trades. Then, θ_1 's allocation is the same as θ_3's from the latter's perspective, leading to bunching these types.[25] In turn, this gives too much rent for type θ_3, making it optimal to *exclude* type θ_1.

Claim 2. *In Example* 3*(a), if only deterministic mechanisms are allowed, it is optimal to offer type* θ_3 *with no disclosure and a posted price* $p(\theta_3) = 6.75$, *type* θ_2 *with full disclosure and a posted price* $p(\theta_2) = 7$, *and to exclude type* θ_1.

The story, however, is different with random allocations. The key is that if θ_1 trades with a small probability (for any signal), this type's allocation becomes unattractive to θ_3. To see this, modify the optimal deterministic mechanism by letting θ_1 trade with a probability $\varepsilon \in [0, \frac{3}{4}]$ and

[24]In the Online Appendix, we fully characterize the optimal random mechanism in several examples.

[25]In Example 3(a), if the seller employs deterministic mechanisms and serves type θ_1, a fixed price $p = 4$, associated with full disclosure is optimal.

adjusting transfers such that truth-telling remains satisfied, as follows:

$$q(\theta_3,x)=1\forall x,\quad q(\theta_2,x)=\mathbb{1}_{x=b},\quad q(\theta_1,x)=\begin{cases}\varepsilon & \text{if} \quad x=g,\\ 0 & \text{if} \quad x=b,\end{cases}$$

$$p(\theta_3)=6.5-\varepsilon,\quad p(\theta_2)=7-2\varepsilon,\quad p(\theta_1)=5 \quad \text{paid conditional on trade occurs.}$$

Then, expected payment by θ_3 and θ_2 reduces by ε; however, that by θ_2 increases by $\frac{5\varepsilon}{2}$. Overall, the seller's revenue increases by $f(\theta_1)\frac{5\varepsilon}{2}-\left[f(\theta_2)+f(\theta_3)\right]\varepsilon=\frac{3\varepsilon}{2}>0$. Therefore, random allocation helps the seller screen effectively distant types (types θ_3 and θ_1), thereby, improving trade surplus extensively as well as the seller's revenue.

Example 3**(b) - Screening signals**: In this example, type θ_2's value varies significantly across states. This makes it optimal to exclude type θ_2 at state x_1, rather than "pooling" the two states under deterministic mechanisms which allow either trade or no trade at any signal realization. Formally, the optimal deterministic mechanism, stated in Claim 3 below, specifies:

$$q(\theta_3,x)=1\forall x,\quad q(\theta_2,x)=q(\theta_1,x)=\mathbb{1}_{x=x_2},$$

which are implemented via full disclosure and a fixed price.

Claim 3. *In Example* 3*(b), if only deterministic mechanisms are allowed, it is optimal to offer full disclosure and a posted price* $p=4$.

Random mechanisms, on the other hand, arm the seller with the flexibility in designing trade probabilities. This helps her screen realized states by allowing trade to happen at a small probability at low states. To see this, revise the optimal deterministic mechanism by letting θ_2 trade with probability $\delta\le\frac{1}{3}$, such that now:

$$q(\theta_3,x)=1\forall x,\quad q(\theta_1,x)=\mathbb{1}_{x=x_2},\quad p(\theta_3)=p(\theta_1)=4,$$

$$\big(q(\theta_2,x),p(\theta_2,x)\big)=\begin{cases}(1,4) & \text{if} \quad x=x_2\\ (\delta,2\delta) & \text{if} \quad x=x_1\end{cases},\quad \text{with } \delta\le\frac{1}{3}.$$

This revised mechanism differs from the optimal deterministic mechanism only in the new trade created with type θ_2 at state x_1. Therefore, as long as this new trade creation preserves incentive compatibility, the seller's revenue increases by $2f(\theta_2)g(x_1)\delta>0$. We show that this is the case in the Online Appendix.

6.2 Screening vs. Bunching

This section generalizes the finding of optimal mechanisms with binary types (Theorem 1) to a general model with finitely many types.

6.2.1 Optimality of screening

Recall that information disclosure can be used to screen the buyer of binary types when it induces a threshold flip of type order. Similarly, information serves as a screening tool in a richer type space under the following notion of type order flip:

Definition 6 (Partition flip of type order)**.**

The partition flip of type order happens if $\mathbb{E}[v(\theta^+, x) \mid x_{\theta^+} \leq x < x_\theta]$ *decreases in* θ.

Under the partition flip of type order, the expected valuations over relevant partitions of states decrease in types. As the relevant partition for a higher type consists of lower states, such a type order flip requires the new information (about the state) to sufficiently dominate the buyer's initial type in driving valuation fluctuations. Indeed, it coincides with the threshold flip notation when there are only two types. In a richer type set, more than one interior threshold is involved under the menu of threshold disclosure $\{\pi^\star_\theta\}_\theta$, leading to relevant partitions of states.

Theorem 2 below states the optimal mechanism under the partition flip of type order, which features discriminatory information and prices.

Theorem 2 (Screening)**.** *Under the partition flip of type order, the optimal allocation is given by* $\mathbb{Q}(\theta, x) = \mathbb{1}_{x \geq \omega^\star(\theta)}$. *A menu of posted prices and threshold disclosures is optimal.*

This result extends Theorem 1(a) to a model with more than two types, following the same logic: when information disclosure matters sufficiently, it helps screen the buyer. The only difference is that the *partition flip* of type order is required here, taking into account interior types.

The proof proceeds by showing that under the partition flip of type order, $\mathbf{M}^\star$ induces truth-telling even if the seller does not observe signals. Therefore, offering $\mathbf{M}^\star$ with the buyer privately observing signals is equivalent to offering a menu of posted prices and threshold disclosures $\{p^\star(\theta), \pi^\star_\theta\}_\theta$, where the posted price is equal to the payment paid after "*good news*" in $\mathbf{M}^\star$: $p^\star(\theta) = p^\star(\theta, s^g)$. This menu helps the seller achieve the upper bound of revenue attained when signals are public signals; hence, it is optimal.

We close this section with an illustrative example.

Example 4. $\Theta = \{h, m, l\}$. *X is a finite subset of* $\mathbb{N}$. *Types and states are equally likely. Valuations are given by* $v(h,x) = x + \Delta_\theta$, $v(m,x) = x$, $v(l,x) = x - \Delta_\theta$. *Accordingly, virtual values are given by* $\phi(h,x) = x + \Delta_\theta$, $\phi(m,x) = x - \Delta_\theta$, $\phi(m,x) = x - 3\Delta_\theta$.

In this example, $v(\theta^+, x) - v(\theta, x) = \Delta_\theta \ \forall x$ and $\Delta_x \equiv v(\theta, \overline{x}) - v(\theta, \underline{x}) = \overline{x} - \underline{x} \ \forall \theta$. In addition, $x_h^\star = \underline{x}$, $x_m^\star = \Delta_\theta$, and $x_l^\star = 3\Delta_\theta$, which implies

$$\begin{aligned}
\mathbb{E}[v(h,x) \mid x_h^\star \leq x < x_m] &= \frac{3\Delta_\theta - 1 + \underline{x}}{2}, \\
\mathbb{E}[v(m,x) \mid x_m \leq x < x_l] &= \frac{4\Delta_\theta - 1}{2}, \\
\mathbb{E}[v(l,x) \mid x_l \leq x \leq \overline{x})] &= \frac{\Delta_\theta + \overline{x} - 1}{2}
\end{aligned}$$

Thus, the partition flip of type order happens if

$$3\Delta_\theta - 1 + \underline{x} \leq 4\Delta_\theta - 1 \leq \Delta_\theta + \overline{x} - 1 \Leftrightarrow \underline{x} \leq \Delta_\theta \leq \Delta_x,$$

which requires the impact of the unknown component to be higher than that of the buyer's type (and is of at least $\underline{x}$). If this is the case, by Theorem 2, it is optimal to screen the buyer's type using different bundles of posted prices and threshold disclosure.

6.2.2 Optimality of bunching

In the binary-type case, the benefit of screening disappears if the threshold disclosure rule $\pi_l^\star$ fails to flip the ranking of willingness to pay by types. A similar story holds with more than two types under a stronger notion of (no) threshold flip of type order:

Definition 7 (Uniformly no threshold flip of type order)**.** *Under uniformly no threshold flip of type order,*

$$\mathbb{E}[v(\theta^+, x \mid x < \hat{x}] \geq \mathbb{E}[v(\theta, x) \mid x \geq \hat{x}] \quad \forall \theta \in \Theta, \forall \hat{x} \in X.$$

In words, this condition satisfies if under *any threshold disclosure* and for *any type* θ: θ^+'s value after "*bad news*" must be higher than θ's after "*good news*". This is more likely to hold when valuation heterogeneity is mainly driven by the buyer's type. For instance, when θ^+'s values are always higher regardless of states, *i.e.*, $v(\theta^+, \underline{x}) \geq v(\theta, \overline{x})$, it is impossible to flip their ranking of valuation after *any rule* of information disclosure, not just the threshold ones.

We are now ready to state the main result of this section.

Theorem 3 (Bunching). *Under uniformly no threshold flip of type order, a posted price, associated with a threshold disclosure, is optimal.*

This result extends Theorem 1(b), carrying the same intuition: when valuation heterogeneity is mainly due to the buyer's types, information about the state becomes inessential for (most types of) the buyer; as a result, its screening function shuts off. The only difference is that no type order flip by *any* threshold disclosure is required here, of which the role is to be explained.

The proof proceeds by solving a relaxed problem considering only deviating behaviors under which all types mimic the lowest type being served. This problem mirrors that for the binary-type case $\Theta = \{h, l\}$, with the lowest type being served representing type l and all the other types echoing type h. The optimality of bunching under *uniformly* no threshold flip of type order follows similar arguments for that in the binary-type setting under no threshold flip by $\pi_l^\star$. The lowest type being served, and thereby, the optimal posted price and threshold disclosure can be explicitly characterized, leveraging the fact that no threshold flip holds *uniformly* regardless of pairs of types and threshold rule.

To end this section, revisit Example 4 for an illustration. In this example, for any $\hat{x} \in X$,

$$\mathbb{E}[v(h,x) \mid x < \hat{x}] - \mathbb{E}[v(m,x) \mid x \geq \hat{x}] = \left(\Delta_\theta + \frac{\hat{x} - 1 + \underline{x}}{2}\right) - \left(m + \frac{\overline{x} + \hat{x}}{2}\right) = \Delta_\theta - \frac{\Delta_x - 1}{2},$$
$$\mathbb{E}[v(m,x) \mid x < \hat{x}] - \mathbb{E}[v(l,x) \mid x \geq \hat{x}] = \left(m + \frac{\hat{x} - 1 + \underline{x}}{2}\right) - \left(l - \Delta_\theta + \frac{\overline{x} + \hat{x}}{2}\right) = \Delta_\theta - \frac{\Delta_x - 1}{2},$$

where, just to recall, Δ_θ and Δ_x measure the impact of the buyer's private type and the unknown component in valuation variations, respectively. Therefore, uniformly no threshold flip of type order occurs if

$$\Delta_\theta \geq \frac{\Delta_x - 1}{2},$$

which requires the buyer's type to be significantly impactful, relative to the unknown component. If this is the case, by (Theorem 3), information is not leveraged to screen the buyer. A single price-threshold disclosure bundle is optimal.

7 INFINITE-TYPE SETTING

All the proofs of our results extend readily if there is a continuum of states. The extension to the infinite-type case, however, is not trivial. Nevertheless, we find that the previous insights remain valid: Section 7.1 shows that a menu of prices and threshold disclosure is optimal under the partition flip the ranking of willingness to pay across cut-off types; and Section 7.2 shows that a fixed price- threshold disclosure bundle is approximately optimal when the type order is almost preserved.

Throughout this section, consider a continuum of types $\Theta = [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$, endowed with the distribution $F(\theta)$. We assume that $F(\theta)$ is differentiable in θ with density $f(\theta)$, and moreover, $v(\theta, x)$ is differentiable in θ. Then, the virtual value in this environment is given by

$$\phi^c(\theta, x) = v(\theta, x) - v_\theta(\theta, x)\frac{1 - F(\theta)}{f(\theta)}.$$

Similar to the finite-type case, we assume that $\phi^c(\theta, x)$ increases in θ and x.

7.1 Optimality of a screening menu

By the monotonicity of the virtual values, each state x is associated with a cut-off type θ_x above (respectively, below) which the buyer's virtual value is non-negative (respectively, negative). Formally,

$$\theta_x \equiv \inf\{\theta \mid \phi^c(\theta, x) \geq 0\}.$$

Moreover, as $\phi^c(\theta, x)$ increases in x, this cut-off type θ_x decreases in x. We use

$$\Theta_x \equiv \{\theta_x\}_{x \in X}$$

to denote the type space consisting of only cut-off types. Even with a continuum of types, there are *finitely* many cut-off types $\{\theta_x\}_{x \in X}$ due to the finiteness of the state space. Accordingly, $\mathbf{M}^\star$ comprises $|\Theta_x|$ options of prices and disclosure rules because each interval of types $[\theta_{x^+}, \theta_x)$ is assigned the same option. Then, the following proposition can be obtained following the proof of Theorem 2 for a type space consisting of only the cut-off types.

Proposition 3. *Fix $\Theta = [\underline{\theta}, \overline{\theta}]$ and $|X| < \infty$. If there is a partition flip of type order within Θ_x, a menu of threshold disclosures and posted prices is optimal.*

This result holds even if there is a continuum of states $X = [\underline{x}, \overline{x}]$ and the valuation function is continuous over states, by approximating an associated finite-state model as the distance between states approaches zero. In this case, the partition flip of type order reduces to the valuation at the cut-off state $v(\theta, x_\theta)$ decreasing in types.[26]

[26]This is the case in, for example, the environments studied in Eső and Szentes (2007) and Wei and Green (2023) under which the valuation function is concave in types and states, and the cross derivative is positive.

7.2 (Approximate) optimality of bunching

When valuations shift smoothly across (a continuum of) types, there are always types whose valuations are sufficiently close to others'. This makes it impossible to preserve the ranking of willingness to pay uniformly across the types. Consequently, the optimality of bunching cannot be derived as an extension of Theorem 3 which shows that under the uniformly no threshold flip of valuation ranking across finitely many types, a fixed price-information bundle is optimal. Nevertheless, we establish the *approximate* optimality of bunching under ε*-uniformly* no threshold flip of type order, formally defined below.

Definition 8 (ε-uniformly no threshold flip of type order)**.** *ε-uniformly no threshold flip of type order occurs if for some $\varepsilon > 0$,*

$$\mathbb{E}[v(\theta+\varepsilon, x) \mid x \leq \hat{x}] \geq \mathbb{E}[v(\theta, x) \mid x \geq \hat{x}] \quad \forall \theta, \hat{x}.$$

The following proposition shows that as ε vanishes, the seller's maximized revenue can be approximated by offering via a fixed price-threshold disclosure bundle. Formally, let R_ε represent the revenue guarantee if the seller offers a single posted price and threshold disclosure rule under the ε-uniformly no threshold flip of type order, we find that:

Proposition 4. $R_\varepsilon \to V(P)$ *as* $\varepsilon \to 0$

Moreover, if there are only two states $\Omega = \{b, g\}$ with $b < g$, we establish the exact optimality of a fixed price and disclosure rule within the class of deterministic mechanisms.

Proposition 5. *Fix $\Theta = [\underline{\theta}, \overline{\theta}]$ and $X = \{b, g\}$. If $v(\theta_b, b > v(\theta_g, g)$ and only deterministic allocations are allowed, a posted price, associated with full disclosure, is optimal.*

The idea of the proof is as follows. With binary states $X = \{b, g\}$, there are only two cut-off types θ_g and θ_b. Hence, the partition flip of type order reduces to $v(\theta_b, b) \leq v(\theta_g, g)$. If this is the case, a menu of prices and threshold disclosures is optimal by Proposition 3. If by contrast, $v(\theta_b, b) \leq v(\theta_g, g)$, the seller adjusts the cut-off types to $\tilde{\theta}_b, \tilde{\theta}_g$ *just enough* to restore the partition flip of type order: $v(\tilde{\theta}_b, b) = v(\tilde{\theta}_g, g)$. In turn, at this boundery of the partition flip, the seller is indifferent between offering a screening menu and a single option of price and information. Put differently, bunching is optimal.

We end this section with a numerical example to illustrate Proposition 5.

Example 5. $v(\theta, x) = 3\theta^2 + 6\theta + x$, $\Theta = [0, 2]$, $X = \{8, 12\}$. *Types and states are likely equally.*

In this example, $\phi(\theta,x) = 3\theta^2+6\theta+x-(6\theta+6)(2-\theta) = 9\theta^2+x-12$. Thus, $\theta_{12}=0$ and $\theta_8=\frac{2}{3}$. Hence, $v(\theta_8,8)=\frac{43}{3}$ and $v(\theta_{12},12)=12$. As $v(\theta_8,8)>v(\theta_{12},12)$, no flip of type order occurs. By Proposition 5, within the class of deterministic mechanism, offering a fixed bundle of price and threshold disclosure to all types is optimal.

8 DISCUSSION

8.1 *Posterior rent and privacy of signals*

As explained in the binary-type model, not observing signals generally hurts the seller due to the presence of the buyer's *posterior* rent. Specifically, implementing the benchmark allocation requires the seller to pay the buyer's *posterior* rent (apart from his *ex ante* rent), making $V(P) < V(\overline{P})$. When valuation shifts smoothly across (infinite) types, the relevance of signal privacy comes from a different reason. Indeed, any allocations *implementable with private signals* can be implemented without generating *posterior* rent to the buyer.[27] Therefore, if the seller fails to achieve the upper bound of revenue $V(\overline{P})$, it is due to an implementability issue. In such a scenario, information design can expand the set of implementable allocations. To illustrate, consider the following example where the benchmark allocation is implementable with private signals only if uninformative experiments are possible.

Example 6. $v(\theta,x)=\theta^2+\theta+x-2$. *Types and states are uniformly distributed over* $\Theta=[0,1]$ *and* $\Omega=[0,3]$.

In this example, $p^\star(\theta,s^g) = -\theta^2+\frac{2}{3}\theta+1$. Moreover, $p^\star(\theta,s^g)$ is a concave function in $[0,1]$ with $p(0,s^g)=1, p(1,s^g)=\frac{2}{3}$. Thus, $p^\star(\overline{\theta},s^g)=\min_\theta p^\star(\theta,s^g)$. Then by Proposition 2, the seller implements the benchmark allocation via $\mathbf{M}^\star$. Suppose the seller provides full disclosure to all types. To implement the benchmark allocation, it must be that for any θ and x, $q(\theta,x)=\mathbb{1}_{x\geq\omega^\star(\theta)}$. For the buyer to report truthfully their states, it is necessary that

$$p(\theta,x)=\begin{cases}\overline{p}(\theta) & \text{if } x\geq\omega^\star(\theta),\\ \underline{p}(\theta) & \text{otherwise.}\end{cases}$$

To prevent the lowest type $\underline{\theta}$ from mimicking some type θ and always report $x<\omega^\star(\theta)$, it must

[27]We omit the formal proof, which extends the arguments in Krähmer and Strausz (2015*a*) to a setting with information design and possibly finitely many states.

be that $\underline{p}(\theta) \geq 0$. Therefore,

$$\int_{x \geq x_\theta} \mu(x) dx p^\star(\theta, s^g) = \int_{x \geq x_\theta} \mu(x) dx \overline{p}(\theta) + \underline{p}(\theta) \int_{x \leq x_\theta} \mu(x) dx \geq \int_{x \geq x_\theta} \mu(x) dx \overline{p}(\theta)$$

where the equality uses the fact that all mechanisms implementing the benchmark allocation share the same expected payment. Thus, $p^\star(\theta) \geq \overline{p}(\theta)$ for all type θ.

Consider $\theta = \frac{1}{3}$, we have $\overline{p}^\star(\frac{1}{3}, s^g) = \frac{10}{9}$, and $v(\frac{1}{3}, x_{\frac{1}{3}}) = \frac{13}{9}$. Thus, $v(\frac{1}{3}, x(\frac{1}{3})) > p^\star(\frac{1}{3}, s^g) \geq \overline{p}(\frac{1}{3})$. Then, if the buyer observes any state $x \in (\overline{p}(\frac{1}{3}), v(\frac{1}{3}, x(\frac{1}{3})))$, it is optimal for him to misreport state $\overline{x}$, receiving the good at a price lower than his valuation. Thus, the benchmark allocation is not implementable under full disclosure.

8.2 Alternative proof for Wei and Green (2023)

Wei and Green (2023) revisit Eső and Szentes (2007)'s "continuous" model, adding a twist that the buyer can walk away after information disclosure. In this section, we solve the former's problem by directly modifying the latter's optimal mechanism.[28]

Under Eső and Szentes (2007)'s optimal mechanism, the seller offers full disclosure and a menu of "information fees" $\hat{c}(\cdot)$ and "strike prices" $\hat{p}(\cdot)$ for the good to implement the benchmark optimal allocation. Thus, $\big(q(\theta), p(\theta)\big) \in \Big\{\big(0, \hat{c}(\theta)\big), \big(1, \hat{c}(\theta) + \hat{p}(\theta)\big)\Big\}$. This menu is a deterministic mechanism. Therefore, following the arguments in the proof of Proposition 1, it is revenue-equivalent to a persuasive-posted price mechanism which offers type θ (i) a binary-signal experiment which sends "*good news*" if $x \geq x_\theta$ and "*bad news*" otherwise, and (ii) a posted price.

$$\tilde{p}(\theta) = \hat{c}(\theta) + \hat{p}(\theta) + \frac{\hat{c}(\theta)\,[1 - \mathscr{Q}(\theta)]}{\mathscr{Q}(\theta)} = \hat{p}(\theta) + \frac{\hat{c}(\theta)}{\mathscr{Q}(\theta)}.$$

In addition, Wei and Green (2023) show that information design leads to reverse price discrimination in the continuous model. This feature can also be obtained by leveraging the properties of Eső and Szentes (2007)'s optimal mechanism. Let $\mathbf{X}(\theta) \equiv \frac{1}{\mathscr{Q}(\theta)}$ represent the inverted trade probability for θ. Then, $\tilde{p}(\theta) = \hat{p}(\theta) + \hat{c}(\theta)\mathbf{X}(\theta)$, and

$$\tilde{p}'(\theta) = \hat{p}'(\theta) + \hat{c}'(\theta)\mathbf{X}(\theta) + \hat{c}(\theta)\mathbf{X}'(\theta) = \hat{c}(\theta)\mathbf{X}'(\theta) < 0,$$

where the second equality uses the fact that under Eső and Szentes (2007)'s optimal mechanism, $\hat{c}(\theta)$ and $\hat{c}(\theta)$ solves $\hat{c}'(\theta) = \hat{p}'(\theta)\mathscr{Q}(\theta) = \hat{p}'(\theta)\frac{1}{\mathbf{X}(\theta)}$, and the last uses $\mathbf{X}'(\theta) < 0$. Thus, $\tilde{p}(\cdot)$ is a decreasing function.

[28] Indeed, this modified mechanism coincides with Wei and Green (2023)'s solution.

8.3 On the number of signals

As we have seen, it is without loss of generality to offer binary-signal experiments with deterministic allocation. This is no longer true when random mechanisms are necessary. When the variations vary significantly across states, a rich menu is needed to screen the states effectively. As a result, binary-signal experiments are not sufficient. In this section, we illustrate this with a simple example where an optimal experiment sends at least three signals to some type.

Example 7. $\Theta = \{t_3, \theta_2, \theta_1\}, X = \{x_1, x_2, x_3, x_4\}$. *Types and states are equally likely.*

$v(\theta, x)$	x_1	x_2	x_3	x_4
θ_3	7	7	7	7
θ_2	0	3	7	7
θ_1	0	0	0	6

In this example, θ_2's valuation varies significantly across states with that at state x_1 being sufficiently low. If restricted to binary-signal experiments, the seller can only separate the state space for type θ_2 into two partitions which, under the optimal mechanism, include $\{x_1, x_2\}$ and $\{x_3, x_4\}$. Armed with three signals, the seller can distinguish a very unfavorable state x_1 from a better one x, fine-tuning the design of allocations. The formal proof is in the Online Appendix.

A PRELIMINARY RESULTS: OMITTED PROOFS

A.1 Proof of Lemma 1

Let $\mathbf{M} \equiv \{\pi_\theta, q(\theta, s), p(\theta, s)\}_{\theta,s}$ be an optimal mechansim. Toward a contradiction, assume that there exists x such that $q(\overline{\theta}, x) < 1$. Then, the seller can improve her revenue by revising type $\overline{\theta}$ contract to $\tilde{\mathbf{C}} \equiv \{\tilde{\pi}_{\overline{\theta}}, \tilde{p}(\overline{\theta})\}$ in which $\tilde{\pi}_{\overline{\theta}}$ provides no information and $\tilde{p}(\overline{\theta})$ is a posted price for the good, given by:

$$\tilde{p}(\overline{\theta}) = \sum_s p(\overline{\theta}, s)\pi_{\overline{\theta}}(s) + \mathbb{E}[v(\theta, x)] - \sum_x v(\overline{\theta}, x) \sum_s q(\overline{\theta}, s)\pi_{\overline{\theta}}(s|x)g(x)$$

To see this, note the following. If type $\overline{\theta}$ buys the good at the price $\tilde{p}(\overline{\theta})$ and no disclosure, he obtains:

$$\mathbb{E}[v(\overline{\theta}, x)] - \tilde{p}(\overline{\theta}) = \sum_x v(\overline{\theta}, x) \sum_s \left[q(\overline{\theta}, s) - p(\overline{\theta}, s)\right]\pi_{\overline{\theta}}(s|x)g(x),$$

which is equal to that under the original mechanism **M**. Moreover, if type θ mimics $\overline{\theta}$, he either does not buy the good to get a zero payoff or buys the good, to obtain

$$
\begin{aligned}
&\mathbb{E}[v(\theta,x)] - \tilde{p}(\overline{\theta}) \\
=&\mathbb{E}[v(\theta,x)] - \mathbb{E}[v(\overline{\theta},x)] + \sum_x\sum_s [v(\overline{\theta},x)q(\overline{\theta},s)\pi_{\overline{\theta}}(s|x)g(x) - \sum_x\sum_s p(\overline{\theta},s)\pi_{\overline{\theta}}(s|x)g(x) \\
=&\mathbb{E}[v(\theta,x)] - \sum_x\sum_s v(\overline{\theta},x)[1-q(\overline{\theta},s)]\pi_{\overline{\theta}}(s|x)g(x) - \sum_x\sum_s p(\overline{\theta},s)\pi_{\overline{\theta}}(s|x)g(x) \\
\leq&\mathbb{E}[v(\theta,x)] - \sum_x\sum_s v(\theta,x)[1-q(\overline{\theta},s)]\pi_{\overline{\theta}}(s|x)g(x) - \sum_x\sum_s p(\overline{\theta},s)\pi_{\overline{\theta}}(s|x)g(x) \\
=&\sum_x\sum_s \left[v(\theta,x)q(\overline{\theta},s) - p(\overline{\theta},s)\right]\pi_{\overline{\theta}}(s|x)g(x),
\end{aligned}
$$

which is exactly type θ's from mimicking $\overline{\theta}$ and report signals truthfully under the **M**. Therefore, $\tilde{\mathbf{C}}$ weakly increases type $\overline{\theta}$'s on-path payoff and weakly reduces the other types' off-path payoff. Consequently, the buyer reveals his true type. While payments by the other types remain unchanged, type $\overline{\theta}$ now pays

$$
\begin{aligned}
\tilde{p}(\overline{\theta}) &= \sum_s p(\overline{\theta},s)\pi_{\overline{\theta}}(s) + \mathbb{E}[v(\theta,x)] - \sum_x v(\overline{\theta},x)\sum_s q(\overline{\theta},s)\pi_{\overline{\theta}}(s|x)g(x) \\
&> \sum_s p(\overline{\theta},s)\pi_{\overline{\theta}}(s),
\end{aligned}
$$

which his expected payment under **M**. This contradicts with **M**. being optimal.

A.2 Proof of Proposition 1

We complete the arguments in the main text by showing that under Case 2: $\underline{p}(\theta) > 0$, the seller's revenue remains weakly higher under the revised mechanism. Fix type θ. if θ mimics type θ', he receives weakly less information and pays weakly higher for each action (buy or do not buy the good). Thus, his off-path payoff is weakly lower under the revised mechanism. On the other hand, by revealing his type and buying the good if and only if s^g is realized, he obtains:

$$
\tilde{U}(\theta) \equiv \sum_x [v(\theta,x) - \tilde{p}(\theta)]\mathcal{Q}(\theta,x)g(x) = \sum_x [v(\theta,x) - \overline{p}(\theta)]\mathcal{Q}(\theta,x)g(x) - \underline{p}(\theta)[1-\mathcal{Q}(\theta)], \tag{1.9}
$$

which is his payoff under the original mechanism $\mathbf{M}^d$. As $\mathbf{M}^d$ induces truth-telling, it follows that θ finds it optimal to reveal his type under the revised mechanism. Moreover, by (1.9) and the fact that the buyer's payoff is non-negative under $\mathbf{M}^d$, $\tilde{U}(\theta) \geq 0$. As a result, the buyer's payoff from buying the good a upon observing s^g, given by $\frac{\tilde{U}(\theta)}{\mathcal{Q}(\theta)}$, is non-negative. Hence, the buyer buys the good after s^g, and accordingly, pays the seller

$$
\tilde{p}(\theta)\mathcal{Q}(\theta) = \overline{p}(\theta)\mathcal{Q}(\theta) + \underline{p}(\theta)[1-\mathcal{Q}(\theta)], \tag{1.10}
$$

which is exactly type θ's expected payment under $\mathbf{M}^d$. Thus, the seller's revenue cannot decrease under the revised mechanism.

A.3 Proof of Proposition 2

Part (a): Given that $\phi(\theta, x_\theta) > 0\ \forall\theta$, the optimal allocation in $(\overline{\mathscr{P}})$ is uniquely given by $\mathbb{Q}(\theta, x) = \mathbb{1}_{x \geq x_\theta}$. Therefore, if there exists $\mathbf{M}$ such that $R_{\mathbf{M}} = V(\overline{\mathscr{P}})$, (i) $\mathbf{M}$ must be deterministic and (ii), upon observing any signal, type θ knows whether $x \geq x_\theta$ or not. Because of (i) and the fact that the buyer pays the same expected payment under $\mathbf{M}$ and $\mathbf{M}^\star$, the buyer pays more to get the good under $\mathbf{M}^\star$. Because of (ii), the buyer receives weakly less information under $\mathbf{M}^\star$. To sum up, the buyer pays more to get the good and gets weakly less information under $\mathbf{M}^\star$. Therefore, if $\mathbf{M}$ is incentive compatible with private signals, so is $\mathbf{M}^\star$. Thus, $R_{\mathbf{M}^\star} = V(\overline{\mathscr{P}})$.

Part (b): *"If":* Suppose $p^\star(\overline{\theta}, s^g) = \min_\theta\{p^\star(\theta, s^g)\}$, we now show that $\mathbf{M}^\star$ induces truth-telling even if the seller does not observe signals. Note that $\mathbf{M}^\star$, as a solution to $(\overline{\mathscr{P}})$, induces truth-telling with public signals. Therefore, it suffices to show that under $\mathbf{M}^\star$, for any type θ and θ', it is not beneficial for θ to report θ' and then either (i) always report s^b, (iii) always report s^g or (ii) always misreporting signals. By always reporting s^b off-path, θ obtains a zero payoff; hence, (i) is not beneficial. If (ii) is beneficial, then θ also benefits from mimicking $\overline{\theta}$ and reporting signals truthfully (type $\overline{\theta}$ always observes s^g and pays the least after s^g), which contradicts $\mathbf{M}^\star$ being incentive compatible with public signals. Now, consider the last deviating behavior. Note that if type θ, who reports θ', prefers to misreport s^g (buys the good) rather than truthfully s^b (and gets a zero payoff), it must be optimal for him to buy the good (or report s^g upon observing this signal. Hence, if (ii) is beneficial, so is (iii), a contradiction,

Part (b): *"Only If":* Suppose $\exists\theta$ such that $p(\overline{\theta}, s^g) > p(\theta, s^g)$. By mimicking θ and always reporting s^g, type $\overline{\theta}$ always gets the good at a lower price $p(\theta, s^g)$. Therefore, θ prefers to misreport θ than truth-telling. Consequently, if the seller offers $\mathbf{M}^\star$ with private signals, she obtains $R_{\mathbf{M}^\star} < V(\overline{\mathscr{P}})$.

B SCREENING VS. BUNCHING: OMITTED PROOFS

B.1 Proof of Claim 1

Let

$$\alpha_l \equiv \max\left\{x' \mid x \leq \hat{x}_l^\star : \mathbb{E}[v(h,x) \mid x > x'] < \mathbb{E}[v(l,x) \mid x > x']\right\},$$
$$\beta_l \equiv \min\left\{x' \mid x \geq \hat{x}_l^\star : \mathbb{E}[v(h,x) \mid x < x'] < \mathbb{E}[v(l,x) \mid x > x']\right\}.$$

If $x^{\star\star}(l) \in [\underline{x}, \alpha_l]$, the seller can do strictly better by offering a threshold disclosure $\tilde{\pi}(l)$ under which (i) the threshold is α_l^+ and (ii) with probability λ, "*good news*" is sent at α_l^+, such that such that $\omega^{\tilde{\pi}}(h, s^g) = \omega^{\tilde{\pi}}(h, s^b)$. Note that λ exists because by definition of α_l ,

$$\mathbb{E}[v(h,x) \mid x > \alpha_l^+] > \mathbb{E}[v(L,x) \mid x > \alpha_l^+],$$
$$\mathbb{E}[v(h,x) \mid x > \alpha_l] < \mathbb{E}[v(l,x) \mid x > \alpha_l].$$

Thus, it must be that $x^{\star\star}(l) > \alpha_l$. By similar arguments, we also have $x^{\star\star}(l) < \beta_l$. Thus, $x^{\star\star}(l) \in (\alpha_l, \beta_l)$. Hence,

$$\mathbb{E}[v(h,x) \mid x > x^{\star\star}(l)] > \mathbb{E}[v(l,x) \mid x > x^{\star\star}(l)],$$

which implies $\omega^{\pi^{\star\star}(l)}(h, s^b) \geq \omega^{\pi^{\star\star}(l)}(l, s^g)$.

B.2 Proof of Theorem 2

The proof leverages Lemma 10 below, which provides two expressions of the price gap between two adjacent types under $\mathbf{M}^\star$.

Lemma 10. *There exist positive functions $\lambda(\theta)$ and $\beta(\theta)$ such that:*

(a) $p^\star(\theta^+, s^g) - p^\star(\theta, s^g) = \Big[\mathbb{E}[v(\theta^+, x) \mid x_{\theta^+} \leq x <] - p^\star(\theta, s^g)\Big]\lambda(\theta), \; \forall\, \theta \geq \underline{\theta}.$

(b) $p^\star(\theta, s^g) - p^\star(\theta^-, s^g) = \Big[\mathbb{E}[v(\theta, \omega) \mid x_\theta \leq x < x_{\theta^-}] - p^\star(\theta, s^g)\Big]\beta(\theta), \quad \forall\, \theta \geq \underline{\theta}^+.$

Proof of Lemma 10. To examine the ranking of $p^\star(\cdot)$, we start with the expected payment $\mathbb{P}(\theta) = p^\star(\theta)\sum_{x \geq x_\theta} \mu(x)$. By its definition (see equations (1.5)),

$$\mathbb{P}(\theta^+) - \mathbb{P}(\theta) = \sum_{x_{\theta^+} \leq x < x_\theta} v(\theta^+, x)\mu(x) \quad \forall \theta \geq \underline{\theta},$$
$$\mathbb{P}(\underline{\theta}) = \sum_{x_{\underline{\theta}} \leq x < x_{\theta^-}} v(\underline{\theta}, x)\mu(x).$$

Using

$$\sum_{x_{\theta^+} \leq x < x_\theta} v(\theta^+, x)\mu(x) = \mathbb{P}^\star(\theta^+) - \mathbb{P}^\star(\theta) = p^\star(\theta^+, s^g)\sum_{x \geq x_{\theta^+}} \mu(x) - p^\star(\theta, s^g)\sum_{x \geq x_\theta} \mu(x) \qquad (1.11)$$

Part (a). Write the RHS of (1.11) as $\sum_{x \geq x_{\theta^+}} \mu(x)[p^\star(\theta^+, s^g) - p^\star(\theta, s^g)] + p^\star(\theta, s^g) \sum_{x_{\theta^+} \leq x < x_\theta} \mu(x).$

Then, we obtain

$$p^\star(\theta^+, s^g) - p^\star(\theta, s^g) = \frac{\sum\limits_{x_{\theta^+} \le x < x_\theta} v(\theta^+, x)\mu(x) - p^\star(\theta), s^g) \sum\limits_{x_{\theta^+} \le x < x_\theta} \mu(x)}{\sum\limits_{x \ge x_{\theta^+}} \mu(x)}$$
$$= \left[\mathbb{E}[v(\theta^+, x) \mid x_{\theta^+} \le x < x_\theta] - p^\star(\theta, s^g)\right] \frac{\sum\limits_{x_{\theta^+} \le x < x_\theta} \mu(x)}{\sum\limits_{x \ge x_{\theta^+}} \mu(x)}$$
$$\propto \mathbb{E}[v(\theta^+, x) \mid x_{\theta^+} \le x < x_\theta] - p^\star(\theta, s^g).$$

Part (b). Write the RHS of (1.11) as $\sum\limits_{x \ge x_\theta} \mu(x)[p^\star(\theta^+, s^g) - p^\star(\theta, s^g)] + p^\star(\theta^+, s^g) \sum\limits_{x_{\theta^+} \le x < x_\theta} \mu(x)$, and the rest followed by similar arguments. □

Armed with Lemma 10, we now show that the highest type pays the lowest price under $\mathbf{M}^\star$. It follows from Lemma 10 that if for all $\theta \ge \underline{\theta}$,

$$\mathbb{E}[v(\theta^+, x) \mid x(\theta^+) \le x < x_\theta] \le \mathbb{E}[v(\theta, x) \mid x_\theta \le x < x_{\theta^-}], \tag{1.12}$$

then the sign of $[p^\star(\theta^+, s^g) - p^\star(\theta, s^g)]$ is decreasing in θ. Moreover, this sign is non-positive because by (6.2.1) for $\underline{\theta}$,

$$\mathbb{E}[v(\underline{\theta}^+, x) \mid x(\theta^+) \le x < x(\underline{\theta})] \le \mathbb{E}[v(\underline{\theta}, x) \mid x(\underline{\theta}) \le x < x(\underline{\theta}^-)] = p^\star(\underline{\theta}, s^g),$$

implying $p^\star(\underline{\theta}^+, s^g) - p^\star(\underline{\theta}, s^g) \le 0$, by part (1) of Lemma 10. Therefore,

$$p^\star(\theta^+, s^g) - p^\star(\theta, s^g) \le 0 \quad \forall \theta \ge \underline{\theta}. \tag{1.13}$$

This implies that $p^\star(\theta, s^g)$ is the lowest price. Then by Proposition 2, $R_{\mathbf{M}^\star} = V(\overline{P})$. Moreover, as $\mathbf{M}^\star$ induces truth-telling with private signals, the seller can simply offer a menu of posted prices and threshold disclosure $\{\pi^\star_\theta, p^\star(\theta)\}_\theta$, where $p^\star(\theta) = p^\star(\theta, s^g)$.

B.3 Proof of Theorem 3

Let L be the lowest type being served under an optimal mechanism. Consider the following relaxed problem ($\mathscr{RP}_L$), under which all types mimics L off-path:

$$
\begin{aligned}
(\mathscr{RP}_L) \quad \max_{(\pi,q,U)} \quad & \sum_{\theta \geq L} \Big[\sum_x \sum_s v(\theta,x) q(\theta, \omega^{\pi_\theta}(\theta,s)) \pi_\theta(s|x)\mu(x) - U(\theta) \Big] f(\theta) \\
s.t. \quad & U(\theta) - U(L) \geq \sum_s \int_{\omega^{\pi_L}(L,s)}^{\omega^{\pi_L}(\theta,s)} q(L,z) dz \pi_L(s) \quad \forall \theta > L \qquad (IC) \\
& U(L) \geq 0 \qquad (IR_L) \\
& q(\theta,\omega) \text{ increases in } \omega. \qquad \text{(MON)}
\end{aligned}
$$

We will show that the solution to this relaxed problem, which features a posted price and a threshold disclosure, solves the original problem. Obviously, (IR_L) and $(IC_{\theta \to L})$ bind for all $\theta > L$ under ($\mathscr{RP}_L$), reducing the seller's relaxed problem to

$$
\begin{aligned}
\max_{q,\gamma} \quad & \sum_\theta \sum_x v(\theta,x) q(\theta,\omega^{\pi_\theta}(\theta,s))\pi_\theta(s|x)\mu(x) f(\theta) - \sum_\theta \sum_x \sum_s \int_{\omega^{\pi_L}(L,s)}^{\omega^{\pi_L}(\theta,s)} q(L,z) dz \pi_L(s|x)\mu(x) f(\theta) \\
s.t. \quad & q(\theta,\omega) \text{ increases in } \omega.
\end{aligned}
$$

Fix π, it is a linear problem in q with (MON) being the only constraint. Thus, the optimal allocation is generally unique, given by

$$
q(L,\omega) = \mathbb{1}_{s \geq \hat{s}_l}, \quad q(\theta,x) = 1 \forall x \forall \theta > L.
$$

Fix $q(L,s) = \mathbb{1}_{s \geq \hat{s}_l}$. The term involving π_L in the seller's objective (revenue) is given by

$$
\begin{aligned}
\mathbf{R}(\gamma_L) \equiv & \sum_x \sum_{\hat{s}_L}^{\bar{s}} v(L,x)\pi_L(s|x) f(L) g(x) \\
& - \sum_{\theta \geq L^+} \Big[\sum_{\hat{s}_L}^{\bar{s}} \big[\omega^{\pi_L}(\theta,s) - \omega^{\pi_L}(L,\hat{s}_L)\big] - \sum_{\theta \geq L^+} \sum_{\underline{s}}^{\hat{s}_L} \max\big\{\omega^{\pi_L}(\theta,s) - \omega^{\pi_L}(L,\hat{s}_L), 0\big\} \Big] \pi_L(s|x) f(\theta) \\
= & \sum_x \sum_{\hat{s}_L}^{\bar{s}} v(L,x)\pi_L(s|x) f(L) g(x) \\
& - \sum_{\theta \geq L^+} \Big[\sum_{\hat{s}_L}^{\bar{s}} \big[\omega^{\pi_L}(\theta,s) - \omega^{\pi_L}(L,\hat{s}_L)\big] - \sum_{\theta \geq L^+} \sum_{\underline{s}}^{\hat{s}_L} [\omega^{\pi_L}(\theta,s) - \omega^{\pi_L}(L,\hat{s}_L), 0] \Big] \pi_L(s|x) f(\theta) \\
= & \sum_x \sum_{\hat{s}_L}^{\bar{s}} v(L,x)\pi_L(s|x) f(L) g(x) - \sum_{\theta \geq L^+} \sum_{\underline{s}}^{\bar{s}} \big[\omega^{\pi_L}(\theta,s) - \omega^{\pi_L}(L,\hat{s}_L)\big] \pi_L(s|x) f(\theta) \\
\equiv & \overline{\mathbf{R}}(\gamma_L).
\end{aligned}
$$

$\overline{\mathbf{R}}(\pi_L)$ is an upper bound of $\overline{\mathbf{R}}(\pi_L)$. We now show that this bound is tight. By replacing all signals $s \geq \hat{s}_L$ with "*good news*" and all signals $s < \hat{s}_L$ with "*bad news*," $\overline{\mathbf{R}}(\gamma_L)$ weakly increases, and $\overline{\mathbf{R}}(\gamma_L)$ reduces to

$$\overline{\mathbf{R}}(\gamma_L) = \omega^{\pi_L}(L, s^g)\Big[\sum_{\theta \geq L^+} f(\theta) + f(L)\pi_L(s^g)\Big] - \sum_{\theta \geq L^+} \mathbb{E}[v(\theta, x)] f(\theta),$$

Let

$$\pi_L^{\star\star} \equiv \underset{\pi_L}{\operatorname{argmax}}\, \omega^{\pi_L}(L, s^g)\Big[\sum_{\theta \geq L^+} f(\theta) + f(L)\pi_L(s^g)\Big] - \sum_{\theta \geq L^+} \mathbb{E}[v(\theta, x)] f(\theta)$$

By the same arguments used for the binary-type case, $\pi_L^{\star\star}$ features a disclosure rule. Next, we find an optimal payment schedule. As optimal allocation is deterministic, without loss of generality to focus on posted-price mechanisms. By (IR_L), $p^{\star\star}(L) = \omega^{\pi_L^{\star\star}}(L, s^g)$. Consider type θ who mimics L. Under no uniformly no threshold flip of type order, $\omega^{\pi_L^{\star\star}}(\theta, s^g) \geq \omega^{\pi_L^{\star\star}}(\theta, s^g) \geq \omega^{\pi_L^{\star\star}}(L, s^g)$. Hence, it is optimal for θ to always buy the good after mimicking L. Hence, by (IR) $p^{\star\star}(\theta) = p^{\star\star}(L)$ for all $\theta > L$. Obviously, this single option of price and information $\{\pi_L^{\star\star}, p^{\star\star}(L)\}$ satisfies ignored constraints and hence, solves the original problem.

Remark 2. *Let $V\big(\mathscr{RP}_\theta\big)$ denote the value of program $\mathscr{RP}_\theta$ in which θ is the lowest type being served. Under no threshold flip of type order, it is optimal to serve only types above (including) L, where L solves $L \in \operatorname{argmax}_\theta V\big(\mathscr{RP}_2(\theta)\big)$.*

C RANDOM MECHANISMS: OMITTED PROOFS

This section provides proofs of Claims 2 and 3, which rely on the following lemma.

Lemma 11. *If only deterministic mechanisms are allowed, it is optimal to offer a menu $\{\alpha(\theta, x), p(\theta)\}_{\theta,x}$ under which each type θ trades with probability $\alpha(\theta, x)$ at state x and pays $p(\theta)$ for the good.*

The proof for Lemma 11 is as follows. If only deterministic mechanisms are allowed, by Proposition 1, it is optimal to offer each type θ receives a posted price $p(\theta)$ for the good and a binary-signal experiment with $S = \{s^g, s^b\}$. Hence, an experiment can be represented by the probability that signal s^g is realized at state x for type θ, $\alpha(\theta, x)$.

C.1 Proof of Claim 2

To characterize the optimal deterministic mechanism, or the optima menu $\{\alpha(\theta, x), p(\theta)\}_{\theta,x}$, consider the following relaxed problem in which (i) only IR condition for θ_1 is kept, and (ii) off the equilibrium path, θ_2 mimics θ_1 and buys the good only after s^g whereas θ_3 mimics θ_1 and always buys the good.

$$\max_{\{p,\alpha\}} \quad \sum_{\theta}\sum_{x} p(\theta)\alpha(\theta,x)\mu(x)f(\theta)$$

$$s.t. \quad \sum_{x}\left[v(\theta_2,x)-p(\theta_2)\right]\alpha(\theta_2,x)\mu(x) \geq \sum_{x}\left[v(\theta_2,x)-p(\theta_1)\right]\alpha(\theta_1,x)\mu(x) \qquad (IC_{21})$$

$$\sum_{x}\left[v(\theta_3,x)-p(\theta_3)\right] \geq \mathbb{E}[v(\theta_3,x)]-p(\theta_1) \qquad (IC_{21})$$

$$\sum_{x}\left[v(\theta_1,x)-p(\theta_1)\right]\alpha(\theta_1,x)\mu(x) \geq 0 \qquad (IR_1)$$

As trading with θ_3 and θ_2 generates no rent for others, these types receive efficient allocations: $\alpha(\theta_3,x_2)=\alpha(\theta_3,x_1)=1, \alpha(\theta_2,x)=\mathbb{1}_{x=g}$.

If $\alpha(\theta_1,x_1)>0$, then reduce $\alpha(\theta_1,x_1)$ and increase $p(\theta_1)$ such that $\sum_x p(\theta_1)\alpha(\theta_1,x)\mu(x)$ remains unchanged. By doing so, the seller's revenue increases. Moreover, no constraints are violated because (i) the right-hand side of (IC_{21}) decreases (as $v(\theta_3,x_1)>0$) and (ii) the right-hand side of (IC_{21}) and left-hand side of (IR_1) remains unchanged (as $v(\theta_2,x_1)=v(\theta_1,x_1)=0$).Thus, $\alpha(\theta_1,x_1)=0$.

If $\alpha(\theta_1,x_2)<1$, then reduces $\alpha(\theta_1,x_2)$ by ε and increases $p(\theta_3)$ by $[v(\theta_3,x_2)-p(\theta_2)]\varepsilon$ and $p(\theta_2)$ by $\frac{[v(\theta_3,x_2)-p(\theta_2)]\varepsilon}{\alpha(\theta_1,x_2)\mu(x_2)}$. By doing so, no constraint is affected while the seller's revenue increases by $[f(\theta_3)+f(\theta_2)][v(\theta_3,x_2)-p(\theta_2)]\mu(x_2)\varepsilon - f(\theta_1)v(\theta_1,x_2)\mu(x_2)\varepsilon = \phi(\theta_1,x_2)f(\theta_1)\mu(x_2)>0$. Thus, $\alpha(\theta_1,x_2)=1$.

If (IR_1) does not bind, increase $p(\theta_1)$ until it binds. This increases the seller's revenue while not violating any constraints. Thus, (IR_1) binds and hence, $p(\theta_1)=v(\theta_1,x_2)=4$.

If (IC_{31}) does not bind, increase $p(\theta_3)$ until it binds. This increases the seller's revenue while not violating any constraints. Thus, (IR_3) binds and hence, $p(\theta_3)=p(\theta_1)=4$.

If (IC_{21}) does not bind, increase $p(\theta_2)$ until it binds. This increases the seller's revenue while not violating any constraints. Thus, (IC_{21}) binds, or $p(\theta_2)=p(\theta_1)=4$.

To sum up, we obtain $\alpha(\theta_3,x_2)=\alpha(\theta_3,x_1)=1, \quad \alpha(\theta_2,x)=\alpha(\theta_1,x)=\mathbb{1}_{x=x_2}$, and $p(\theta_3)=p(\theta_2)=p(\theta_1)=5$. The seller's revenue is $4\cdot\left[f(\theta_3)+[f\theta_2)+f(\theta_1)]\mu(x_2)\right]=\frac{10}{3}$. As $p(\theta_3)=4$, type θ_3 buys the good at any state. Thus, the maximized revenue can be obtained via a posted price of 4 and full disclosure to all types.

C.2 *Proof of Claim* 3

To characterize the optimal deterministic mechanism, or the optima menu $\{\alpha(\theta,x), p(\theta)\}_{\theta,x}$, consider the following relaxed problem in which (i) only IR condition for θ_1 is kept, and (ii) off the equilibrium path, m mimics θ_1 and buys the good only after s^g whereas θ_3 mimics θ_1 and always buys the good.

$$\max_{\{p,\alpha\}} \quad \sum_\theta \sum_x p(\theta)\alpha(\theta,x)\mu(x)f(\theta)$$

$$s.t. \quad \sum_x \left[v(\theta_2,x) - p(\theta_3)\right]\alpha(\theta_2,x)\mu(x) \geq \mathbb{E}[v(\theta_3,x)] - p(\theta_2) \qquad (IC_{32})$$

$$\sum_x \left[v(\theta_2,x) - p(\theta_2)\right]\alpha(\theta_2,x)\mu(x) \geq \sum_x \left[v(\theta_2,x) - p(\theta_1)\right]\alpha(\theta_1,x)\mu(x) \qquad (IC_{21})$$

$$\sum_x \left[v(\theta_1,x) - p(\theta_1)\right]\alpha(\theta_1,x)\mu(x) \geq 0 \qquad (IR_l)$$

If $\alpha(\theta_1,x_1) > 0$, reduce $\alpha(\theta_2,b)$ and and increase $p(\theta_1)$ such that $p(\theta_1)\sum_x \alpha(\theta_1,x)\mu(x)$ remains unchanged, and increase $p(\theta_2)$ such that (IC_{21}) remains satisfied. By doing so, no constraints are affected, whereas the seller's revenue strictly increases. Thus, $\alpha(\theta_1,x_1) = 0$.

Note that to ensure that type θ_2's on-path payoff is non-negative, it is necessary that $v(\theta_2,x_2) \geq p(\theta_2)$. Hence, if $\alpha(\theta_2,x_2) < 1$, by increasing $\alpha(\theta_2,x_2)$, we strictly improve the seller's revenue while not violating any constraints. Thus $\alpha(\theta_2,x_2) = 1$.

If $\alpha(\theta_1,x_2) < 1$. Then, set $p(\theta_1) = v(\theta_1,x_2) = 4$, increase $\alpha(\theta_1,x_2)$ by ε and reduce $p(m)$ and $p(h)$ by $\frac{\varepsilon}{\sum_x \alpha(\theta_2,x)\mu(x)}$. Under this change, no constraints are violated. Moreover, the seller's revenue increases by $f(\theta_1)4\varepsilon - f(\theta_3)\frac{\varepsilon}{\sum_x \alpha(\theta_2,x)\mu(x)} - f(\theta_2)\varepsilon > 0$. Thus $\alpha(\theta_1,x_2) = 1$.

If (IR_1) does not bind, we can increase $p(\theta_1)$ up to it becoming binding, thereby increasing the seller's revenue without violating any constraints. Thus, (IR_1) binds. Given that $\alpha(\theta_1,x_1) = 0$, we thus have $p(\theta_1) = v(\theta_1,x_2)$.

If (IC_{32}) does not bind, increase $p(\theta_3)$ until it binds. This increases the seller's revenue while not violating any constraints. Thus, (IC_{21}) binds, and hence, $p(\theta_3) = p(\theta_2)$.

If (IC_{21}) does not bind, increase $p(\theta_2)$ until it binds. This increases the seller's revenue while not violating any constraints. Thus, (IC_{21}) binds. Given that $\alpha(\theta_2,x_2) = \alpha(\theta_1,x_2) = 1$, $\alpha(\theta_1,x_1) = 0$

and $p(\theta_1) = v(\theta_1, x_2)$, this implies

$$[v(\theta_2, x_1) - p(\theta_2)]\alpha(\theta_2, x_1)\mu(x_1) + [v(\theta_2, x_2) - p(\theta_2)]\mu(x_2) = [v(\theta_2, x_2) - v(\theta_1, x_2)]\mu(x_2)$$

$$\Leftrightarrow p(\theta_2) = \frac{v(\theta_2, x_1)\alpha(\theta_2, x_1)\mu(b) + v(\theta_2, x_2)\mu(x_2) - [v(\theta_2, x_2) - v(\theta_1, x_2)]\mu(x_2)}{\alpha(\theta_2, x_1)\mu(x_1) + \mu(x_2)}. \quad (1.14)$$

Then, the objective problem of the relaxed problem becomes

$$\Big[f(\theta_3) + f(\theta_2)[\alpha(\theta_2, x_1)\mu(x_1) + \mu(x_2)]\Big] \frac{v(\theta_2, x_1)\alpha(\theta_2, x_1)\mu(x_1) + v(\theta_2, x_2)\mu(x_2) - [v(\theta_2, x_2) - v(\theta_1, x_2)]\mu(x_2)}{\alpha(\theta_2, x_1)\mu(x_1) + \mu(x_2)}$$

$$\equiv H(\alpha(\theta_2, x_1))$$

Under the specification in Example 3(b), the relaxed problem becomes

$$\max_{\alpha(\theta_2, x_1)} H(\alpha(\theta_2, x_1)) \equiv (\alpha(\theta_2, x_1) + 3)\frac{(\alpha(\theta_2, x_1) + 2)}{\alpha(\theta_2, x_1) + 1} = \alpha(\theta_2, x_1) + 2 + \frac{2(\alpha(\theta_2, x_1) + 2)}{\alpha(\theta_2, x_1) + 1}$$

Thus, $H'(\alpha(\theta_2, x_1)) = 1 - \frac{1}{(\alpha(\theta_2, x_1)+1)^2}$ and $H"(\alpha(\theta_2, x_1)) = \frac{2}{(\alpha(\theta_2, x_1)+1)^3} > 0$. Therefore, $H(\alpha(\theta_2, x_1))$ is a convex function. Moreover, $R(0) = R(1) = 6$. Thus, $\alpha(\theta_2, x_1) = 0$ is optimal. This implies that $p(\theta_3) = p(\theta_2) = p(\theta_1) = 4$. Hence, a posted price $p = 4$ and full disclosure is an optimal deterministic mechanism.

D INFINITE TYPES: OMITTED PROOFS

D.1 Proof of for Proposition 3

We first solve the seller's benchmark problem with public signals when $\Theta = [\underline{\theta}, \overline{\theta}]$. With $\mathbb{P}(\theta)$ and $\mathbb{Q}(\theta, x)$ representing the expected payment and allocation over signals, this problem writes:

$$(\overline{P}^c) \quad \sup_{\mathbb{P},\mathbb{Q}} \quad \int_\theta \mathbb{P}(\theta) dF(\theta)$$

$$s.t. \quad \forall \theta, \theta' : \quad \sum_x v(\theta, x)\mathbb{Q}(\theta, x)\mu(x) - \mathbb{P}(\theta) \geq \sum_x v(\theta, x)\mathbb{Q}(\theta', x)\mu(x) - \mathbb{P}(\theta') \quad (1.15)$$

$$\sum_x v(\theta, x)\mathbb{Q}(\theta, x)\mu(x) - \mathbb{P}(\theta) \geq 0. \quad (1.16)$$

By the Envelope condition, (1.15) implies $U'(\theta) = \sum_x v_\theta(\theta, x)\mathbb{Q}(\theta, x)\mu(x) \; \forall \theta \geq \tilde{\theta}_g$. By integration by parts,

$$U(\theta) = U(\tilde{\theta}_g) + \int_{\tilde{\theta}_g}^{\theta} \sum_x v_\theta(\theta', x)\mathbb{Q}(\theta', x)\mu(x) d\theta'. \quad (1.17)$$

Consider a relaxed problem which keeps (2.2) and the partition constraint for the lowest type $\underline{\theta}$. Using $U(\underline{\theta}) = 0$ at optimum, this relaxed problem becomes:

$$\sup_{\pi,q} \int_\theta \phi^c(\theta,x)\mathbb{Q}(\theta,x)\mu(x)]\,dF(\theta),$$

where $\phi^c(\theta,x) \equiv v(\theta,x) - v_\theta(\theta,x)\frac{1-F(\theta)}{f(\theta)}$. As $\phi^c(\theta,x)$ increases in θ and x, it is optimal to se $\mathbb{Q}(\theta,x) = \mathbb{1}_{\theta\geq\theta_x}$ or equivalently, $\mathbb{Q}(\theta,x) = \mathbb{1}_{x\geq x_\theta}$. Fix an arbitrary $x \in X$. For any $\theta \in [\theta_x, \theta_{x^-}]$, $\mathbb{Q}(\theta) = \mathbb{Q}(\theta_x)$ and $\mathbb{P}(\theta) = \mathbb{P}(\theta_x)$. Payments are backed out using $U(\tilde{\theta}) = 0$ and (2.2), given by:

$$\mathbb{P}(\theta) = \sum_{x\geq x_\theta} v(\theta,x)\mu(x) - \int_\theta \sum_{x\geq x_{\theta'}} v_\theta(\theta',x)\mu(x)\,d\theta'$$

Moreover,

$$\int_{\theta_x}^{\theta_x^-} \sum_{x\geq x_{\theta'}} v_\theta(\theta',x)\mu(x)\,d\theta' = \int_{\theta_x}^{\theta_x^-} \sum_{x\geq x_{\theta_x}} v_\theta(\theta',x)\mu(x)\,d\theta' = \sum_{x\geq x_{\theta_x}} \int_{\theta_x}^{\theta_x^-} v_\theta(\theta',x)\mu(x) = \sum_{x\geq x_{\theta_x}} [v(\theta_x^-,x) - v(\theta_x,x)]\mu(x),$$

which implies

$$\mathbb{P}(\theta) = \sum_{x\geq x_\theta} v(\theta,x)\mu(x) - \sum_{\theta_x\leq\theta} \sum_{x\geq x_{\theta_x}} [v(\theta_{x^-},x) - v(\theta_x,x)]\mu(x).$$

Therefore, $\mathbf{M}^\star$ for this problem consists of Θ_x options $\{\pi_\theta^\star, p^\star(\theta,s^g), p^\star(\theta,s^b)\}_{\theta\in\Theta_x}$. Then, following similar arguments in the proof of Theorem 2 for the type space Θ_x, we get

$$p^\star(\theta_{\underline{x}},s^g) = \min_{\theta\in\Theta} p^\star(\theta,s^g)$$

under the partition flip of type order within Θ_x. As a result, $\mathbf{M}^\star$ induces truth-telling even if the seller does not observe signals. Then, offering $\mathbf{M}^\star$ is equivalent to offering a menu of threshold disclosures and posted prices $\{\pi_\theta^\star, p^\star(\theta)\}_{\theta\in\Theta_x}$, where $p^\star(\theta) = p^\star(\theta,s^g)$. Hence, this menu helps the seller achieve the upper bound of revenue $V(\overline{P}^c)$; hence, it is optimal.

D.2 Proof of for Proposition 4

Suppose it is optimal to exclude all types below L, or $q(\theta,x) = 1$ for all x and $\theta < L$. Then, the seller's revenue must be weakly lower than that obtained from selling to the buyer whose types is distributed by $\hat{f}$ over Θ, where $\hat{f}(\theta) = f(\theta)\ \forall\theta \notin [L, L+\varepsilon]$, $\hat{f}(\theta) = 0\ \forall\theta \in [L, L+\varepsilon)$, and $\hat{f}(L+\varepsilon) = \int_{\theta=L}^{\theta=L+\varepsilon} f(\theta)d\theta$. Let $(\hat{P})$ represent the seller's problem when $\theta \sim \hat{f}$ and $V(\hat{P})$ the corresponding value. Consider the following relaxed problem of $(\hat{P})$ where all types mimic $L+\varepsilon$ off the equilibrium path:

$$(\mathscr{RP}_{L+\varepsilon}) \quad \max_{(\pi,q,U)} \quad \sum_{\theta \geq L+\varepsilon} \sum_x \sum_s p(\theta, \omega^{\pi_\theta}(\theta,s))\pi_\theta(s|x)\mu(x)\hat{f}(\theta)$$

$$s.t. \quad U(\theta) - U(L+\varepsilon) \geq \sum_s \int_{\omega^{\pi_L}(L+\varepsilon,s)}^{\omega^{\pi_L}(\theta,s)} q(L+\varepsilon, z)dz\pi_{L+\varepsilon}(s) \quad \forall \theta > L+\varepsilon \quad (IC_{\theta \to L+\varepsilon})$$

$$U(L+\varepsilon) \geq 0 \quad (IR_{L+\varepsilon})$$

$$q(\theta,\omega) \text{ increases in } \omega. \quad \text{(MON)}$$

By the same arguments as the proof of Theorem 3, a posted price $\hat{p}_{L+\varepsilon}$, associated with a threshold disclosure $\hat{\pi}_{L+\varepsilon}$, solves this relaxed problem. Note that $(\hat{\pi}_{L+\varepsilon}, \hat{p}_{L+\varepsilon})$ does not necessary solve the original problem. In case it does, the seller's revenue is the value of problem $(\mathscr{RP}_{L+\varepsilon})$, denoted by $V((\mathscr{RP}_{L+\varepsilon}))$. Let R_ε represent the seller's revenue if she offers $(\hat{\pi}_{L+\varepsilon}, \hat{p}_{L+\varepsilon})$ (regardless of whether it solves the original problem or not). Then,

$$\begin{aligned} R_\varepsilon &\geq V(\mathscr{RP}_{L+\varepsilon}) - \mathbb{E}[v(L+2\varepsilon,x)] \int_{L+\varepsilon}^{L+2\varepsilon} \hat{f}(\theta)d\theta \\ &\geq V(\hat{P}) - \mathbb{E}[v(L+2\varepsilon,x)] \int_{L+\varepsilon}^{L+2\varepsilon} \hat{f}(\theta)d\theta \end{aligned}$$

Therefore,

$$\lim_{\varepsilon \to 0} R_\varepsilon \geq V(\hat{P}) - \lim_{\varepsilon \to 0} \mathbb{E}[v(L+2\varepsilon,x)] \int_{L+\varepsilon}^{L+2\varepsilon} \hat{f}(\theta)d\theta = V(\hat{P})$$

On the other hand, $\lim_{\varepsilon \to 0} R_\varepsilon \leq V(\hat{P})$. Therefore, $\lim_{\varepsilon \to 0} R_\varepsilon = V(\hat{P})$.

D.3 Proof of Proposition 5

Let $\mathbf{M}$ be an arbitrary optimal (deterministic) mechanism, which is, without loss of generality, a menu of trade probabilities and posted prices $\mathbf{M} = \{p(\theta), \alpha(x,\theta)\}_{x,\theta}$. Let $\tilde{\theta}_b \equiv \inf\{\theta \mid \alpha(\theta,x) = 1 \forall x\}$ represents the lowest type who receives an efficient allocation under $\mathbf{M}$, and $\tilde{\theta}_g \equiv \inf\{\theta \mid \alpha(\theta,x) > 0 \text{ for some } x\}$ be the lowest type being served. With $\tilde{\Theta} \equiv \{\theta \mid \theta \geq \tilde{\theta}_b\}$, $\mathbf{M}$ must solve the following problem:

$$(\mathscr{P}) \quad \sup_{p,\alpha} \quad \int_\theta p(\theta)dF(\theta)$$

$$s.t. \quad \alpha(\theta,x) = 1 \quad \forall x, \theta \geq \tilde{\theta}_b$$

$$\sum_x \left[v(\theta,x) - p(\theta)\right] \alpha(\theta,x)\mu(x) \geq \sum_x [v(\theta,x) - p(\theta')]\alpha(\theta',x)\mu(x) \quad \forall \theta,\theta' \in \tilde{\Theta}$$

$$\sum_x \left[v(\theta,x) - p(\theta)\right] \alpha(\theta,x)\mu(x) \geq 0 \quad \forall \theta \in \tilde{\Theta}.$$

By IR condition for $\tilde{\theta}_g$, $p(\tilde{\theta}_g) \le v(\tilde{\theta}_g, g)$. Consider $\theta \in [\tilde{\theta}_b, \overline{\theta}]$. If $p(\theta) > v(\tilde{\theta}_g, g)$, then θ prefers to mimic $\tilde{\theta}_g$ and always buy the good at a lower price. Thus, to incentivize θ to reveal his type, it must be that

$$p(\theta) \le v(\tilde{\theta}_g, g) \quad \forall \theta \in [\tilde{\theta}_b, \overline{\theta}].$$

Suppose $v(\tilde{\theta}_b, b) > v(\tilde{\theta}_g, g)$. Then, $\exists \hat{\theta}$ such that for any $\theta' \in [\hat{\theta}, \tilde{\theta}_b]$,

$$v(\theta', b) \ge v(\tilde{\theta}_g, g) \ge p(\overline{\theta}).$$

It then induces θ' to mimic $\overline{\theta}$ and always buy the good. By doing so, he gets the good a higher expected surplus at a lower price. Therefore, it must be that

$$v(\tilde{\theta}_b, b) \le v(\tilde{\theta}_g, g).$$

By Envelope condition, for the buyer to report truthfully his type, it is necessary that $U'(\theta) = \sum_x v_\theta(\theta, x)\alpha(\theta, x)\mu(x)$. Then, by integration by parts,

$$U(\theta) = U(\tilde{\theta}_g) + \int_{\tilde{\theta}_g}^{\overline{\theta}} \sum_x v_\theta(\theta, x)\alpha(\theta, x)\mu(x)\, d\theta. \tag{1.18}$$

Consider a relaxed problem that keeps only the IR condition for $\tilde{\theta}_g$ and the necessary envelope condition for truth-telling. Using (1.18) and the fact that $U(\tilde{\theta}_g) = 0$ at optimum, this relaxed problem reduces to

$$\sup_q \int_\theta \phi^c(\theta, x) q(\theta, x)\mu(x) f(\theta) \quad s.t. \quad q(\theta, x) = 1 \quad \forall x, \forall \theta \ge \tilde{\theta}_b,$$

where $\phi^c(\theta, x) \equiv v(\theta, x) - v_\theta(\theta, x)\frac{1-F(\theta)}{f(\theta)}$. Solving this point-wise maximization problem yields

$$q(\theta, x) = \begin{cases} 1 & \text{if} \quad \theta \ge \min\{\theta_b, \tilde{\theta}_b\} \\ \mathbb{1}_{x=g} & \text{if} \quad \max\{\theta_g, \tilde{\theta}_g\} \le \theta \le \min\{\theta_b, \tilde{\theta}_b\}. \end{cases}$$

Prices are pinned down using binding constraints, given by

$$p(\theta) = \begin{cases} v\left(\min\{\theta_b, \tilde{\theta}_b\}, b\right)\mu(b) + v\left(\max\{\theta_g, \tilde{\theta}_g\}, g\right)\mu(g) & \text{if} \quad \theta \ge \min\{\theta_b, \tilde{\theta}_b\} \\ v\left(\max\{\theta_g, \tilde{\theta}_g\}, g\right) & \text{if} \quad \max\{\theta_g, \tilde{\theta}_g\} \le \theta \le \min\{\theta_b, \tilde{\theta}_b\} \end{cases}$$

Consider $\theta \ge \min\{\theta_b, \tilde{\theta}_b\}$ and $\theta' \in \left[\max\{\theta_g, \tilde{\theta}_g\}, \min\{\theta_b, \tilde{\theta}_b\}\right]$. As $v(\tilde{\theta}_b, b) \le v(\tilde{\theta}_g, g)$, we have

$$v\left(\min\{\theta_b, \tilde{\theta}_b\}, b\right) \le v\left(\max\{\theta_g, \tilde{\theta}_g\}, g\right),$$

which implies $p(\theta) \leq p(\theta')$. Thus, this two-option menu of prices and threshold disclosure induces participation and truth-telling. As **M** solves the original problem, by definition of $\tilde{\theta}_b$ and $\tilde{\theta}_g$, it must be that

$$\min\{\theta_b, \tilde{\theta}_b\} = \tilde{\theta}_b, \quad \max\{\theta_g, \tilde{\theta}_g\} = \tilde{\theta}_g.$$

Suppose $v\left(\min\{\theta_b, \tilde{\theta}_b\}, b\right) < v\left(\max\{\theta_g, \tilde{\theta}_g\}, g\right)$, then $p(\theta) < p(\theta')$. Then, it is optimal to set $\hat{\theta} \equiv \inf\{\theta \geq \theta_g \mid v\left(\min\{\theta_b, \tilde{\theta}_b\}, b\right) \leq v(\max\{\hat{\theta}, \tilde{\theta}_g\}, g)\}$ as the lowest type being served, a contradiction. Therefore,

$$v\left(\min\{\theta_b, \tilde{\theta}_b\}, b\right) = v\left(\max\{\theta_g, \tilde{\theta}_g\}, g\right).$$

This implies that all types receive the same price. Moreover, all types $\theta \geq \min\{\theta_b, \tilde{\theta}_b\}$ always buy the good regardless of signal realization. Thus, it is optimal to offer full disclosure for all types.

Chapter 2

Buyer's Optimism, Information Design, and Price Discrimination

1 INTRODUCTION

In markets for experience goods, consumers do not know their willingness to pay prior to consumption, especially if they are newcomers to the market. However, they may form their own beliefs regarding possible payoffs before making purchasing decisions. This process can potentially induce diversity in consumers' beliefs. An optimistic buyer, who receives (possibly biased) good reviews, thinks that his valuation is more likely to be high, whereas a pessimistic one assigns excessive weights to lower valuations. For example, a student who is about to buy a new iPad for study purposes might underrate the distractions that an iPad can cause to himself. Consequently, he is over-optimistic about his valuation for an iPad. On the other hand, there are many ways by which the seller can provide additional information that helps the buyer correct his belief. For instance, software suppliers usually offer trial versions of the product to their consumers. They can also just provide product guides or some kind of informative advertising. Hence, the seller's strategies include not only pricing but also information disclosure.

Some natural questions arise in such situations. How could the seller screen the buyer's degrees of optimism using price-information bundles? Will the seller practice price and/or information discrimination? Should the seller offer information free of charge? How does the presence of information design and/or biased priors shape the optimal selling mechanism?

This paper aims to answer these questions in a simple monopolistic screening setup. Formally, our model features a single buyer who faces uncertainty about the product's match value (assumed to be either high or low), about which the seller can design information. From the seller's perspective, the buyer holds a biased belief about the likelihood that the product matches with him. The buyer's degree of optimism is his private information (his type). The seller designs a menu of prices and disclosure policies to maximize her revenue.

Our main result is regarding the interaction of the buyer's optimism, information design, and price discrimination. First, we examine the two benchmarks without either biased priors or information design. The first benchmark assumes that the seller does not control information, and hence, fails to refine the buyer's belief. As a result, she designs prices based on the buyer's private (*ex ante*) valuation (calculated based on his biased prior). A posted price is optimal in such a screening setup. In the second benchmark, the seller can provide information but the buyer's belief is equal to the seller's. If this belief is publicly known, the seller can fully extract the surplus by charging a price equal to the buyer's (unconditional) expected valuation. In case the buyer's prior belief is his private information, we show that a posted price also maximizes the seller's revenue under a mild condition. Overall, neither the diversity in the (biased) priors nor information design suffices to trigger price discrimination.

We then establish that the optimal mechanism features a menu of information-price bundles under the *simultaneous* presence of information design and the buyer's biasedness. In such a scenario, the seller can "bet" with the buyer about his posterior valuation (after information disclosure). This creates a new source of revenue: the fictional surplus due to non-common priors, apart from the widely known virtual surplus. Information disclosure generates a positive fictional surplus by allowing for a high trading price with a relatively pessimistic buyer. It, by contrast, triggers a negative fictional surplus with a relatively optimistic buyer. However, it does not mean that an optimistic (*resp.*, pessimistic) buyer should receive no (*resp.*, full) information disclosure. What determines the optimal mechanism is not only the fictional surplus gain (trade at high state/high price) but also the virtual surplus loss (no trade at low state) due to information disclosure.

Balancing this trade-off, the optimal mechanism follows a cut-off structure: optimistic types (compared to the cut-off, not the seller's prior) receive no disclosure and always buy the good at a reduced price; whereas pessimistic types get full information and buy only if his valuation is high, at the cost of paying a relatively high price (that is exactly his true value). At the cut-off type, the (marginal) fictional surplus gain is equal to the (marginal) virtual surplus loss due to full disclosure.

This finding implies that information and price discrimination are beneficial in screening the buyer's degrees of optimism. The optimistic buyer, who cares more about price than information, picks the no information/ low price option, whereas the pessimistic one opts to learn the match value and buy at a higher price. Moreover, despite the non-traditional ingredients in our model such as non-common priors, the optimal mechanism features the well-known "no distortion at the top and no rent at the bottom" with sufficiently optimistic types receiving an efficient allocation and pessimistic ones not marginally gaining from full information.

Our model assumes that information is provided free of charge. It is then natural to ask whether charging information fees strictly improves the seller's revenue. We find that this is not the case. The key is that information fees are always paid regardless of signal realizations, making it separable and independent of the buyer's prior (his type). Then, in a similar vein to the revenue equivalence theorem, the seller's revenue can be determined entirely by the allocation and prices (of the good).[1] Therefore, as long as the optimal mechanism can be solved via a commonly used relaxed problem which replaces incentive compatibility (IC) constraints by its well-known envelope condition as in our binary-(match) value setting, offering free information is optimal. In Section 6.2, we provide an example where there are three possible values from consumption and information fees are necessary.

1.1 Related literature.

First, our paper is related to the literature in behavioral economics that studies optimal contracts under the presence of consumers' biases in estimating potential payoffs.[2] The previous literature, however, does not accommodate information design and instead, lets payoff uncertainty be resolved fully and exogenously. Similar to us, Grubb (2009) studies situations where consumers assign wrong weights to their possible *ex post* valuations. Relaxing the common prior assumption in Courty and Li (2000), he incorporates consumers' overconfidence whose prior narrowly concentrates around the mean and mainly focuses on characterizing the optimal contract under complete information. By contrast, our model features consumers who put too much weight on high/low valuations. Similar kinds of "optimism/pessimism" have been observed in Eliaz and Spiegler (2008) for a monopolistic screening model in which consumers assign excessive weights to the states of nature associated with their large gains from trade.

[1] Due to the correlation between the buyer's prior belief and the additional information, we cannot get rid of prices (of the good) in the buyer's marginal rent.

[2] See Eliaz and Spiegler (2008) for a detailed literature review.

They find that the diversity in consumers' degree of optimism is *necessary* for price discrimination. Instead, we focus on the *combination* of consumers' optimism and information design in activating price discrimination.

Second, in line with the Bayesian persuasion framework following Kamenica and Gentzkow (2011), we impose no restrictions on the seller's information structures. The most related work is **?**, which studies Bayesian persuasion with heterogeneous priors.[3] They establish a surprising result that even when the prior difference is in the direction that benefits the Sender, she may still prefer to disclose information. The disclosure rule under our optimal mechanism shares a similar spirit with, however, very different driving forces. In their *pure* persuasion model, information is valuable whenever it is possible to design a lottery where the Sender is more optimistic than the Receiver about more beneficial actions. In our *joint* price and information design problem, information disclosure also affects pricing and the buyer's rent (disclosure is private). The optimal disclosure rule (as part of the optimal mechanism) trades off the gains and losses induced for the fictional and virtual surplus.[4]

Finally, our paper is close to the literature on joint mechanism and information design. Early contributions include Lewis and Sappington (1994), Johnson and Myatt (2006), Eső and Szentes (2007). Recent papers allow for general information structures, (see, for example, Li and Shi (2017), Guo et al. (2022), Smolin (2023), Wei and Green (2023), **?** as the most relevant ones to this paper). All previous works employ the common prior assumption, completely shutting down the impact of information disclosure (and prices) via the fictional surplus channel. Wei and Green (2023) find that information design about a payoff-relevant state, which is *independent* with the buyer's private information (*e.g.*, a taste shock), activates reverse price discrimination. In our model, the buyer's private information is his (biased) prior belief and more importantly, non-common priors are, in many cases, necessary for price discrimination. Relatedly, Guo et al. (2022) considers the buyer holding a private but *unbiased* prior belief, establishing the optimality of interval disclosure with binary types and infinite values. Focusing on a binary-value (but infinite-type) setting in which a take-or-leave offer of price and information is optimal without the buyer's optimism, we shed light on how non-common priors shape the optimal selling mechanism. See Section 6.1 for a detailed discussion.

[3]See also the online appendix **?** for a multiple-agent Bayesian persuasion problem with heterogeneous priors. Guo and Shmaya (2019), in their discussion section, also study how their results extend to cases where the players share no common prior.

[4]I thank Odilon Camara for his suggestion on the connection with their paper.

Environment: A buyer considers whether to buy a product from a monopolist. His utility from consumption (valuation), denoted by v, is ex-ante unknown. Moreover, his valuation can be either high (H) or low (L), depending on whether the product matches his need. Formally, $v \in \{L, H\}$ with $0 \leq L < H$.[5]

From the seller's perspective, the buyer holds a private and biased belief about his valuation, modelled by his type $\theta \in [0,1]$. Specifically, the buyer of type θ thinks that the good fits him with probability θ. There is a continuum of types, distributed over the interval $\Theta = [0,1]$ by $F(\theta)$ that admits a density $f(\theta)$.

The seller knows that the product fits the buyer with a subjective probability θ_S. From her perspective, type $\theta > \theta_S$ is relatively optimistic with higher θ being more optimistic, whereas type $\theta < \theta_S$ is relatively pessimistic with lower θ being more pessimistic.

We impose the usual monotone hazard rate assumption.

Assumption 4 (Monotone hazard rate)**.** *At any* $\theta \in \Theta$, $\dfrac{f(\theta)}{1-F(\theta)}$ *increases in* θ.

Selling mechanism: The seller designs prices and additional information about the good that helps the buyer refine his belief *privately* (only the buyer observes the signal). Information is modelled using the concept of a statistical experiment $E \equiv (S, \pi)$ that consists of two parts

- A signal space S, and
- A likelihood function π that maps each state (valuation) to a distribution of signals: $\pi : v \to \Delta(S)$.

As the buyer's type is his prior belief, it correlates with the distribution of signals through the buyer's Bayesian updating process.

A selling mechanism is a menu of prices and experiments, denoted by $\{(p(\theta), E(\theta))_\theta\}$. For now, we assume that information is provided free of charge. We show that charging information fees does not improve the seller's revenue in Section 5.2.

Timing: The timing of interactions is as follows:

1. The seller offers a selling mechanism.

[5] The results extend to the case with $L < 0 < H$ under some mild conditions on the parameters.

2. The buyer learns his type θ and decides to accept or reject the offer. In case of acceptance, he reports a type $\hat{\theta}$.
3. The buyer receives his price-experiment bundle $\left(p(\hat{\theta}, E(\hat{\theta})\right)$.
4. The buyer *privately* observes a signal s generated from experiment $E(\hat{\theta})$ and decides whether to buy the good at price $p(\hat{\theta}$.

3 SELLER'S PROBLEM

Given that the buyer's action space is binary (either buying the good or not), Lemma 12 below shows that it is without loss of optimality to focus on the binary-signal experiments.

Lemma 12. *It is without loss of generality to restrict to binary-signal experiments where the signal space consists of two signals, "buy" and "not buy".*

The proof (omitted) is standard, following the "revelation principle" argument by which it is without loss of generality to assume that each signal represents a recommended action (See **?**).

It then follows from Lemma 12 that an experiment can be represented by trade probabilities, $\{q(\theta, v)\}_{\theta, v}$, with which signal "buy" is sent to each type θ at each state v. The seller's problem thus reduces to finding the optimal menu $\{p(\theta), q(\theta, H), q(\theta, L)\}_\theta$ where q satisfies the following feasibility condition:

$$\forall \theta: \quad 0 \leq q(\theta, H), q(\theta, L) \leq 1, \tag{FC}$$

and as is common in private persuasion problems, the seller faces two kinds of constraints. First, the obedience constraints (OB) ensure that the buyer, having reported his type truthfully, follows the recommended signals. Second, the truth-telling constraints (IC) incentivize the buyer to report his type truthfully.

Obedience constraints: To make the buyer obedient, his posterior valuation must be (i) higher than the price of the good after signal "buy", and (ii) lower than the price after signal "not buy". Formally, for type $\theta \in [0, 1]$,

$$\mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{"buy"}] \equiv \frac{\theta q(\theta, H) H + (1-\theta) q(\theta, L) L}{\theta q(\theta, H) + (1-\theta) q(\theta, L)} \geq p(\theta), \tag{OB^b}$$

$$\mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{"not buy"}] \equiv \frac{\theta \left[1 - q(\theta, H)\right] H + (1-\theta) \left[1 - q(\theta, L)\right] L}{\theta \left[1 - q(\theta, H)\right] + (1-\theta) \left[1 - q(\theta, L)\right]} \leq p(\theta). \tag{OB^{nb}}$$

where we use the shorthand $\mathbb{E}_\theta[.]$ to denote the expectation calculated based on type θ's prior: $\text{Prob}(v = H) = \theta$.

Truth-telling constraints: Note that the seller needs to handle double deviations where the buyer first lies about his type and then disobeys the signals. Let

$$\pi^f(\theta,\theta') \equiv \theta q(\theta',H)\left[H - p(\theta')\right] + (1-\theta) q(\theta',L)\left[L - p(\theta')\right],$$

represent the payoff of type θ buyer who reports θ' and then follows recommended signals. With abuse of notation, let $\pi^f(\theta) \equiv \pi^f(\theta,\theta)$.

If disobeying the signals, the buyer can either (i) *always buy* regardless of the signals and obtain $\mathbb{E}_\theta[v] - p(\theta')$, (ii) *never buy* regardless of the signals to obtain a zero payoff; or (iii) *do the opposite* of the signals and achieve $\theta\left[1-q(\theta',H)\right]\left[H-p(\theta')\right] + (1-\theta)\left[1-q(\theta',L)\right]\left[L-p(\theta')\right]$. Let $\pi^d(\theta,\theta')$ represent the payoff for type-θ buyer, who misreports θ' and then disobeys signals. Then,

$$\pi^d(\theta,\theta') \equiv \max\left\{\theta\left[1-q(\theta',H)\right]\left[H-p(\theta')\right] + (1-\theta)\left[1-q(\theta',L)\right]\left[L-p(\theta')\right], \mathbb{E}_\theta[v] - p(\theta'), 0\right\}.$$

IC constraints write:

$$\forall \theta,\theta': \quad \pi(\theta) \geq \max\left\{\pi^f(\theta,\theta'), \pi^d(\theta,\theta')\right\}, \tag{IC}$$

To sum up, the seller's problem is formally expressed as follows

$$\begin{aligned} \max_{p,q} \quad & \int_0^1 \left[\theta_S q(\theta,H) + (1-\theta_S) q(\theta,L)\right]\left[p(\theta) - c\right] dF(\theta) \\ s.t. \quad & (FC),\ (OB^b),\ (OB^{nb}),\ (IC). \end{aligned} \tag{P}$$

It is worth noting that the objective function is calculated based on the seller's prior, while constraints are formulated based on the buyer's.

4 TWO BENCHMARKS

In this section, we examine the two benchmarks in which (i) the buyer's belief is unbiased with $\theta_S = \theta$ and (ii) information design is not feasible. We show that in either problem, a take-it-or-leave-it offer is optimal.

Buyer is unbiased: First, we consider the case when the buyer's belief coincides with the seller's and hence, is no longer the buyer's private information. In this case, the seller fully extracts the surplus by charging a price equal to the buyer's expected value from buying the good. Formally:

Proposition 6 (No biased priors)**.** *Suppose* $\theta_S = \theta$ *and* θ *is commonly known by everyone. Then, it is optimal to offer no disclosure and a posted price* $p = \mathbb{E}_\theta[v] = \theta H + (1-\theta)L$*.*

A slightly more general case when the buyer's belief is "private and unbiased" is studied in Section 6.1 where we show that a fixed price, associated with full disclosure, is optimal under a mild condition.

No information design: Next, suppose that information design is not feasible. Without information control, the seller designs a menu of prices and trade probabilities $\{p(\theta), q(\theta)\}_\theta$ to screen the buyer's level of optimism. Because the seller cannot provide information to refine the buyer's belief, her revenue depends on the buyer's perspective. Moreover, given that the buyer is risk-neutral, it is without loss of generality to assume that the buyer's valuation is $v(\theta) = \theta H + (1-\theta)L = L + \theta(H-L)$. The seller's problem then becomes standard with the buyer privately knowing his valuation $v(\theta)$. Then, a posted price, which equals the ex-ante value of a cutoff type, is optimal.

Proposition 7 (No information design)**.** *It is optimal for the seller to offer a posted price* $p^\star_{bm} = \theta^\star_{bm} H + (1-\theta^\star_{bm})L$*, where* $\theta^\star_{bm}$ *is the cutoff type that solves:*

$$\theta^\star_{bm} = \inf\left\{\theta \mid L + \theta(H-L) - (H-L)\frac{1-F(\theta)}{f(\theta)} - c \geq 0\right\}. \tag{2.1}$$

Proof. Formally, the seller's problem can be written as:

$$\begin{aligned} \max_{p,q} \quad & R = \int_\theta (p(\theta)-c)q(\theta)\,dF(\theta) \\ \text{s.t.} \quad \forall \theta, \theta' : \quad & v(\theta)q(\theta) - p(\theta) \geq v(\theta)q(\theta') - p(\theta') \qquad (IC) \\ & v(\theta)q(\theta) - p(\theta) \geq 0. \qquad (IR) \end{aligned}$$

Using the standard technique in mechanism design (particularly the necessary envelope condition for truth-telling), we obtain the following relaxed problem of the seller's problem:

$$\max_q = \int_\theta \left\{v(\theta) - v'(\theta)\frac{1-F(\theta)}{f(\theta)} - c\right\} q(\theta)\,dF(\theta)$$

Under Assumption 4, the virtual surplus of type θ, given by $v(\theta) - v'(\theta)\frac{1-F(\theta)}{f(\theta)} - c = L + \theta(H-L) - (H-L)\frac{1-F(\theta)}{f(\theta)} - c$, increases in θ. Thus, the optimal allocation follows the following cutoff rule:

$$q(\theta) = \begin{cases} 1 & \text{if } \theta \geq \theta^\star_{bm} \\ 0 & \text{otherwise} \end{cases}, \quad \text{where } \theta^\star_{bm} = \inf\left\{\theta \mid L + \theta(H-L) - (H-L)\frac{1-F(\theta)}{f(\theta)} - c \geq 0\right\}.$$

This allocation can be implemented using a posted price: $p^\star_{bm} = v(\theta^\star_{bm}) \equiv \mathbb{E}_{\theta^\star_{bm}}[v]$. □

In this section, we show that the simultaneous presence of information design and the buyer's optimism induces the optimal mechanism featuring both information and price discrimination.

5.1 Optimal mechanism

To characterize the optimal mechanism, we solve a relaxed problem which considers only necessary conditions for truth-telling and obedience. First, the truth-telling condition (IC) implies that for any θ, θ',

$$\pi^f(\theta) \geq \pi^f(\theta, \theta'),$$
$$\Leftrightarrow \theta q(\theta, H)[H - p(\theta)] + (1-\theta) q(\theta, L)\left[L - p(\theta')\right] \geq \theta q(\theta', H)\left[H - p(\theta')\right] + (1-\theta) q(\theta', L)\left[L - p(\theta')\right].$$

Using the famous Envelope theorem, the marginal rent for the buyer is given by

$$[\pi^f(\theta)]' = q(\theta, H)\left[H - p(\theta)\right] - q(\theta, L)\left[L - p(\theta)\right]. \tag{2.2}$$

It is worth noting that this expression for the buyer's marginal rent differs from what is commonly seen. By the famous revenue equivalence theorem, it should be expressed fully by the allocation terms. Here, it also involves the payment term $p(\theta)$, which is technically due to the correlation between the *ex ante* type θ and *ex post* valuation v.

Using (2.2) and intergration by parts, we obtain $\int_0^1 \pi(\theta) dF(\theta) = \int_0^1 \pi'(\theta)\left[1 - F(\theta)\right] d\theta + \pi(0)$. Then, the seller's relaxed problem becomes

$$\max_{\{p,q\}} \int_0^1 \Bigg[\theta q(\theta, H)(H - c) + (1-\theta) q(\theta, L)(L - c) - \pi'(\theta) \frac{1 - F(\theta)}{f(\theta)}$$
$$+ (\theta_S - \theta)\left[q(\theta, H) - q(\theta, L)\right]\left[p(\theta) - c\right]\Bigg] dF(\theta) - \pi(0). \tag{2.3}$$

The following lemma shows that obedience constraints require that trade probabilities increase in values.

Lemma 13. *Under any mechanism that satisfies the obedience constraints, $q(\theta, L) \leq q(\theta, H)$ for any θ.*

Proof of Lemma 13. If only signal "buy" is sent, then $q(\theta, L) = q(\theta, H) = 1$. If only signal "not buy" is sent, then $q(\theta, L) = q(\theta, H) = 0$. If each signal is sent with a strictly positive probability, then by (OB^b) and (OB^{nb}),

$$\mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“buy"}] \geq p(\theta) \geq \mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“not buy"}],$$

which implies

$$\begin{aligned} & \mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“buy"}] \geq \mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“not buy"}] \\ \Leftrightarrow\ & \frac{\theta q(\theta, H) H + (1-\theta) q(\theta, L) L}{\theta q(\theta, H) + (1-\theta) q(\theta, L)} \geq \mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“not buy"}] \\ \Leftrightarrow\ & \frac{\theta \left[q(\theta, H) - q(\theta, L)\right]}{q(\theta, L)} \geq \frac{\mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“not buy"}] - L}{H - \mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“not buy"}]}. \end{aligned} \tag{2.4}$$

Because $H \geq \mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“buy"}]$ and $\mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“not buy"}] \geq L$, the right-hand side of (2.4) is non-negative. Hence, the left-hand side of (2.4) must be non-negative, which means $q(\theta, H) \geq q(\theta, L)$. □

The next lemma establishes upper and lower bounds on the price for the good.

Lemma 14. *Under any mechanism that satisfies the obedience constraints with $q(\theta, L) < q(\theta, H) \leq 1$, then $L \leq p(\theta) \leq H$ for any θ.*

The proof is straightforward. As $q(\theta, L) < 1$, signal "not buy" is sent with a strictly positive probability. By (OB^{nb}), $p(\theta) \geq \mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“not buy"}] \geq L$. In addition, as $q(\theta, H) > 0$, signal "buy" is sent with a strictly positive probability. By (OB^b), $p(\theta) \leq \mathbb{E}_\theta[v \mid \hat{\theta} = \theta, \text{“buy"}] \leq H$.

Using the new objective function (2.3), lemmas 13 and 14, and the fact that $\pi(0) = 0$ at optimum, we obtain the following relaxed problem:

$$\begin{aligned} \max_{\{p,q\}} \quad & \int_0^1 \Bigg[\theta q(\theta, H)(H-c) + (1-\theta) q(\theta, L)(L-c) - \pi'(\theta) \frac{1-F(\theta)}{f(\theta)} \\ & + (\theta_S - \theta)\left[q(\theta, H) - q(\theta, L)\right]\left[p(\theta) - c\right]\Bigg] dF(\theta) \end{aligned} \tag{RP}$$

$$\begin{aligned} s.t. \quad \forall\theta: \quad & 0 \leq q(\theta, L) \leq q(\theta, H) \leq 1 \\ & L \leq p(\theta) \leq H \quad \text{if} \quad q(\theta, L) < q(\theta, H). \end{aligned}$$

Delving into the seller's marginal revenue, it is composed of two components (i) the usual *virtual surplus* measured by the total surplus minus the buyer's rent and (ii) the *fictional surplus* due to non-common priors, as follows:

$$\underbrace{\theta q(\theta,H)(H-c)+(1-\theta)q(\theta,L)(L-c)-\pi'(\theta)\frac{1-F(\theta)}{f(\theta)}}_{\text{virtual surplus}}+\underbrace{(\theta_S-\theta)\left[q(\theta,H)-q(\theta,L)\right]\left[p(\theta)-c\right]}_{\text{fictional surplus}},$$

where $\pi'(\theta)=\Big[q(\theta,H)\left[H-p(\theta)\right]-q(\theta,L)\left[L-p(\theta)\right]\Big]$.

If the buyer is relatively optimistic (*i.e.*, $\theta_S-\theta\leq 0$), the seller is more likely to garble information, reducing the difference between $q(\theta,H)$ and $q(\theta,L)$ and thereby, eroding the negative impact of *fictional surplus*. By contrast, if the buyer is relatively pessimistic (*i.e.*, $\theta_S-\theta\leq 0$), the seller is more willing to disclose information, increasing the difference between $q(\theta,H)$ and $q(\theta,L)$, and thereby, enhancing the positive impact of *fictional surplus*.

In addition, the seller also needs to consider how information disclosure affects the virtual surplus. To set intuition on its joint impact via the two channels, focus on full information disclosure (which is indeed without loss). By offering full information, the seller persuades the (pessimistic) buyer at a higher price at $v=H$, improving *fictional surplus* by

$$\Delta_{FS}=(\theta_S-\theta)(H-c).$$

At the same time, full disclosure lead to no purchase at $v=L$, reducing the *virtual surplus* (relative to no disclosure/always sending signals "buy") by

$$\Delta_{VS}=(1-\theta)(L-c)-(H-L)\frac{1-F(\theta)}{f(\theta)}.$$

The optimal mechanism trades off *fictional surplus* gain and the *virtual surplus* loss, following a cut-off structure as described in Theorem 4. Formally, let

$$\mathscr{H}(\theta)\equiv\Delta_{FS}-\Delta_{VS}=(\theta_S-\theta)(H-c)-(1-\theta)(L-c)+(H-L)\frac{1-F(\theta)}{f(\theta)}$$

represent the net surplus gain from full disclosure. Under Assumption 4, $\mathscr{H}(\theta)$ decreases in θ. Hence, there exists $\theta^\star$ such that $\mathscr{H}(\theta)\geq 0\Leftrightarrow\theta\leq\theta^\star$. Then:

Theorem 4. *The optimal mechanism follows a cutoff rule under which:*

- *All types* $\theta\in(\theta^\star,\overline{\theta}]$ *receive no information and a posted price* $p)=\mathbb{E}_{\theta^\star}[v]$.
- *All types* $\theta'\in[0,\theta^\star]$ *receive full information and a posted price* $p'=H$.

See Appendix A for a formal proof. By Theorem 4, the optimal menu of price-information bundles consists of two options: (i) types below the cut-off receive full information and a relatively higher price, and (ii) types above the cut-off get no information but enjoy a lower price. Note that in general, $\theta^\star$ differs from θ_S. This implies that it is generally suboptimal to offer full (*resp.*, no) disclosure to *any* type who is pessimistic (*resp.*, optimistic) from the seller's perspective.

The optimal mechanism is illustrated in the following firgure

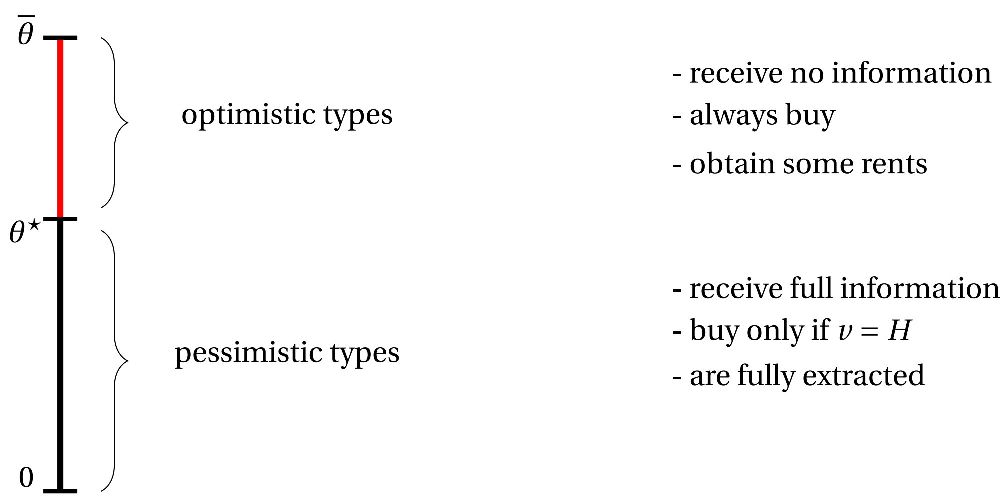

Figure 2.1: Optimal mechanism

It is also worth noting that the solution induces the well-known "no rent at the bottom and no distortion at the top" with the optimistic types receiving an efficient allocation (always trade) and the pessimistic ones fully extracted (he buys (only if the match value is high) at the price equal to his true value). To end this section, we provide a simple example for an illustration.

Example 8. *$\theta \sim U[0,1]$ and $\theta_S = \frac{\theta^2}{2}$. Valuations are such that: $H = 1$ and $L = c = 0$. Then, $\theta^\star \equiv \max_\theta\{\theta \mid H(\theta) \leq 0\} = \frac{3-\sqrt{5}}{2}$. By Theorem 4, under the optimal mechanism, any type $\theta \geq \frac{3-\sqrt{5}}{2}$ receives no disclosure and a posted price $p = \mathbb{E}_{\theta^\star}[v] = \frac{3-\sqrt{5}}{2}$, whereas any type $\theta' < \frac{3-\sqrt{5}}{2}$ receives full disclosure and a posted price $p' = H = 1$.*

5.2 Optimality of zero information fees

The main model assumes that information is provided free of charge. Accordingly, the buyer can walk away without buying the good and paying anything after information disclosure. In this section, we show that this assumption is without loss of optimality. Suppose now the seller offers a menu of $\{c(\theta), p(\theta), q(\theta, H), q(\theta, L)\}_\theta$ to screen the buyer where $c(\theta)$ type θ's information fee. Then, his on-path and off-path payoffs now become:

$$\hat{\pi}^f(\theta) = -c(\theta) + \pi^f(\theta),$$
$$\hat{\pi}^f(\theta,\theta') = -c(\theta') + \pi^f(\theta,\theta'),$$
$$\hat{\pi}^d(\theta,\theta') = -c(\theta') + \pi^d(\theta,\theta').$$

First, as information fees have been paid before information disclosure, the seller faces the same obedience constraints (OB^b) and (OB^{nb}) as in the setup with free information. Therefore, lemmas (13) and (14) continue to hold. Second, there are also IC constraints:

$$\forall \theta, \theta': \quad \hat{\pi}(\theta) \geq \max\left\{\hat{\pi}^f(\theta,\theta'), \hat{\pi}^d(\theta,\theta')\right\}. \tag{$\widehat{IC}$}$$

Finally, the choice of information fees, $c(\theta)$, are subject to interim IR constraints:

$$\forall \theta: \quad \hat{\pi}(\theta) \geq 0. \tag{$\widehat{IR}$}$$

Therefore, the seller's problem can be written as follows:

$$\max_{c,p,q} \quad \int_0^1 c(\theta)dF(\theta) + \int_0^1 \Big[\big[\theta_S q(\theta,V) + (1-\theta_S) q(\theta,L)\big]\big[p(\theta)-c\big]\Big] dF(\theta) \tag{$\widehat{P}$}$$
$$s.t. \quad (FC), (OB^b), (OB^{nb}), (\widehat{IC}), \text{ and } (\widehat{IR}).$$

Let $V(P)$ (*resp.,* $V(\widehat{P})$) represent the value of problem (P) (*resp.,* $(\widehat{P})$). Note that with $c(\theta) = 0$ $\forall\theta$, the two problems (P) and $(\widehat{P})$ coincide. Therefore, the value of $(\widehat{P})$ is an upper bound of its counterpart: $V(P) \leq V(\widehat{P})$. Interestingly, as we will show, this upper bound is tight.

The key is that information fees do not affect the buyer's marginal rent: $\hat{\pi}'(\theta) = \pi'(\theta)$. Consequently, by integration by parts, $\int_0^1 \hat{\pi}(\theta) = \int_0^1 \pi'(\theta)[1-F(\theta)]d\theta + \hat{\pi}(0)$. Using this, lemmas (13) and (14), and the fact that $\hat{\pi}(0) = 0$ at optimum, the seller's relaxed problem *with* information fees writes:

$$\max_{\{p,q\}} \quad \int_0^1 \Big[\theta q(\theta,H)(H-c) + (1-\theta)q(\theta,L)(L-c) - \pi'(\theta)\frac{1-F(\theta)}{f(\theta)}$$
$$+ (\theta_S - \theta)\big[q(\theta,H) - q(\theta,L)\big]\big[p(\theta)-c\big]\Big] dF(\theta) \tag{$\hat{RP}$}$$
$$s.t. \quad \forall\theta: \quad 0 \leq q(\theta,L) \leq q(\theta,H) \leq 1$$
$$L \leq p(\theta) \leq H \quad \text{if} \quad q(\theta,L) < q(\theta,H).$$

Therefore, $(\hat{RP})$ coincides with (RP). As a consequence,

$$V(P) \leq V(\widehat{P}) \leq V(\hat{RP}) = V(RP) = V(P),$$

where the last equality follows from the fact that a solution of (RP) solves (P). Thus, $V(P) = V(\widehat{P})$. This establishes the optimality of zero information fees.

Proposition 8. *The seller does not benefit from using information fees, i.e., $V(P) = V(\widehat{P})$.*

6 DISCUSSION

6.1 Alternative modeling

In our model, ex-ante heterogeneity is introduced via non-common priors. One might wonder whether this could be viewed as a reduced-form version of a model where the buyer, prior to interacting with the seller, privately receives some information about his valuation. In this section, we show that this is not true. First, we formally define this alternative setting, called "model B", as below:

Model B: *The buyer faces uncertainty about his valuation $v \in \{L, H\}$. Both the buyer and the seller initially share a common prior about the distribution of valuations: $Prob(v = H) = \mu_0$. Then, the buyer privately receives a signal θ about his valuation. Upon observing such a signal θ, the buyer and the seller agree that $Prob(v = H) = \theta$.*[6]

We emphasize that in model B, while the parties share a common (posterior) belief, only the buyer observes his signal θ. In other words, the buyer's belief is *private* and *unbiased.* As a result, the seller's maximization problem uses the buyer's belief (which is also the seller's) to formulate the objective function (and constraints), as folows:

$$\begin{aligned} \max_{\{p,q\}} \quad & \int_0^1 \left[\theta q(\theta, H) + (1-\theta) q(\theta, L)\right] \left[p(\theta) - c\right] dF(\theta) \\ s.t. \quad & (FC), (OB^b), (OB^{nb}), (IC). \end{aligned} \tag{$\tilde{P}$}$$

A striking difference between the seller's problems in our setting and model B is that in the latter, the *fictional* surplus disappears. This can significantly reshape the optimal mechanism. To illustrate, let us solve model B in a special case with $L = c = 0$. In this case, the (expected)

[6] Model B can be seen as the *binary*-state and *infinite*-type version of Guo et al. (2022).

total surplus is θH. The seller can extract all the surplus by offering full disclosure and a price of H for the good. Therefore, *bunching* is revenue maximizing.

By contrast, when the seller holds a subjective belief θ_S, Theorem 4 shows that the optimal mechanism is, in general, a *screening* menu of prices and information. Specifically, with $L = c = 0$, $\mathcal{H}(\theta) = H\Big(\theta - \frac{1-F(\theta)}{f(\theta)} + \theta_S\Big)$. Then, the optimal menu of price-information bundles has two items, separating types above and below the cut-off type $\theta^\star = \inf\Big\{\theta \mid \theta - \frac{1-F(\theta)}{f(\theta)} + \theta_S \geq 0\Big\}$.

Note that the optimality of bunching in model B holds beyond this special case with $L = c = 0$. Formally:

Proposition 9. *Consider model B. If* $\dfrac{(1-\theta)L}{H-L} \leq \dfrac{1-F(\theta)}{f(\theta)}$*, then it is optimal to offer full disclosure and a fixed price* $p = H$ *to all types.*

The condition stated in Proposition 9 satisfies when, for example, the valuation gap $H - L$ sufficiently large or $L \leq c$.

6.2 Multiple valuations

A limitation of our work is its focus on binary valuations. With more than two types, partial disclosure could be necessary, which significantly complicates the trade-off between the fictional surplus gain and the virtual surplus loss (and vice versa) due to information disclosure. Consequently, the characterization of the optimal mechanism becomes challenging and is out of scope of this paper.

Regarding the optimality of zero information fees, the sort of "revenue-equivalence" argument holds regardless of the valuation space: as long as local downward deviations are sufficient, information fees do not affect the buyer's marginal payoff and thereby, the seller's revenue. With more than two values, however, the validity of this "first-order approach", and thereby, the optimality of zero information fees is not guaranteed.[7] Indeed, in the following example with three possible valuations, postive information fees are part of the optimal mechanism.

Example 9. $v \in \{-1, 1, 3\}$ *and* $\theta \in \{\theta_1, \theta_2\}$*. Type* θ_2*'s prior belief is given by* μ_2*:*

$$\mu_2(-1) = \mu_2(1) = \mu_2(3) = \frac{1}{3}.$$

[7] See **?** for an example of binding global and upward IC constraints in a joint mechanism and information design problem with general state and type space.

Type θ_1's prior belief is given by μ_1:

$$\mu_1(-1) = \frac{2}{3}, \quad \mu_1(1) = 0, \quad \mu_1(3) = \frac{1}{3}.$$

The seller's belief μ_S is equal to type θ_2's, that is, $\mu_S = \mu_2$. Thus, type θ_2 is unbiased, where as θ_1 is pessimistic from the seller's perspective. The seller does not know the buyer's true type, but knows that the two types are likely equal.

Let R_M be the seller's revenue obtained from a menu $M \equiv \{c(\theta), p(\theta)q(\theta, v)\}_{\theta,v}$. We have:

$$\begin{aligned} R_M =& f(\theta_2)\left[\sum_v \mu_S(v)q(\theta_2, v)p(\theta_2) + c(\theta_2)\right] + f(\theta_1)\left[\sum_v \mu_S(v)q(\theta_1, v)p(\theta_1) + c(\theta_1)\right] \\ =& \underbrace{f(\theta_2)\left[\sum_{v\geq 0} \mu_2(v)q(\theta_2, v)p(\theta_2) + c(\theta_2)\right] + f(\theta_1)\left[\sum_v \mu_1(v)q(\theta_2, v)p(\theta_1) + c(\theta_1)\right]}_{\equiv R_M^1} \\ &+ \underbrace{f(\theta_1)\left[\mu_S(-1) - \mu_1(-1)\right]q(\theta_1, -1)p(\theta_1) + f(\theta_1)\left[\mu_s(1) - \mu_1(1)\right]q(\theta_1, 1)p(\theta_1)}_{\equiv R_M^2}, \end{aligned}$$

where R_M^1 is the revenue obtained from M under no biasedness $\mu_S(v) = \mu_\theta(v)$, and R_M^2 represents the fictional surplus due to non-common prior between the seller and type θ_1. Note that R_M^1 is bounded by the (expected) total surplus from the buyer's perspective. The seller can fully extract this surplus by offering the following mechanism, called $M^\star$. Under $M^\star$, type θ_2 receives $q(v,\theta_2) = \mathbb{1}_{\{v\geq 1\}}$, and $\left(c(\theta_2), p(\theta_2)\right) = \left(\frac{4}{3}, 0\right)$; and type θ_1 receives $q(v,\theta_1) = \mathbb{1}_{\{v\geq 1\}}$, and $\left(c(\theta_1), p(\theta_1)\right) = (0, 3)$.

Note also that $\mu_s(-1) - \mu_1(-1) = \frac{-1}{3} < 0$, $\mu_s(1) - \mu_1(1) = \frac{1}{3} > 0$, $0 \leq q(\theta_1, v) \leq 1$, $p(\theta_2) \leq 3$; therefore, R_M^2 is maximized at $q(\theta_1, -1) = 0$, $q(\theta_1, 1) = 1$, and $p(\theta_1) = 3$. Thus, $M^\star$ also maximizes R_M^2 and hence, is optimal.

Now, suppose that the seller offers type θ_2 information free of charge, or $\tilde{c}(\theta_2) = 0$. Then, to extract all surplus from type θ_2, it is necessary to modify the price for this type to $\tilde{p}(\theta_2) = 2$. In turn, this triggers type θ_1 to mimic type θ_2 because his deviating payoff is $\mu_1(3)\left[3 - \tilde{p}(\theta_2)\right] = \frac{1}{3} > 0$. Therefore, it is impossble for the seller to extract the full surplus measured by the buyer's prior with zero information fees.

7 CONCLUDING REMARKS

The common prior assumption has been extensively employed in economic theory, often for technical convenience. Instead of following this routine, we introduce non-common priors in a joint mechanism and information design problem. This uncovers a new trade-off (between the virtual and fictional surplus) for designing information disclosure and how non-common priors shape optimal mechanisms.

There are several follow-up questions. First, we establish the optimality of free information when the optimal mechanism can be solved via a standard relaxed problem. A natural question is then under which environments, this approach remains valid. Another direction is to characterize the optimal mechanism under agents' optimism about some payoff-relevant state in other environments such as auctions or collective-decision making.

Alternatively, one could consider a robustness approach in a monopolistic screening environment. How would the seller design a robustly optimal mechanism when she has little knowledge about the buyer's prior belief and would like to maximize the worst-case revenue? Such a question has been studied in both the Bayesian persuasion and mechanism design literatures separately, but not in their intersection. We leave these potential extensions for future research.

A PROOF OF THEOREM 4

Proof. The proof proceeds by first solving the relaxed problem (RP). We then verify that its solution satisfies ignored constraints, and hence, solves the original problem.

Step 1: In this step, we solve the relaxed problem (RP). Using $[\pi(\theta)]' = \left[q(\theta,H)\left[H-p(\theta)\right] - q(\theta,L)\left[L-p(\theta)\right]\right]\frac{1-F(\theta)}{f(\theta)}$, ($RP$) becomes

$$
\begin{aligned}
\max_{\{p,q\}} \quad & \int_0^1 \Big\{\theta q(\theta,H)(H-c) + (1-\theta)q(\theta,L)(L-c) + (\theta_S-\theta)\left[q(\theta,H)-q(\theta,L)\right]\left[p(\theta)-c\right] \\
& \qquad - \left[q(\theta,H)\left[H-p(\theta)\right] - q(\theta,L)\left[L-p(\theta)\right]\right]\frac{1-F(\theta)}{f(\theta)}\Big\} dF(\theta) \\
s.t. \quad \forall\theta: \quad & 0 \le q(\theta,L) \le q(\theta,H) \le 1 \\
& L \le p(\theta) \le H \quad \text{if} \quad q(\theta,L) < q(\theta,H).
\end{aligned}
$$

Let

$$
\begin{aligned}
R(q(\theta,v),p(\theta)) \equiv & \theta q(\theta,H)(H-c) + (1-\theta)q(\theta,L)(L-c) + (\theta_S-\theta)\left[q(\theta,H)-q(\theta,L)\right]\left[p(\theta)-c\right] \\
& - \left[q(\theta,H)\left[H-p(\theta)\right] - q(\theta,L)\left[L-p(\theta)\right]\right]\frac{1-F(\theta)}{f(\theta)}.
\end{aligned}
$$

denote the point-wise objective function. Note that

$$\frac{\partial R(q(\theta,v),p(\theta))}{\partial p(\theta)} = -\Big(\theta - \frac{1-F(\theta)}{f(\theta)} - \theta_S\Big)\left[q(\theta,H) - q(\theta,L)\right].$$

Let $\mathcal{K}(\theta) \equiv \theta - \dfrac{1-F(\theta)}{f(\theta)} - \theta_S$. By Assumption 4, $\dfrac{1-F(\theta)}{f(\theta)}$ decreases in θ. Therefore, $\mathcal{K}(\theta)$ increases in θ. Consequently, there exists $\hat{\theta}$ such that $\mathcal{K}(\theta) \geq 0 \Leftrightarrow \theta \geq \hat{\theta}$. Consider the following two cases:

Case 1: $\theta \geq \hat{\theta}$. Then, either (i) $q(\theta,L) = q(\theta,H)$, or (ii) $q(\theta,L) < q(\theta,H)$, $p(\theta) = L$.

- Case 1(i) $q(\theta,H) = q(\theta,L) = q(\theta)$. the point-wise objective becomes

$$\begin{aligned} R(q(\theta,v),p(\theta)) &= \left[\theta H + (1-\theta)L - c - (H-L)\frac{1-F(\theta)}{f(\theta)}\right] q(\theta) \\ &= \left[L - c + (H-L)\Big(\theta - \frac{1-F(\theta)}{f(\theta)}\Big)\right] q(\theta). \end{aligned}$$

  Since $\mathcal{K}(\theta) \equiv \theta - \dfrac{1-F(\theta)}{f(\theta)} - \theta_S \geq 0$ with $\theta \geq \hat{\theta}$ and $\theta_S \geq 0$, it must be that $\theta - \frac{1-F(\theta)}{f(\theta)} \geq 0$. Thus, $R(q(\theta,v),p(\theta))$ is maximized at

$$q(\theta,L) = q(\theta,H) = q(\theta) = 1. \tag{2.5}$$

  This solution leads to

$$R(q(\theta,v),p(\theta)) = L - c + (H-L)\Big(\theta - \frac{1-F(\theta)}{f(\theta)}\Big) \equiv R_1(\theta).$$

- Case 1(ii): $0 < q(\theta,L) < q(\theta,H)$, $p(\theta) = L$. In this case,

$$\begin{aligned} R(q(\theta,v),p(\theta)) =& \Big[\theta(H-c) + (\theta_S - \theta)[L-c] - (H-L)\frac{1-F(\theta)}{f(\theta)}\Big] q(\theta,H) + (1-\theta_S)(L-c)q(\theta,L) \\ <& \Big[\theta(H-c) + (\theta_S - \theta)[L-c] - (H-L)\frac{1-F(\theta)}{f(\theta)}\Big] q(\theta,H) + (1-\theta_S)(L-c)q(\theta,H) \\ =& \Big[\theta H + (1-\theta)L - c - (H-L)\frac{1-F(\theta)}{f(\theta)}\Big] q(\theta,H) \\ \leq& \theta H + (1-\theta)L - c - (H-L)\frac{1-F(\theta)}{f(\theta)} \\ =& R_1(\theta). \end{aligned}$$

Therefore, in case 1, the solution is given by (2.5).

Case 2: $\theta < \hat{\theta}$. Then, either (i) $q(\theta, L) = q(\theta, H)$, or (ii) $q(\theta, L) < q(\theta, H)$, $p(\theta) = H$.

1. Case 2(i): By similar arguments to Case 1(i), the solution is given by (2.5)
2. Case 2(ii): $q(\theta, L) < q(\theta, H)$, $p(\theta) = H$. Then,

$$R(q(\theta,v),p(\theta)) = \theta_S q(\theta,H)(H-c) + \Big[\theta H + (1-\theta)L - c - (H-L)\frac{1-F(\theta)}{f(\theta)} - \theta_S(H-c)\Big] q(\theta,L),$$

Let

$$\begin{aligned}\mathscr{H}(\theta) &\equiv \theta_S(H-c) - \theta H - (1-\theta)L + c + (H-L)\frac{1-F(\theta)}{f(\theta)} \\ &= \theta_S(H-c) - (L-c) - (H-L)\Big(\theta - \frac{1-F(\theta)}{f(\theta)}\Big).\end{aligned}$$

By Assumption 4, $\frac{1-F(\theta)}{f(\theta)}$ decreases in θ. Hence, $\mathscr{H}(\theta)$ decreases in θ. Consequently, there exists $\theta^\star$ such that $\mathscr{H}(\theta) \geq 0 \Leftrightarrow \theta \leq \theta^\star$.

First, consider $\theta > \theta^\star$, then $\mathscr{H}(\theta) < 0$ and

$$\begin{aligned}R(q(\theta,v),p(\theta)) =& \theta_S q(\theta,H)(H-c) + \Big[(1-\theta_s)(L-c) - (H-L)\big(\theta_S - \theta + \frac{1-F(\theta)}{f(\theta)}\big)\Big] q(\theta,L) \\ <& \theta_S(H-c) + \Big[(1-\theta_s)(L-c) - (H-L)\big(\theta_S - \theta + \frac{1-F(\theta)}{f(\theta)}\big)\Big] \\ =& \theta H + (1-\theta)L - c - (H-L)\frac{1-F(\theta)}{f(\theta)} f(\theta) \\ =& R_1(\theta),\end{aligned}$$

which is what was obtained in Case 2(i).

Second, consider $\theta \leq \theta^\star$, then $\mathscr{H}(\theta) \geq 0$. Therefore, then $R(q(\theta,v),p(\theta))$ is maximized at $q(\theta, L) = 0$ and $q(\theta, H) = 1$, leading to

$$R(q(\theta,v),p(\theta)) = \theta_S(H-c) \equiv R_2(\theta).$$

For this solution to outperform that in Case 2(i) which is given by (2.5), we need

$$R_2(\theta) \geq R_1(\theta) \Leftrightarrow R_2(\theta) - R_1(\theta) = \theta_S(H-c) - \theta H - (1-\theta)L + c + (H-L)\frac{1-F(\theta)}{f(\theta)} \equiv \mathscr{H}(\theta) \geq 0, \tag{2.6}$$

which is true with $\theta \leq \theta^\star$.

To sum up,

- If $\theta > \theta^\star$, then in either case 1 or case 2, it is optimal to set $q(\theta, L) = q(\theta, H)$. In this case, θ receives no disclosure (he always observes signal "buy"). We choose $p(\theta) = \mathbb{E}_{\theta^\star}[v] \equiv \theta^\star H + (1-\theta^\star)L$ for all $\theta \geq \theta^\star$ so that the solution satisfies all constraints in the original problems (see Step 2 of the proof).
- If $\theta \leq \theta^\star$, then it is optimal to set $q(\theta, L) = 0$, $q(\theta, H) = 1$ and $p(\theta) = H$. In this case, θ receives full disclosure and is fully extracted (he observes signal "buy" if and only if his valuation is high).

Step 2: In this step, we show that the menu with two options (i) full disclosure coupled with a posted price of H for type $\theta \geq \theta^\star$ and (ii) no disclosure associated with a posted price of $\mathbb{E}_{\theta^\star}[v]$ for $\theta' > \theta^\star$ induces truth-telling and obedience. Note that any type choosing the second option buys only if observing signal "buy" (or $v = H$) and earns a zero payoff. Consider $\theta \leq \theta^\star$. If he takes the first option, it is optimal for him not to buy the good and get nothing. Therefore, θ prefers to reveal his type and buys the good if and only if his value is H. Next, consider $\theta' > \theta^\star$, by taking the first option, he pays less than his expected valuation to (always) get the good. As a result, it is optimal for him to reveal his type and obey recommended signals to enjoy a strictly positive rent. □

B PROOF OF PROPOSITION 9

Proof. If $L \leq c$, then the seller can extract all the surplus by providing full disclosure and a posted price $p = H$ to all types. Next, suppose $L \geq c$, then we can apply Theorem 4 for a special case with $\theta_S = \theta$. Note that with $\theta_S = \theta$,

$$\mathcal{H}(\theta) = (\theta_S - \theta)(H-c) - (1-\theta)(L-c) + (H-L)\frac{1-F(\theta)}{f(\theta)} = (H-L)\frac{1-F(\theta)}{f(\theta)} - (1-\theta)(L-c).$$

Therefore, if $(1-\theta)(L-c) \leq (H-L)\frac{1-F(\theta)}{f(\theta)}$ for all θ, $\mathcal{H}(\theta) \geq 0$ for all θ. As a result, full disclosure, coupled with a posted price equal to H is optimal.

□

Chapter 3

Auction Design with Heterogeneous Priors

1 INTRODUCTION

The common knowledge assumptions have been challenged by many papers in the literature on robust mechanism design (see our detailed discussion in the related literature section). In an influential work, Chung and Ely (2007) consider an auction environment where the seller has little idea about each bidder's belief about the other bidders' valuations. They show that, for some specification of the bidders' beliefs (formally identified by a type space), a dominant-strategy auction mechanism is revenue-maximizing among all Bayesian incentive compatible auction mechanisms, even if the seller knows that that type space governs the bidders' beliefs (*Baysian foundation* for a dominant-strategy mechanism). As a consequence, in case the seller does not know which type space governs their beliefs, a dominant-strategy auction mechanism is max-min optimal (*Maximin foundation*).

The interim belief that each bidder of each type must have in this critical type space is special, and seems very different from what any common-prior type space would imply. Indeed, Chung and Ely (2007) provide a counterexample such that a dominant-strategy auction mechanism cannot be Bayesian-founded if the bidders' type space must be one of the common-prior type spaces.

A natural question is "how far" this Chung-Ely's type space is relative to those given by some common prior. To investigate this question, we examine the class of types spaces which are induced by *(ϵ-) heterogeneous priors*. Namely, each player (seller and each bidder) possesses a prior distribution about the value distribution before their values being drawn, which can be

ε-different from each other (in the metric similar to the one considered by Madarász and Prat (2017) and Carroll (2017)). Then, each bidder's value is drawn, making him Bayesian update his own prior conditional on his own value. Clearly, with $\varepsilon = 0$, the model reduces to the standard common-prior case, and hence, no foundation. With a large enough ε, it is natural to think that the critical type space of Chung-Ely can be captured, and hence a foundation exists. We show that, in fact, the critical type space of Chung-Ely can be represented by a type space induced by ϵ-heterogenous priors, *for any* $\varepsilon > 0$, no matter how small it is. Therefore, with any $\varepsilon > 0$, the dominant-strategy auction mechanism is Bayesian (and hence Maximin) founded.

This result seems counter-intuitive, given that full surplus extraction is possible when $\varepsilon = 0$ and values are sufficiently correlated according to the common prior (see Crémer and McLean (1988*a*)). The reason why a small distance $\varepsilon > 0$ between the priors can disrupt this possibility is explained via a simple example in Section 4. The basic intuition is that by introducing small perturbations in the support of the prior beliefs, it becomes possible to significantly upset the correlation structure and, thereby, the interim beliefs of bidders. Hence, even if a bidder's prior is close to the others' (in particular, to the seller's), it does not mean that their "interim beliefs" are close to each other. In fact, they can be so flexible that any small (but positive) heterogeneity in their priors can result in very different interim beliefs. This gives room to construct a type space induced by ε-different priors, for any $\varepsilon > 0$, to represent Chung and Ely (2007)'s critical type space.[1]

Although the original result of Chung and Ely (2007) suggests that the dominant-strategy approach would be reasonable in case the seller has *very little* idea about the bidders' information (for example, when there have not been similar items auctioned), it is sometimes informally argued that, if rich data is available about past similar auctions, it might be more difficult to justify the dominant-strategy approach, as both the seller and bidders would have a more precise idea about the true value distribution. In practice, the players typically have *some* information about past similar auctions, though they never have an *exact* common prior. In this sense, it is important to investigate the "boundary" of Chung-Ely's argument: With which class of type spaces (related to which information of the bidders about past similar auctions) the dominant-strategy approach has a Chung-Ely foundation? The result of our paper contributes to a better understanding of this question by examining (possibly small) heterogeneity in the players' priors.

[1] In Appendix D, we provide a counterexample in which the dominant-strategy auction mechanism is *not* founded when interim beliefs are *close* enough.

While the optimal dominant-strategy mechanism is *belief-free*, it hinges on the (true) *joint distribution* of bidders' valuations. Theorem 6 shows that if the seller is confident that her prior is ε-close to the true distribution of valuations, the highest revenue guarantee converges to that under the best dominant-strategy mechanism when the seller knows the true distribution. Note that if the seller naively offers an optimal dominant mechanism based on her own prior, even when it is very close to the truth, some bidders may strictly prefer to imitate those whose values are far from theirs. However, we show that by appropriately reducing the transfers of such a naively optimal dominant-strategy mechanism (by an amount as a function of ε), incentive compatibility is restored in a stronger sense: it is dominant for bidders to reveal their values as much as possible (by reporting the value that is closest to their true value in that mechanism). As a consequence, the highest revenue guarantee can be approximated via this *transfer-reducing mechanism*.

1.1 Related literature

This paper contributes to the growing literature on robust mechanism design (see, for example, Bergemann and Morris (2005), Chung and Ely (2007), Chen and Li (2018), and Yamashita and Zhu (2022) as the most relevant ones to this paper). These papers consider the situation where the agents' beliefs can be arbitrarily different from each other (and from the principal's, if the principal has a prior). For example, as aforementioned, Chung and Ely (2007) identifies a type space with heterogeneous priors with which one of the optimal Bayesian mechanisms is a dominant-strategy mechanism. Thus, if the seller has little idea about the bidders' beliefs, then the worst-case-minded seller has a justification to offer a dominant-strategy mechanism. See Chen and Li (2018) for its generalization to non-auction environments. We show that, even if the seller has a much better idea about the bidders' beliefs in that their priors are arbitrarily close to each other and also to the seller's (and that being their common knowledge), essentially the same conclusion is obtained. In this sense, our result strengthens that of Chung and Ely (2007).

Our notion of prior perturbations is related to the (various) notions of "local robustness" in the literature. For example, in Lopomo et al. (2021) where each agent's type is associated with a set of "fully overlapping"[2] interim beliefs, a mechanism is robust if it is implementable for every

[2] Roughly speaking, this "fully overlapping" requirement means that nearby types share a sufficiently rich set of beliefs. A focal special case is when the set is an arbitrarily small neighborhood around a fixed belief. See also Lopomo et al. (2022) where they derive the necessary and sufficient conditions for full extraction in this setting.

possible interim belief. They find that robustness is hard to achieve even when this set is arbitrarily small. As another example, Ollár and Penta (2017) propose a general form of restrictions directly on the agents' interim beliefs, and show that, when the set of possible interim beliefs is small in an appropriate sense, much more permissive results are possible. Our result shows that the *ex ante* belief restriction does not imply their *interim* restriction, and hence they lead to very different results. In this sense, our notion of uncertainty may be interpreted as "ex-ante-local" uncertainty. Jehiel et al. (2012) consider a related notion of local uncertainty in terms of interim beliefs, but in a generic multi-dimensional interdependent-value environment. They show that, if the principal's goal is to implement some belief-invariant social choice function, then the same kind of an impossibility result is obtained as in Jehiel et al. (2006) (where the latter paper considers ex post implementation, and in this sense allows for global robustness). Our environment is with private values, and the seller's goal is revenue maximization rather than a social choice function implementation.[3]

In a single-agent environment, Madarász and Prat (2017) consider a situation where the principal is aware that the true distribution of the agent's type can be ε-different from what the seller has in mind. Carroll (2017) generalizes their notion of ε-closeness in the context of a (single-agent) multi-dimensional screening problem.[4] As far as we are aware, ours is the first paper that generalizes their notions of closeness to a multi-agent environment. Importantly, with multiple agents, it is not only the principal who is uncertain about the true distribution, but also the agents enjoy uncertainties about the true distributions and the others' beliefs. On the other hand, relative to Madarász and Prat (2017) and Carroll (2017), we focus on a single-good private-value auction, with which the agents' payoff structures satisfy the single-crossing conditions.

There has been some work on mechanism design with heterogeneous priors. For example, in **?**, a consumer assigns excessive weights to the states of nature associated with their large gains from trade. They find that non-common priors can be necessary for price discrimination. Grubb (2009) studies a situation where a consumer assigns wrong weights to their possible valuations (narrowly concentrates around the mean), relative to the seller's prior. He mainly focuses on characterizing the optimal contract under complete information. Our paper intro-

[3] Hence, in principle, the seller might find it optimal to use a mechanism that induces a highly belief-dependent outcome. Put differently, the set of feasible mechanisms in our case is larger than those that implement a belief-invariant social choice function.

[4] See also Bergemann and Schlag (2011). Carroll and Meng (2016) considers the local robustness in a single-agent moral-hazard environment.

duces heterogeneous priors in the auction context (i.e., with multiple agents rather than a single representative agent).

2 AUCTION ENVIRONMENT

A seller wants to sell an indivisible good. There are N risk-neutral bidders with private values. Each bidder $i \in \{1,\dots,N\}$ knows his own valuation $v_i \in \mathbb{R}$. An allocation is denoted by $(q,p) = (q_i,p_i)_{i=1}^N$, where $q_i \in [0,1]$ denotes the probability that bidder i obtains the good, and $p_i \in \mathbb{R}$ denotes his payment to the seller. An allocation is feasible if $\sum_i q_i \le 1$. Given (q_i,p_i), i's payoff is given by $v_i q_i - p_i$.

The players (the seller and the bidders) enjoy *heterogeneous priors* for the distribution of the bidders' values v. Specifically, let $g \in \Delta(\mathbb{R}^N)$ be the seller's prior, which has a finite support represented by $\{\gamma, 2\gamma, \dots, K\gamma\}^N$ (following Chung and Ely (2007)) for some $K \in \mathbb{N}$ and $\gamma > 0$ for notational simplicity. Throughout the paper, we assume that g satisfies the single-crossing virtual value condition (Chung and Ely (2007)). For each $i \ne j$ and v, let $\gamma_i(v)$ be i's virtual valuation:

$$\gamma_i(v) = v_i - \gamma \frac{1 - G_i(v)}{g(v)},$$

where $G_i(v) = \sum_{v_i' \le v_i} g(v_i', v_{-i})$.

Assumption 5. *For each $i \ne j$, and each v_i, v_i', v_{-i} with $v_i' > v_i$:*

$$\begin{aligned} \gamma_i(v_i, v_{-i}) \ge 0 \quad &\Rightarrow \quad \gamma_i(v_i', v_{-i}) > 0 \\ \gamma_i(v_i, v_{-i}) \ge \gamma_j(v_i, v_{-i}) \quad &\Rightarrow \quad \gamma_i(v_i', v_{-i}) > \gamma_j(v_i', v_{-i}). \end{aligned}$$

As shown in Chung and Ely (2007), it is satisfied if g exhibits affiliation and monotone hazard-rates. In this sense, it may be considered a mild assumption.

As opposed to the standard exact-common-prior model where not only the seller but every bidder i believes g (and that itself being common knowledge), we allow the possibility that they enjoy heterogeneous priors: For each i, let $h_i \in \Delta(\mathbb{R}^N)$ be bidder i's prior, which again has a finite support for simplicity (but potentially with a different support from g and from h_j, $j \ne i$). Note that bidder i knows his own value v_i at the time he plays an auction mechanism. That his prior is h_i implies that his belief about the others' values is based on h_i conditional on his v_i.

We assume that the seller has limited knowledge as to "how distant" each bidder i's h_i could be from the seller's prior g. This distance may be interpreted as the level of the seller's confidence in his own information.[5]

Our notion of distance is based on Madarász and Prat (2017) and Carroll (2017):

Definition 9. *Two distributions* μ *and* $\hat{\mu}$ *are* ε*-close to each other if* $V = \text{supp}(\mu)$ *and* $\hat{V} = \text{supp}(\hat{\mu})$ *can be partitioned into disjoint measurable sets* $\{V^1, ..., V^r\}$ *and* $\{\hat{V}^1, ..., \hat{V}^r\}$ *respectively such that, for each* $k \in \{1, ..., r\}$*:*

1. $\mu(V^k) = \hat{\mu}(\hat{V}^k)$*, and*
2. $d(v, \hat{v}) \leq \varepsilon$ *for any* $(v, \hat{v}) \in V^k \times \hat{V}^k$*,*

where $d(v, \hat{v})$ *represents the Euclidean distance between* v *and* $\hat{v}$*.*

A collection of distributions $\{\mu_1, \ldots, \mu_K\}$ *is* ε*-close to each other if any pair* μ_i, μ_j *are* ε*-close to each other as above.*

Example 10. *We illustrate the closeness of two distributions in the following example with* $N = 2$*. Let* g *be the distribution represented as follows:*

$g(v_1, v_2)$	$v_2 = 1$	$v_2 = 2$
$v_1 = 1$	$\frac{1}{3}$	$\frac{1}{6}$
$v_1 = 2$	$\frac{1}{6}$	$\frac{1}{3}$

Table 3.1: Distribution g

and let f *be represented as follows:*

$f(v_1, v_2)$	$v_2 = 1 - \varepsilon$	$v_2 = 1$	$v_2 = 2 - \varepsilon$	$v_2 = 2$
$v_1 = 1 - \varepsilon$	$\frac{1}{3}$			
$v_1 = 1$			$\frac{1}{6}$	
$v_1 = 2 - \varepsilon$		$\frac{1}{6}$		
$v_1 = 2$				$\frac{1}{3}$

Table 3.2: Distribution f

[5] This interpretation implies a related but different question: what if the seller's prior g is different from the true value distribution? For now, we assume that the seller is confident in his own g as the true value distribution, but we study the case where the seller fears the possibility that g is wrong. See Section 6.

Then, according to the definition above, f and g are $(\varepsilon\sqrt{2})$-close to each other. Figures 1 and 2 illustrate f and g in the (v_1, v_2)-space.

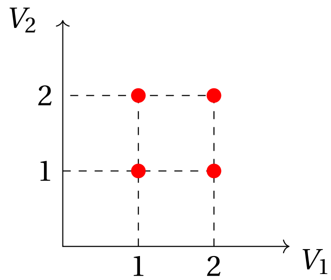

Figure 3.1: Distribution g

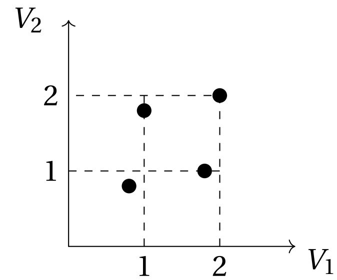

Figure 3.2: Distribution f

The seller believes that $(g, h_1, \ldots, h_N)$ are ε-close to each other for some given $\varepsilon > 0$, and that any such combination of h_i's is possible. This uncertainty makes the seller cautious in designing an auction mechanism.

3 AUCTION MECHANISM

Given the concern about the above prior heterogeneity, the seller can design a "robust" mechanism in a certain sense. One of the possible approaches is to design a *dominant-strategy* auction mechanism, where each bidder has a dominant action given each v_i regardless of the other bidders' behavior. Such a mechanism can guarantee some level of expected revenue regardless of each bidder's belief about the opponents' values and their (higher-order) beliefs; in particular, regardless of each i's prior h_i.

Another possibility is to try to extract each bidder's information (about each i's h_i, for example) in order to design a more profitable auction mechanism. Indeed, in the standard exact common-prior environment where $g = h_i$ for all i is common knowledge, if $g(= h_i)$ satisfies a certain correlation structure, the seller can extract the entire surplus (Cremer-McLean), while the optimal dominant-strategy mechanism leaves a non-negligible rent to the winning bidder. Even if h_i can be different from g, if the seller knows that they cannot be too far from each other, it may be natural to expect that better mechanisms than dominant-strategy mechanisms exist.

3.1 Notation

An auction mechanism is represented by $(M,q,p)=(M_i,q_i,p_i)_{i=1}^N$, where: each M_i is a set, $M=\prod_{i=1}^N M_i$, $q_i:M\to[0,1]$ with $\sum_i q_i(m)\le 1$ for all $m\in M$, and $p_i:M\to\mathbb{R}$. An interpretation is that, given mechanism (M,q,p), each bidder is asked to simultaneously choose any $m_i\in M_i$; and given a chosen vector $m=(m_1,\dots,m_N)\in M$, allocation $(q_i(m),p_i(m))_{i=1}^N$ is executed. A feasible mechanism must contain some element $\phi_i\in M_i$ for each i such that $q_i(\phi_i,m_{-i})=p_i(\phi_i,m_{-i})=0$ for any $m_{-i}\in M_{-i}$, representing the idea of i's individual rationality requirement.

3.2 Dominant-strategy auction mechanism

We first introduce *dominant-strategy auction mechanisms.*

Definition 10. *Mechanism (M,q,p) admits a dominant-strategy equilibrium if there exists $\sigma_i(v_i)$ for each i, $v_i\in\mathbb{R}$ such that, for each m_i,m_{-i}:*

$$\begin{aligned} v_i q_i(\sigma_i(v_i),m_{-i})-p_i(\sigma_i(v_i),m_{-i}) &\ge v_i q_i(m_i,m_{-i})-p_i(m_i,m_{-i}) \\ v_i q_i(\sigma_i(v_i),m_{-i})-p_i(\sigma_i(v_i),m_{-i}) &\ge 0. \end{aligned}$$

Mechanism Γ guarantees expected revenue R in dominant strategy if Γ admits a dominant-strategy equilibrium $\sigma=(\sigma_i)_{i=1}^N$ such that

$$\sum_v [\sum_i p_i(\sigma(v))]g(v)\ge R.$$

Let R^D denote the best revenue guarantee in dominant strategy. That is, for any $R<R^D$, there is a mechanism which guarantees R in dominant strategy.

3.3 Bayesian auction mechanism

In order to define the other standard concept of Bayesian equilibrium, we need further information about the bidders' higher-order beliefs (such as what each bidder believes about the others' values, and about the others' beliefs about it, etc.). In this paper, we consider the simplest possible alternative: Each bidder i believes h_i as his first-order belief, and that fact itself is common knowledge (i.e., trivial higher-order beliefs).[6]

[6]Formally, the type space we consider in Definition 11 is in the class of the known-own-payoff-type type space (Bergemann and Morris (2005)), denoted by $(T_i,\hat{v}_i,\hat{\beta}_i)_{i=1}^N$. For each i, let (i) $T_i=\text{supp}\{v_i|\exists v_{-i};\ h_i(v_i,v_{-i})>0\}$, (ii)

Given that $(h_i)_{i=1}^N$ is common knowledge among the bidders, a Bayesian equilibrium in a mechanism is naturally defined as follows.

Definition 11. *Mechanism* Γ *admits a Bayesian equilibrium given* $(h_i)_{i=1}^N$ *if there exists* $\sigma_i(v_i)$ *for each* $i, v_i \in \mathbb{R}$ *such that, for each* m_i*:*

$$\begin{aligned} & \sum_{v_{-i}}[v_i q_i(\sigma_i(v_i), \sigma_{-i}(v_{-i})) - p_i(\sigma_i(v_i), \sigma_{-i}(v_{-i}))] h_i(v_i, v_{-i}) \\ \geq & \sum_{v_{-i}}[v_i q_i(m_i, \sigma_{-i}(v_{-i})) - p_i(m_i, \sigma_{-i}(v_{-i}))] h_i(v_i, v_{-i}). \end{aligned}$$

Mechanism Γ *guarantees expected revenue* R_ε *in Bayesian equlibrium if, for any* $(h_i)_{i=1}^N$ *such that* $(g, (h_i)_{i=1}^N)$ *are* ε*-close to each other,* Γ *admits a Bayesian equilibrium* $\sigma = (\sigma_i)_{i=1}^N$ *given* $(h_i)_{i=1}^N$ *such that*

$$\sum_v [\sum_i p_i(\sigma(v))] g(v) \geq R.$$

Let $R_\varepsilon^\star$ denote the best revenue guarantee in Bayesian equilibrium. That is, for any $R_\varepsilon < R_\varepsilon^\star$, there is a mechanism which guarantees R_ε in Bayesian equilibrium.

Obviously, the best revenue guarantee in dominant strategy is weakly lower than that in Bayesian equilibrium: For any ε,

$$R_\varepsilon^\star \geq R^D.$$

Recall that, in case g exhibits certain correlation (as specified in Crémer and McLean (1988*a*)) and $\varepsilon = 0$, the expected revenue in Bayesian equilibrium is very different from that in dominant strategies (i.e., $R_0^\star > R^D$). On the contrary, we show that, as long as ε is strictly positive, no matter how small it is, the guaranteed revenue in Bayesian equilibrium coincides with that in dominant strategies (i.e., $R_\varepsilon^\star = R^D$).

$\hat{v}_i : T_i \to \mathbb{R}$ be an identity map (i.e., $\hat{v}_i(t_i) = t_i$ for all $t_i \in T_i$), and (iii) $\hat{\beta}_i : T_i \to \Delta(T_{-i})$ is consistent with h_i in the sense that:

$$\hat{\beta}_i(t_{-i}|t_i) = \frac{h_i(t_i, t_{-i})}{\sum_{t'_{-i}} h_i(t_i, t'_{-i})}.$$

Our modelling choice may be justified as follows. First, even if one prefers other specifications, they would probably include this common-knowledge possibility as one of the possible situations; Second, as a related point, our approach would make the departure from the standard exact-common-knowledge model minimal. Given that our result is basically a negative result, this minimalistic choice makes the conclusion strongest.

We prove this claim in Section 5, followed by a motivating example in Section 4, explaining why the problem with $\varepsilon > 0$ can be very different from that with $\varepsilon = 0$.

4 MOTIVATING EXAMPLE

We employ Example 1 to illustrate the seller's revenue loss if he adopts the optimal mechanism *without* taking into account the possibility of prior heterogeneity. More precisely, imagine that the seller wrongly assumes that g is the common prior, while each bidder i actually has a different prior $h_i \neq g$. We will show that the seller's revenue loss does not vanish even when g and each h_i get closer in the sense of our distance.

Assume that the seller's benchmark distribution g is as illustrated in Table 3.5. If the seller believes that g is the common prior, then as in Crémer and McLean (1988*a*), the optimal mechanism is a combination of a second-price auction (SPA) and side-bets, which extracts the full surplus as his expected revenue ($\frac{5}{3}$). The following table corresponds to one such mechanism (it only shows bidder 1's allocation; bidder 2's is symmetric), where "NP" stands for non-participation:

$(q_1(v), t_1(v))$	NP	$v_2 = 1$	$v_2 = 2$
NP	$(0,0)$	$(0,0)$	$(0,0)$
$v_1 = 1$	$(1,0)$	$(\frac{1}{2}, \frac{1}{2} - \frac{1}{3})$	$(0; \frac{2}{3})$
$v_1 = 2$	$(1,0)$	$(1, 1 - \frac{1}{3})$	$(\frac{1}{2}, 1 + \frac{2}{3})$

Table 3.3: Outcomes from a SPA and side-bets

where the red parts in the transfers come from the side-bets. Each bidder's expected payment is $\frac{1}{3}\left(\frac{1}{2} - \frac{1}{3} + 1 + \frac{2}{3}\right) + \frac{1}{6}\left(1 - \frac{1}{3} + \frac{2}{3}\right) = \frac{5}{6}$, and therefore, the expected revenue is $\frac{5}{3}$, which is exactly the ex-ante total surplus.

Now consider the case where each bidder i's prior h_i is ε-close to but different from g. One might conjecture that, if the above mechanism is appropriately perturbed so that the bidders' participation and incentive constraints are satisfied *with strict inequality* (more specifically, with the strictness in the order of ε), then a similar level of expected revenue may be guaranteed. In particular, as $\varepsilon \to 0$, that guaranteed revenue converges to the full-surplus revenue again.

This conjecture is false. To explain the key idea, suppose that each h_i coincides with f in Table 3.2, while g is, as assumed by the seller, the true distribution of values. Even though f and g are ε-close to each other as *priors*, they are very different in terms of their induced conditional distributions, that is, each bidder's *interim* belief given his value. Given $f(= h_i)$, bidder i with any v_i essentially *knows* the other bidder's value. Therefore, in the above Crémer-McLean mechanism, truth-telling (or more precisely, reporting the values closest to their true values) is no longer an equilibrium.

For example, bidder i with $v_i = 1$ puts probability 1 on bidder $-i$'s having $v_{-i} = 2-\varepsilon$, and vice versa. They play a " complete-information" equilibrium where bidder $-i$ bids 2 and *bidder i does not participate in the auction.* Similarly, bidder i with $v_i = 2$, putting probability 1 on $v_{-i} = 2$, does not participate in the auction either. Therefore, as long as g is the true distribution (which only assigns positive probabilities on $v_i \in \{1,2\}$), no one participates in the auction, yielding 0 revenue.

Note that this property does not depend on the exact value of $\varepsilon > 0$. Therefore, the seller's expected revenue in this mechanism would be far below the first-best surplus.

5 MAIN RESULT

In this section, we show that $R^D = R^\star_\varepsilon$ for any $\varepsilon > 0$.

Theorem 5. *For any $\varepsilon > 0$, we have:*

$$R^D = R^\star_\varepsilon.$$

The proof is in Appendix A, and proceeds as follows. The key intuition is that, even if $\varepsilon(> 0)$ is arbitrarily small, it is always possible to find a specific prior h_i of each bidder i such that, after Bayesian updating observing i's own value v_i, his "interim belief" about the others' values is very different from the one where i's prior is g (i.e., the case with $\varepsilon = 0$). Moreover, this interim belief structure is such that the seller finds it optimal to offer a dominant-strategy auction mechanism *even if he knows that that h_i is each bidder's prior.* This last property is building on the original work by Chung and Ely (2007), while our more concise proof is building on Chen and Li (2018).[7]

Recall that the original result of Chung and Ely (2007) shows that an auction seller finds it optimal to offer a dominant-strategy auction mechanism to bidders if the seller has very little idea

[7] Specifically, the rationalizing interim beliefs are constructed such that each binding constraint in the seller's problem under the class of Bayesian-strategy mechanisms is a weighted sum of binding constraints under dominant-strategy mechanisms. Moreover, the weights coincide with the optimal Lagrangian multipliers for the latter, resulting in the two problems sharing the same value.

as to the bidders' belief structure, and hence any interim belief structure is deemed possible. Our result suggests that their result is relevant not only when the seller literally has very little idea about the bidders' information, but also when the seller and bidders have close (but heterogeneous) priors.

6 POSSIBLE MISSPECIFICATION OF g

So far, the seller assumes that his prior g is the true distribution of the bidders' valuations, although he thinks it possible that the bidders' priors are "ε-different" from g. However, if we interpret this ε as the seller's level of confidence in his g, it may also be natural to allow for the seller to worry about the possibility that g is not the true distribution of valuations.

Formally, let f represent the true distribution of v, the bidders' value profile. The seller does not know f, while he thinks that his prior g is a reasonable approximation of f (and each bidder's prior h_i). Based on the idea that $\varepsilon(>0)$ represents the seller's confidence in his g, we assume that $(f, g, (h_i)_{i=1}^N)$ are ε-close to each other.

To explain the subtlety, consider the optimal dominant-strategy mechanism if g is indeed the true prior (which guarantees R^D). Typically, some incentive compatibility constraints are binding in this mechanism. Thus, if $f(\neq g)$ is the true prior with $\text{supp}(f) \neq \text{supp}(g)$, some bidders may find it strictly optimal to make a type report that is far from his true type.[8]

Nevertheless, we show that an appropriately modified version of the mechanism, which we call a *transfer-reducing mechanism*, guarantees the same level of expected revenue even if $f \neq g$, as ε vanishes. The key of the construction is, by reducing the transfers of the mechanism (by an appropriate amount as a function of ε), the mechanism can now make all the incentive constraints satisfied in a stronger sense, so that even if g and f have (ε-)different supports, each agent finds it dominant to report the value that is closest to his true value in that mechanism. Although the revenue must be smaller, as $\varepsilon \to 0$, this revenue loss vanishes.

Definition 12. *Mechanism Γ guarantees expected revenue R_ε in dominant-strategy equilibrium in g's ε-neighborhood if, for any f that is ε-close to g, Γ admits a dominant-strategy equilibrium $\sigma = (\sigma_i)_{i=1}^N$ such that*

$$\sum_v [\sum_i p_i(\sigma(v))] f(v) \geq R_\varepsilon.$$

[8] In Appendix B, we observe that such a global deviation under misspecification is the norm rather than the exception.

Let R^D_ε represent the optimal revenue guarantee in dominant strategy if $(f, g, (h_i)_{i=1}^N)$ are ε-close to each other.

Theorem 6. $R^D_\varepsilon \to R^D$ *as* $\varepsilon \to 0$.

Madarász and Prat (2017) show that, in a general single-agent mechanism design environment, a similar approximation result is possible by their *profit-participation mechanism* even without single-crossing conditions. That is, as the seller's benchmark distribution converges to the true distribution, their optimal expected revenues also converge. Its basic idea is to make the agent "biased in favor of the principal" so that any (even non-local) deviation due to misspecification only increases the principal's payoff. Our proof generalizes their result to a multi-agent environment, but in a single-crossing payoff environment. The single-crossing property seems important for this continuity result with multiple agents. To explain this, it is worth noting that (a naive adaptation of) their profit-participating mechanism may not work in our multi-agent setup. This is because, under that mechanism, each agent might have an incentive to deviate globally (i.e., a value far from his true one is reported), which in turn distorts other agents' reporting strategies. Consequently, it is not certain that the vanishing revenue loss is obtained. Our transfer-reducing mechanism prevents such global deviations by ensuring that it is a dominant strategy for agents to report the value closest to their true values.

7 CONCLUSION

In this paper, we consider the private-value auction setting where the true distribution of bidders' valuations is unknown. The seller and each bidder, however, know its approximation. In this framework, we have shown that the dominant-strategy mechanism secures the seller with the highest revenue guarantee. Besides, if the seller is restricted to using a dominant-strategy mechanism, we have characterized the transfer reducing mechanism that helps the seller to obtain a vanishing loss as the estimates by her and the bidders get close to the truth.

There are several follow-up questions. Firstly, when restricting to dominant-strategy mechanisms, our proof works only if the bidders' payoff functions satisfy the single-crossing condition. Although this property holds for a wide range of mechanism design problems, there are cases where it does not hold, such as multi-unit auctions. In such situations, our proposed mechanism may not work.

Another natural direction is to characterize the optimal robust mechanisms in non-auction environments[9] or with common/interdependent values.[10] We leave these potential extensions for future research.

[9] Chen and Li (2018) generalize the foundation result of Chung and Ely (2007) to some private-value non-auction environments. We conjecture that our approach would work in those environments, establishing the worst-case optimality of dominant-strategy mechanisms.

[10] Yamashita and Zhu (2022) generalize the foundation result of Chung and Ely (2007) to an interdependent-value auction environment. We conjecture that our approach would work in those environments. However, they suggest that general interdependent-value models may not admit the same sort of foundation result, and in those cases, it is an open question how the approximate worst-case optimal mechanism would look like.

A PROOF OF THEOREM 5

We construct each bidder i's prior, h_i, as follows.

For each i, let $V_i = \{v_i \mid \exists v_{-i},\ g(v_i, v_{-i}) > 0\}$ denote the set of i's possible values in the true distribution g, and denote it by $V_i = \{v_i^1, \ldots, v_i^m, \ldots, v_i^M\}$ so that $v_i^m < v_i^{m+1}$. Define $\hat{V}_i = \{v_i + \varepsilon \mid v_i \in V_i\}$ as the "shifted" version of V_i by ε. Also, we denote $V_{-i} = \{v_{-i}^1, \ldots, v_{-i}^l, \ldots, v_{-i}^L\}$, without any ordering on them (i.e., arbitrary labelling will do). Define $h_i(\cdot)$ so that: for each m, l,

$$h_i(v_i^m, v_{-i}^l) = x_i \frac{\tau_i^\star(v_{-i}^l | v_i^m)}{\tau_i^\star(v_{-i}^1 | v_i^m)}$$

(recall $v_i^m \in V_i$), and

$$h_i(v_i^m + \varepsilon, v_{-i}^l) = g(v_i^m, v_{-i}^l) - h_i(v_i^m, v_{-i}^l)$$

(recall $v_i^m + \varepsilon \in \hat{V}_i$), where

$$\tau_i^\star(v_{-i} | v_i) \equiv \frac{\sum_{\hat{v}_i \geq v_i} g(\hat{v}_i, v_{-i})}{\sum_{v_{-i}} \sum_{\hat{v}_i \geq v_i} g(\hat{v}_i, v_{-i})}, \quad x_i = \min_{k,m} \frac{\tau_i^\star(v_{-i}^1 | v_i^m)}{\tau_i^\star(v_{-i}^l | v_i^m)} g(v_i^m, v_{-i}^l).$$

The following table illustrates our construction:

$h_i(.,.)$	v^1_{-i}	v^2_{-i}	...	v^L_{-i}
v_i^1	x_i	$x_i \dfrac{\tau_i^\star(v^2_{-i} \mid v_i^1)}{\tau_1^\star(v^1_{-i} \mid v_i^1)}$	...	$x_i \dfrac{\tau_i^\star(v^L_{-i} \mid v_i^1)}{\tau_1^\star(v^1_{-i} \mid v_i^1)}$
$v_i^1+\varepsilon$	$g(v_i^1, v^1_{-i}) - x_i$	$g(v_i^1, v^2_{-i}) - x_i \dfrac{\tau_i^\star(v^2_{-i} \mid v_i^1)}{\tau_i^\star(v^1_{-i} \mid v_i^1)}$	...	$g(v_i^1, v^L_{-i}) - x_i \dfrac{\tau_i^\star(v^L_{-i} \mid v_i^1)}{\tau_1^\star(v^1_{-i} \mid v_i^1)}$
v_i^2	x_i	$x_i \dfrac{\tau_i^\star(v^2_{-i} \mid v_i^2)}{\tau_i^\star(v^1_{-i} \mid v_i^2)}$	...	$x_i \dfrac{\tau_i^\star(v^L_{-i} \mid v_i^2)}{\tau_i^\star(v^1_{-i} \mid v_i^2)}$
$v_i^2+\varepsilon$	$g(v_i^2, v^1_{-i}) - x_i$	$g(v_i^2, v^2_{-i}) - x_i \dfrac{\tau_i^\star(v^2_{-i} \mid v_i^2)}{\tau_i^\star(v^1_{-i} \mid v_i^2)}$	...	$g(v_i^2, v^L_{-i}) - x_i \dfrac{\tau_1^\star(v^L_{-i} \mid v_i^2)}{\tau_i^\star(v^1_{-i} \mid v_i^2)}$
...	...	...	...	...
v_i^M	x_i	$x_i \dfrac{\tau_i^\star(v^2_{-i} \mid v_i^M)}{\tau_i^\star(v^1_{-i} \mid v_i^M)}$	...	$x_i \dfrac{\tau_i^\star(v^L_{-i} \mid v_i^M)}{\tau_i^\star(v^1_{-i} \mid v_i^M)}$
$v_i^M+\varepsilon$	$g(v_i^M, v^1_{-i}) - x_i$	$g(v_i^M, v^2_{-i}) - x_i \dfrac{\tau_i^\star(v^2_{-i} \mid v_i^M)}{\tau_i^\star(v^1_{-i} \mid v_i^M)}$	...	$g(v_i^M, v^L_{-i}) - x_i \dfrac{\tau_i^\star(v^L_{-i} \mid v_i^M)}{\tau_i^\star(v^1_{-i} \mid v_i^M)}$

First, the choice of x_i guarantees that $h_i(v) \geq 0$ for all $v \in (V_i \cup \hat{V}_i) \times V_{-i}$. It is also immediate that h_i and g are ε-close to each other, because $h_i(v_i^m + \varepsilon, v^l_{-i}) = g(v_i^m, v^l_{-i}) - h_i(v_i^m, v^l_{-i})$, and from this equation, we can also easily see that $\sum_v h_i(v) = 1$.

Next, we show that under this construction of bidders' beliefs, $\bar{R}^\star_\varepsilon \leq R^D$. This, combined with the fact that $\bar{R}^\star_\varepsilon \geq R^D$ completes the proof. Note that under V, the shifted valuations $\{v_i + \varepsilon\}_{v_i \in V}$ are never realized. Therefore, the seller's problem under the class of dominant-strategy mechanisms is defined entirely on V, as follows:

$$(P^D) \qquad \bar{R}^D = \sup_{(q,p)} \quad \mathbb{E}_{v\sim g}[\sum_i p_i(v)] \equiv \sum_{v\in V}\sum_i p_i(v)g(v)$$

$$s.t. \quad \forall i, \quad \forall v_i, v_i' \in V_i, \quad \forall v_{-i} \in V_{-i}:$$

$$v_i q_i(v) - p_i(v) \geq 0$$

$$v_i q_i(v) - p_i(v) \geq v_i q_i(v_i', v_{-i}) - p_i(v_i', v_{-i})$$

$$q_i(v) \geq 0; \quad \sum_i q_i(v) \leq 1$$

Let $\{q^D(v), p^D(v)\}_v$ denote the solution of (P^D). By a standard result, only local downward IC constraints and IR constraints for the lowest type bind. Therefore, there exist multipliers $\{\lambda_i^D(v)\}_v$ associated with those constraints such that $\{q^D(v), p^D(v), \lambda_i^D(v)\}_v$ maximizes the following Lagrangian function:

$$\mathscr{L}^D \equiv \sum_{i,v} p_i(v)g(v) + \sum_{i,v_{-i}} \lambda_i^D(v_i^1, v_{-i})\left[v_i^1 q_i(v_i^1, v_{-i}) - p_i(v_i^1, v_{-i})\right]$$
$$+ \sum_{i,v_{-i}} \sum_{v_i^m \geq v_i^2} \lambda_i^D(v_i^m, v_{-i})\Big[\left[v_i^m q_i(v_i^m, v_{-i}) - p_i(v_i^m, v_{-i})\right] - \left[v_i^m q_i(v_i^{m-1}, v_{-i}) - p_i(v_i^{m-1}, v_{-i})\right]\Big]$$

over the domain $(q,p) \in \mathbb{Q} \times \mathbb{R}$ where $\mathbb{Q} \equiv \{q_i(v) \geq 0; \sum_i q_i(v) \leq 1\}$.

Note that there are no restrictions imposed on payments. Therefore, at optimum:

$$\frac{\partial \mathscr{L}^D}{\partial p(v_i^M, v_{-i})} = 0 \Leftrightarrow \lambda_i^D(v_i^M, v_{-i}) = g(v_i^M, v_{-i}),$$

$$\frac{\partial \mathscr{L}^D}{\partial p(v_i^m, v_{-i})} = 0 \Leftrightarrow \lambda_i^D(v_i^m, v_{-i}) = \lambda_i^D(v_i^{m+1}, v_{-i}) + g(v_i^m, v_{-i}) \quad \forall 1 \leq m < M$$

Thus, we have for all (v_i^m, v_{-i}):

$$\lambda_i^D(v_i^m, v_{-i}) = \sum_{\hat{v}_i \geq v_i^m} g(\hat{v}_i, v_{-i}) \tag{3.1}$$

Similarly, the seller's problem under the class of Bayesian-strategy mechanisms is also defined entirely on V, as follows:

$$
\begin{aligned}
(P^B) \quad \bar{R}^\star_\varepsilon = \sup_{(q,p)} \quad & \mathbb{E}_{v\sim g}[\sum_i p_i(v)] \equiv \sum_{v\in V}\sum_i p_i(v)g(v) \\
s.t. \quad & \forall i, \quad \forall v_i, v_i' \in V_i, \quad \forall v_{-i} \in V_{-i}: \\
& \sum_{v_{-i}} h_i(v_{-i}|v_i)[v_i q_i(v) - p_i(v)] \ge 0 \\
& \sum_{v_{-i}} h_i(v_{-i}|v_i)[v_i q_i(v) - p_i(v)] \ge \sum_{v_{-i}} h_i(v_{-i}|v_i)[v_i q_i(v_i', v_{-i}) - p_i(v_i', v_{-i})] \\
& q_i(v) \ge 0; \quad \sum_i q_i(v) \le 1
\end{aligned}
$$

Consider its relaxed problem, denoted by (RP^B) where only local downward IC constraints and IR constraints for the lowest type are considered and hence, must be binding.

Recall our construction for i's prior:

$$
h_i(v_{-i}|v_i^m) = \frac{h_i(v_{-i}, v_i^m)}{\sum_{v_{-i}} h_i(v_{-i}, v_i^m)} \equiv \frac{x_i \dfrac{\tau_i^\star(v_{-i}^l|v_i^m)}{\tau_i^\star(v_{-i}^1|v_i^m)}}{\sum_{v_{-i}} x_i \dfrac{\tau_i^\star(v_{-i}^l|v_i^m)}{\tau_i^\star(v_{-i}^1|v_i^m)}} = \tau_i^\star(v_{-i}^l|v_i^m) = \frac{\sum_{\hat{v}_i \ge v_i^m} g(\hat{v}_i, v_{-i})}{\sum_{v_{-i}} \sum_{\hat{v}_i \ge v_i^m} g(\hat{v}_i, v_{-i})} \tag{3.2}
$$

(3.1) and (3.2) imply that:

$$
h_i(v_{-i}|v_i^m) = \frac{\lambda_i^D(v_i^m, v_{-i})}{\sum_{v_{-i}} \lambda_i^D(v_i^m, v_{-i})}
$$

Therefore, each (binding) constraint in (RP^B) is a weighted sum of (binding) constraints in (P^D), with the weight being the corresponding optimal Lagrangian multiplier for the latter. Then, it can be verified that these two problems have the same values. Note that the value of (RP^B) is obviously an upper bound of that under the original problem (P^B). Hence, we obtain: $\bar{R}^\star_\varepsilon \le R^D$. Therefore, $\bar{R}^\star_\varepsilon = R^D$.

B GLOBAL DEVIATION IN OPTIMAL DOMINANT-STRATEGY MECHANISMS WITH A MISSPECIFIED SUPPORT

Recall the standard properties of the optimal dominant-strategy mechanism under the assumption that $g = f$:

1. All the local downward IC constraints bind, i.e, for any $k \geq 2$, any $\hat{v}_{-i}$, and any $s^k \in supp(S)$:

$$s^k q(s^k, \hat{v}_{-i}) - p(s^k, \hat{v}_{-i}) = s^k q(s^{k-1}, \hat{v}_{-i}) - p(s^{k-1}, \hat{v}_{-i})$$

where $s^{k-1} = \max\{s \in supp(S) \mid s < s^k\}$.

2. Allocation is monotone, i.e., $q_i(s^k, v_{-i}) \leq q_i(s^{k'}, v_{-i})$ if $k < k'$.

If it is possible that f and g are (ε-close to but) different from each other, then global deviations would typically be relevant.

Proposition 10. *Fix ε and g. In the optimal dominant-strategy mechanism assuming g is the true prior, there exists f that is ε-close to g such that, if $v \sim f$, then a bidder does not find it optimal to report the value that is closest to his true valuation.*

Proof. Let $S = \text{supp}(g)$. Let f be such that some i's value $v_i = s^k - \varepsilon$ is supported. Then, he prefers reporting s^{k-1} to reporting s^k, even though s^k is closer to v_i than s^{k-1}. This is because:

$$\begin{aligned}(s^k - \varepsilon)q(s^k, \hat{v}_{-i}) - p(s^k, \hat{v}_{-i}) &= s^k q(s^k, \hat{v}_{-i}) - p(s^k, \hat{v}_{-i}) - \varepsilon q(s^k, \hat{v}_{-i}) \\ &\leq s^k q(s^{k-1}, \hat{v}_{-i}) - p(s^{k-1}, \hat{v}_{-i}) - \varepsilon q(s^{k-1}, \hat{v}_{-i}) \\ &= (s^k - \varepsilon)q(s^{k-1}, \hat{v}_{-i}) - p(s^{k-1}, \hat{v}_{-i})\end{aligned}$$

where the inequality follows from the local $\text{IC}_{k,k-1}$ constraint and the monotonicity constraint. Moreover, if $q(s^k, \hat{v}_{-i}) > q(s^{k-1}, \hat{v}_{-i})$, the inequality then becomes strict. That is:

$$(s^k - \varepsilon)q(s^k, \hat{v}_{-i}) - p(s^k, \hat{v}_{-i}) < (s^k - \varepsilon)q(s^{k-1}, \hat{v}_{-i}) - p(s^{k-1}, \hat{v}_{-i})$$

Consequently, the agent whose value is $v_i = s^k - \varepsilon$ strictly prefers to report his valuation as s^{k-1} instead of his closest type s^k. □

C PROOF OF THEOREM 6

Let $(q^\star(\cdot), p^\star(\cdot))$ represent the optimal dominant-strategy mechanism under the assumption that g is the true prior. Let $V = \text{supp}(g)$, and let $V_i = \{v_i \in \mathbb{R} \mid \exists v_{-i};\ (v_i, v_{-i}) \in V\}$ denote its i-th coordinate. We also denote $v_i^+ \equiv \min\{s \in V_i \mid s > v_i\}$ and $v_i^- \equiv \max\{s \in V_i \mid s < v_i\}$.

Fix $\delta > 0$, which is sufficiently small. The *δ-transfer reduction mechanism* of $(q^\star(\cdot), p^\star(\cdot))$ has the same message space and the winning-probability function as the optimal dominant-strategy mechanism, but the price is smaller by δ.

For each $v_i \in V_i$, truth-telling is still dominant-strategy incentive compatible, but now in a stronger sense: for bidder i whose value is δ-close to $v_i \in V_i$, it is dominant for him to report v_i in the δ-transfer-reduction mechanism. Note that under the original mechanism, for all $\sigma(v_{-i})$ and $v_i' < v_i$, we have $q_i^\star(v_i, \sigma(v_{-i})) \geq q_i^\star(v_i', \sigma(v_{-i}))$ and:

$$v_i q_i^\star(v_i, \sigma(v_{-i})) - p_i^\star(v_i, \sigma(v_{-i})) \geq v_i q_i^\star(v_i', \sigma(v_{-i})) - p_i^\star(v_i', \sigma(v_{-i}))$$

which means:

$$\begin{aligned}
&(v_i - \delta) q_i^\star(v_i, \sigma(v_{-i})) - [p_i^\star(v_i, \sigma(v_{-i})) - \delta q_i^\star(v_i, \sigma(v_{-i}))] \\
\geq &(v_i - \delta) q_i^\star(v_i', \sigma(v_{-i})) - [p_i^\star(v_i', \sigma(v_{-i})) - \delta q_i^\star(v_i', \sigma(v_{-i}))]
\end{aligned}$$

By single crossing property and $\hat{v}_i \geq v_i - \delta$ ($\hat{v}_i$ is δ-close to $v_i \in V_i$), we thus have:

$$\hat{v}_i q_i^\star(v_i, \sigma(v_{-i})) - [p_i^\star(v_i, \sigma(v_{-i})) - \delta q_i^\star(v_i, \sigma(v_{-i}))] \geq \hat{v}_i q_i^\star(v_i', \sigma(v_{-i})) - [p_i^\star(v_i', \sigma(v_{-i})) - \delta q_i^\star(v_i', \sigma(v_{-i}))]$$

for all $\sigma(v_{-i})$ and $v_i' < v_i$, i.e.,

$$v_i = \operatorname*{argmax}_{v_i' \leq v_i} \Big[\hat{v}_i q_i^\star(v_i', \sigma(v_{-i})) - [p_i^\star(v_i', \sigma(v_{-i})) - \delta q_i^\star(v_i', \sigma(v_{-i}))] \Big] \quad \forall \sigma(v_{-i}) \tag{3.3}$$

Note also that under the original mechanism, for all $\sigma(v_{-i})$ and $v_i' > v_i^+$, we have $q_i^\star(v_i', \sigma(v_{-i})) \geq q_i^\star(v_i^+, \sigma(v_{-i})) \geq q_i^\star(v_i, \sigma(v_{-i}))$, and:

$$v_i^+ q_i^\star(v_i^+, \sigma(v_{-i})) - p_i^\star(v_i^+, \sigma(v_{-i})) = v_i^+ q_i^\star(v_i, \sigma(v_{-i})) - p_i^\star(v_i, \sigma(v_{-i}))$$

which means:

$$\begin{aligned}
&(v_i^+ - \delta) q_i^\star(v_i^+, \sigma(v_{-i})) - [p_i^\star(v_i^+, \sigma(v_{-i})) - \delta q_i^\star(v_i^+, \sigma(v_{-i}))] \\
= &(v_i^+ - \delta) q_i^\star(v_i, \sigma(v_{-i})) - [p_i^\star(v_i, \sigma(v_{-i})) - \delta q_i^\star(v_i, \sigma(v_{-i}))]
\end{aligned}$$

By single crossing property and $v_i^+ - \delta \geq \hat{v}_i$ ($\hat{v}_i$ is δ-close to $v_i \in V_i$), this implies that for all $\sigma(v_{-i})$:

$$\begin{aligned}
&\hat{v}_i q_i^\star(v_i^+, \sigma(v_{-i})) - [p_i^\star(v_i^+, \sigma(v_{-i})) - \delta q_i^\star(v_i^+, \sigma(v_{-i}))] \\
\leq &\hat{v}_i q_i^\star(v_i, \sigma(v_{-i})) - [p_i^\star(v_i, \sigma(v_{-i})) - \delta q_i^\star(v_i, \sigma(v_{-i}))]
\end{aligned} \tag{3.4}$$

Moreover, for all $\sigma(v_{-i})$ and $v_i' > v_i^+$:

$$v_i^+ q_i^\star(v_i', \sigma(v_{-i})) - p_i^\star(v_i', \sigma(v_{-i})) \leq v_i^+ q_i^\star(v_i^+, \sigma(v_{-i})) - p_i^\star(v_i^+, \sigma(v_{-i}))$$

which means:

$$
\begin{aligned}
&(v_i^+ - \delta) q_i^\star(v_i', \sigma(v_{-i})) - [p_i^\star(v_i', \sigma(v_{-i})) - \delta q_i^\star(v_i', \sigma(v_{-i}))] \\
\leq &(v_i^+ - \delta) q_i^\star(v_i^+, \sigma(v_{-i})) - [p_i^\star(v_i^+, \sigma(v_{-i})) - \delta q_i^\star(v_i^+, \sigma(v_{-i}))]
\end{aligned}
$$

By single crossing property and $v_i^+ - \delta \geq \hat{v}_i$ ($\hat{v}_i$ is δ-close to $v_i \in V_i$), we thus have for all $\sigma(v_{-i})$ and $v_i' > v_i^+$:

$$
\begin{aligned}
&\hat{v}_i q_i^\star(v_i', \sigma(v_{-i})) - [p_i^\star(v_i', \sigma(v_{-i})) - \delta q_i^\star(v_i', \sigma(v_{-i}))] \\
\leq &\hat{v}_i q_i^\star(v_i^+, \sigma(v_{-i})) - [p_i^\star(v_i^+, \sigma(v_{-i})) - \delta q_i^\star(v_i^+, \sigma(v_{-i}))]
\end{aligned}
\tag{3.5}
$$

Combining (3.4) and (3.5), we obtain:

$$
v_i = \underset{v_i' \geq v_i}{\operatorname{argmax}} \Big[\hat{v}_i q_i^\star(v_i', \sigma(v_{-i})) - [p_i^\star(v_i', \sigma(v_{-i})) - \delta q_i^\star(v_i', \sigma(v_{-i}))] \Big] \quad \forall \sigma(v_{-i}) \tag{3.6}
$$

Then, (3.3) and (3.6) imply that for bidder i whose value is δ-close to $v_i \in V_i$, it is dominant for him to report v_i in the δ-transfer-reduction mechanism. We take $\delta = \varepsilon$ then. Although we omit the details, it can also be shown that his ex post individual rationality is satisfied.

By construction, the ε-transfer-reduction mechanism collects the same amount of transfer from each type of each agent less at most ε. Therefore, *if* $g = f$, then the expected revenue in the ε-transfer-reduction mechanism, denoted by $R'_\varepsilon(g)$, is not lower than $R^D - N\varepsilon$:

$$
R'_\varepsilon(g) \geq R^D - N\varepsilon.
$$

Even if f is different from g, it remains true that each bidder with each type finds it dominant to report his closest type in the same ε-transfer-reduction mechanism. Therefore, denoting by $R'_\varepsilon(f)$ the expected revenue of the same ε-transfer-reduction mechanism but with distribution f, by continuity we obtain:

$$
\lim_{\varepsilon \to 0} |R'_\varepsilon(g) - R'_\varepsilon(f)| = 0,
$$

and therefore:

$$
\lim_{\varepsilon \to 0} |R'_\varepsilon(g) - \inf_{f|\varepsilon\text{-close to } g} R'_\varepsilon(f)| = 0,
$$

By Theorem 1:

$$
R^D = R^\star_\varepsilon \geq \inf_{f|\varepsilon\text{-close to } g} R'_\varepsilon(f)
$$

Therefore:

$$0 \le R^D - \inf_{f|\varepsilon\text{-close to } g} R'_\varepsilon(f) \le R'_\varepsilon(g) + N\varepsilon - \inf_{f|\varepsilon\text{-close to } g} R'_\varepsilon(f),$$

where the right-hand side converges to 0 as $\varepsilon \to 0$, implying:

$$\inf_{f|\varepsilon\text{-close to } g} R'_\varepsilon(f) \to R^D,$$

as $\varepsilon \to 0$. We complete the proof by noticing that $R^D_\varepsilon \in [\inf_{f|\varepsilon\text{-close to } g} R'_\varepsilon(f), R^D]$.

D A COUNTEREXAMPLE WITH SUFFICIENTLY CLOSE INTERIM BELIEFS

This section provides a simple example in which (i) interim beliefs induced by the seller's and bidders' priors are close enough to each other, and (ii) dominant-strategy mechanisms are not Bayesian founded.

Example 11. *There are two bidders $i \in \{1,2\}$. The seller knows the true prior. Her prior $g(v_1, v_2)$ is given by*

$g(v_1, v_2)$	$v_2 = 1$	$v_2 = 2$
$v_1 = 1$	$\frac{1}{3}$	$\frac{1}{6}$
$v_1 = 2$	$\frac{1}{6}$	$\frac{1}{3}$

Table 3.4: Seller's prior

Thus, g induces the following interim belief

$g(v_{-i} \mid v_i)$	$v_{-i} = 1$	$v_{-i} = 2$
$v_i = 1$	$\frac{2}{3}$	$\frac{1}{3}$
$v_i = 2$	$\frac{1}{3}$	$\frac{2}{3}$

Table 3.5: Interim belief induced by g

Moreover, each bidder i's interim belief after observing his own valuation, $h_i(v_{-i} \mid v_i)$ is such that $h_i(v_{-i} \mid v_i) \in [g(v_{-i} \mid v_i) - \varepsilon, g(v_{-i} \mid v_i) + \varepsilon]$ for all i, v_i and v_{-i}, for some $\varepsilon > 0$.

In this example, if $\varepsilon = 0$, the seller can fully extract surplus by offering a SPA associated with side-bets (as explained in the motivating example).

Now we show that with small $\varepsilon > 0$, the seller can *almost* fully extract the surplus by reducing each bidder's expected payment under the side-bets by ε. The following table corresponds to this mechanism (it only shows bidder 1's allocation; bidder 2's is symmetric), where "NP" stands for non-participation:

$(q_1(v), t_1(v))$	NP	$v_2 = 1$	$v_2 = 2$
NP	$(0,0)$	$(0,0)$	$(0,0)$
$v_1 = 1$	$(1,0)$	$(\frac{1}{2}, \frac{1}{2} - \frac{1+3\varepsilon}{3})$	$(0; \frac{2-3\varepsilon}{3})$
$v_1 = 2$	$(1,0)$	$(1, 1 - \frac{1+3\varepsilon}{3})$	$(\frac{1}{2}, 1 + \frac{2-3\varepsilon}{3})$

Table 3.6: Outcomes from a SPA and side-bets

where the red parts in the transfers come from the side-bets. Let x_i denote bidder i's payment under side-bets. Let $x_i(v_{-i})$ denote bidder i's payment under the side-bets (the red parts in the table) when the other bidder's value is v_{-i}. Bidder i's conditional expected payments under the side bet (after observing v_i) are given by

$$\begin{aligned}
\mathbb{E}[x_i \mid v_i = 1] &= h_i(1 \mid 1)x_i(1) + h_i(2 \mid 1)x_i(2) \\
&\leq [g_i(1 \mid 1) - \varepsilon]x_i(1) + [g_i(2 \mid 1) + \varepsilon]x_i(2) \\
&= -(\frac{2}{3} - \varepsilon)\frac{1+3\varepsilon}{3} + (\frac{1}{3} + \varepsilon)\frac{2-3\varepsilon}{3} \\
&= 0, \\
\mathbb{E}[x_i \mid v_i = 2] &= h_i(1 \mid 2)x_i(1) + h_i(2 \mid 2)x_i(2) \\
&\leq [g_i(1 \mid 2) - \varepsilon]x_i(1) + [g_i(2 \mid 2) + \varepsilon]x_i(2) \\
&= -(\frac{1}{3} - \varepsilon)\frac{1+3\varepsilon}{3} + (\frac{2}{3} + \varepsilon)\frac{2-3\varepsilon}{3} \\
&= \frac{1}{3}.
\end{aligned}$$

Moreover, note that bidder i's interim expected payment, conditional on $v_i = 1$ (*resp.*, $v_i = 2$) from the SPA is given by 0 (*resp.*, $\frac{1}{3}$). Therefore, i's interim expected payoff from the SPA combined with the side-bets is non-negative. Consequently, it is optimal for bidders to report their value truthfully. As the reduction in expected payments by bidders is proportional to ε, the seller's revenue is close to the full surplus as ε approaches zero, which is strictly higher than that obtained from the best dominant-strategy mechanism.[11] Hence, we obtain no foundation for the use of a dominant-strategy mechanism in this example.

[11] A second-price auction (without a reserve price) is an optimal dominant-strategy mechanism, which generates a revenue of $\frac{4}{3}$ for the seller, whereas the full expected surplus is $\frac{5}{3}$.

Chapter 4

First Best Implementation with Costly Information Acquisition

1 INTRODUCTION

In most mechanism design problems, there is a collection of agents who have exogenously given private information, and there is a principal who desires to implement a social choice rule by designing a mechanism which incentivizes the agents to reveal their information.

In many practical problems, however, the agents' private information is often a consequence of their own (possibly costly) *information acquisition*. For example, bidders in an oil-tract auction (Wilson, 1969) may conduct test drills; bidders in a spectrum auction may conduct market research; voters in a presidential election may investigate the candidates' past political activities; members of a hiring committee may study the job applicant's background in order to see whether he is fit for the job.

Importantly, in such situations, a mechanism in place does not only affect each agent's incentive to report the acquired information truthfully, but also affects his choice of *what kind of information* to acquire. In this sense, the properties of desirable mechanisms could potentially be very different from those which only guarantee truth-telling incentives for a *given* information structure.

Although this issue is already relevant in single-agent environments,[1] the degree of complexity

[1]Mensch (2020) studies a mechanism design problem with a single agent. See also Section 1.1.

is even higher in multi-agent environments: in principle, flexibility of each agent's information acquisition does not only mean flexibility in terms of his signal's informativeness about the payoff-relevant state, but also means flexibility in terms of his signal's informativeness about his opponents' signals. This issue of higher-order information and beliefs distinguishes multi-agent from single-agent environments. Modeling the dependence of the cost of information acquisition on higher-order information is a challenging task. In this paper, we assume that an agent's information acquisition cost only depends on his signal's informativeness about the payoff-relevant state, but not about the other agents' signals. In particular, it is costless to acquire a signal that is independent from the payoff-relevant state.[2] For example, imagine a situation where agents (e.g., telecom companies who buy spectrum) have to acquire information from data providers (e.g., market research firms) operating on a competitive market for data. Each data provider generates signals about the payoff-relevant state (e.g., demand conditions in the mobile services market). Competition among the data providers forces them to price their data at the cost of production, which in turn depends on the informativeness of their data. The agents can then decide to make their signals perfectly correlated by strategically choosing the same data provider, or less than perfectly correlated by choosing different data providers. In both cases, the agents will pay the same price for the same informativeness, and hence the cost of information acquisition will be independent of the correlation structure among signals.[3]

We consider a model with four or more agents. The principal and the agents share a common prior about the payoff-relevant state, and none of them has any private information at the beginning. We show that there exists a mechanism which allows the principal to implement any social choice rule at zero information acquisition cost to the agents. The key idea is that the mechanism recommends each agent to choose a special information acquisition action, which satisfies the *individually-uninformative-but-aggregately-revealing* property of Zhu (2021) (and each agent finds it optimal to obey this recommended action). The *individually-uninformative* part means that each agent's signal on its own is independent from the payoff-relevant state, which guarantees that his information cost is zero. The *aggregately-revealing* part means that the principal, by observing all the agents' reports — in fact, any two of them — can correctly identify the true payoff-relevant state. The fact that only two are enough, together with the fact

[2] The literature on cost of information proposes and discusses a variety of possible cost functions (see Section 1.1), but it seems to be universally accepted that uninformative signals about the state of the world are costless.

[3] Of course, one can come up with cases where other cost specifications seem more reasonable (e.g., more positive correlation is more costly, or less costly). We discuss a range of possible alternative assumptions in our concluding remarks, see Section 5.

that there are four or more agents, enables the principal to detect any unilateral deviation. It thus establishes the incentive compatibility of the mechanism.

1.1 Related Literature

In the literature on information acquisition in mechanism design, we usually consider restricted and/or less flexible spaces of information (see, for example, Bergemann and Välimäki (2002) for efficient mechanism design, Shi (2012) and Bikhchandani and Obara (2017) for optimal auction design, and Persico (2004), Gerardi and Yariv (2008), Gershkov and Szentes (2009), and Zhao (2016) for committee design with information acquisition[4]).

Mensch (2020) studies mechanism design with a single agent's flexible and costly information acquisition, building on the rational inattention framework (Sims (2003)).[5] Flexible and costly information acquisition is also considered by Gleyze and Pernoud (2020) who study a mechanism design problem with transferable utility and private values, in which agents acquire costly information on their own preferences and the preferences of other agents, and by Ravid et al. (2020) who study a bilateral trade model with costly information acquisition by the buyer. Flexible but not costly information acquisition is considered by Roesler and Szentes (2017) in the context of buyer-optimal information in monopoly pricing,[6] by Bergemann et al. (2017) and Brooks and Du (2021) in the context of seller-pessimal information in common-value auctions, and by Yamashita (2018) in private-value auctions. All these papers feature a single entity, "nature", who chooses the information structure (of one or multiple agents). In contrast to that, in our model each agent acquires information in a decentralized manner, which leads to a very different conclusion.

The information structure we employ was proposed in the context of mechanism design by Zhu (2021), who studies information disclosure by a mechanism designer. It builds on the idea of the one-time pad, an unbreakable encryption method (Shannon, 1949).[7]

[4]Restricting to the class of conservative rules, Li (2001) solves for the optimal degree of conservatism in committee design. The optimally chosen conservative rule outperforms the *ex post* optimal rule.

[5]Mensch (2020) also considers a multiple-agent extension of his model, but restricts attention to symmetric mechanisms in an independent private values setting, in which agents can acquire information about their own values, but cannot acquire any information about others' values.

[6]See also Condorelli and Szentes (2020), though they also consider non-information changes of the agent's private information distribution.

[7]See also Krähmer (2020) and Krähmer (2021) in the context of information disclosure in mechanism design and strategic communication respectively where the randomization of information structures is allowed to keep

This key information structure makes the agents' acquired information statistically dependent. In quasi-linear environments, Crémer and McLean (1988*b*) show that the principal can extract full surplus from the agents who share a correlated prior. Although the extreme positivity of the results is a common feature of our paper and theirs, the two problems are quite different. First, our paper does not assume quasi-linearity. Second, their side-bet mechanism exploits an exogenously given correlated signal structure, and it is not clear if such a signal structure can be induced in equilibrium given some reasonable space of information acquisition actions.[8] In our case, the resulting information structure is an equilibrium outcome, even though each agent can potentially acquire information independently from the others' signals.

In non-quasi-linear environments, such as collective decision-making in committees, the first best outcome is generally not implementable under the commonly imposed restrictions on information acquisition technologies. For example, Li (2001) and Persico (2004), assuming that the agents have access to conditionally independent signals, show that the first best outcome is not attainable. In contrast to the previous results, we show that correlated information acquisition helps to implement the first best outcome.

There is a growing literature on the cost of flexible information in decision environments (see for example Sims (2003), Matejka and McKay (2015), Caplin and Dean (2015), and Pomatto et al. (2020)). Usually the main focus is on the cost of acquiring more or less precise information about a payoff-relevant state, and its relationship with a single decision-maker's optimal choice. The framework, however, has been applied in multi-player problems, e.g. in coordination games (Yang (2015); Morris and Yang (2021); Denti (2020)). In particular, Denti (2020) proposes a model of unrestricted information acquisition in games, in which, as in our paper, the players can endogenously learn about a payoff-relevant state and actions of other players.

the single agent (sender) uninformative; Kalai et al. (2010), Renou and Tomala (2012), Renault et al. (2014) in the context of games of communication network. Peters and Troncoso-Valverde (2013) apply this idea in mechanism-design games with multiple principals, and Liu (2015) applies it in his concept of individually uninformative correlating device. Our construction is most directly related to Zhu (2021).

[8]Bikhchandani (2010) shows that, indeed, an agent in the Crémer-McLean mechanism may have a strong incentive of acquiring information about others.

2 MODEL

2.1 Setup

There is a principal and $I \geq 4$ agents, and a finite set of payoff-relevant states Θ. Each agent i's payoff is denoted $u_i(d,\theta)$, when a social decision $d \in D$ is selected in state θ.[9] For example, in an auction, d is a vector of bidders' winning probabilities and their expected payments, and each u_i is quasi-linear in the payment part. Later, each agent's payoff *net* his information acquisition cost is considered as his objective.

At the beginning, neither the principal nor any of the agents know θ. The agents can acquire costly information about θ by generating private signals, possibly correlated with each other, whereas the principal cannot acquire any information about θ. Each agent has access to a sufficiently large set of possible signal realizations S_i. In principle, S_i (in particular, its size) may be a part of i's choice, but assuming exogenous S_i is without loss of generality as long as $|S_i| \geq |\Theta|$.

To model information acquisition, we introduce a space of states of nature $X = [0,1]$ with a typical element x, equipped with a Borel σ-algebra and a uniform probability measure .[10] We assume there is a commonly known measurable function $\Theta : X \to \Theta$ mapping the states of nature to the payoff-relevant states. This function induces a common prior on the payoff-relevant states as follows: $\mu_0(\theta) \equiv \int_0^1 1_{\{\Theta(x)=\theta\}} dx$ for each $\theta \in \Theta$. Agent i's information acquisition action is a measurable function $\sigma_i : X \to S_i$, such that, once x (and hence $\theta = \Theta(x)$) is realized, then i observes $s_i = \sigma_i(x)$. Let Σ_i denote the set of all such measurable functions, defining i's information acquisition action space. Note that any profile of information acquisition actions σ induces a joint distribution over payoff-relevant states and signal realizations, which we denote by $\alpha \in \Delta(\Theta \times S)$. When we want to make its dependence on σ more explicit, we write α_σ.

We assume that information acquisition is fully private in the sense that neither the principal nor any other agent observes which information acquisition action i takes and which signal realization is observed by agent i. Agent i's objective is the *net* payoff $u_i(d,\theta) - c_i(\sigma_i)$, where σ_i represents i's information acquisition action. We assume the information acquisition cost function of agent i has the following properties:

Assumption 6. Properties of information acquisition cost.

[9] We can endow the principal with his own payoff function $u_0(d,\theta)$, though it is not necessary.

[10] Taking a richer space of states of nature would not change our results. See also Gentzkow and Kamenica (2017) who use a similar approach in the context of multi-sender Bayesian persuasion.

1. $c_i(\sigma_i) \geq 0$ *for any* σ_i.

2. $c_i(\sigma_i) = 0$ *if* σ_i *and* Θ *are stochastically independent.*

The second property makes sure that agent i pays nothing as long as he learns nothing about the payoff-relevant state from his signal. This property is usually assumed in the context of single-player information acquisition.[11] For example, in the literature on rational inattention, the cost of information acquisition is often assumed to be proportional to the reduction in "relative entropy" (which measures the informativeness of a signal about the state). There, our second property is satisfied, because any signal that is stochastically independent from Θ leaves the relative entropy unchanged, and is therefore costless.

With multiple players, even if a signal is uninformative about the payoff-relevant state, it could be informative about other players' signals, which is the key to our result. Our study can be interpreted as investigating the consequence of this assumption (seemingly quite natural in single-agent environments) in multi-agent mechanism design environments.

2.2 Mechanism

The principal faces both hidden action and hidden information of each agent. The principal commits to a mechanism at the *ex ante* stage in order to control the agents' incentives. More specifically, following the literature, we let the principal (i) send a message privately to each agent before his information acquisition action, and (ii) collect a message privately from each agent after the agent has observed a signal realization. Formally, a mechanism comprises $(R, \rho; M, \delta)$ where $R = (R_i)_{i=1}^I$ and $M = (M_i)_{i=1}^I$; R_i denotes the set of messages that the principal can send to each agent i; M_i denotes the set of messages that each agent i can send to the principal; $\rho \in \Delta(R)$ is a distribution over the principal's messages, and $\delta : R \times M \to D$ denotes the decision rule.

The timing of the game is summarized as follows:

$t = 0$: $x \sim U(0, 1)$ is drawn but no one observes it.

$t = 1$: The principal designs a mechanism $(R, \rho; M, \delta)$.

$t = 2$: After observing the mechanism and receiving $r_i \in R_i$, each agent i privately chooses his information acquisition action $\sigma_i \in \Sigma_i$.

[11] See the literature on cost of information, such as Sims (2003), Matejka and McKay (2015), Caplin and Dean (2015), and Pomatto et al. (2020).

$t = 3$: Each agent i privately observes $s_i = \sigma_i(x)$, and privately sends $m_i \in M_i$ to the principal.

$t = 4$: The principal executes $d = \delta(r, m)$ where $m = (m_i)_{i=1}^I$.

Because no agent observes the other agents' actions or information (even noisily) at all, we consider Nash equilibrium as a solution concept. Then, applying the revelation principle of Forges (1986), we focus on *direct* mechanisms where (i) the principal directly recommends an information-acquisition action to each agent, and each agent directly reports a signal to the principal, and (ii) each agent finds it optimal to obey the recommended action and truthfully report his signal.[12]

Formally, a direct mechanism comprises $((\sigma_i)_{i=1}^I, (S_i)_{i=1}^I, \delta)$, where the principal recommends $\sigma_i \in \Sigma_i$ privately to each agent i,[13] and executes $\delta(s) \in D$ if the agents report $s = (s_i)_{i=1}^I \in S = \times_{i=1}^I S_i$. A direct mechanism is *incentive compatible* if it satisfies the following constraints: for any $\sigma_i' \in \Sigma_i$ and $\tau_i : S_i \to S_i$,

$$\sum_{\theta, s_i, s_{-i}} (u_i(\delta(s_i, s_{-i}), \theta)\alpha_{\sigma_i, \sigma_{-i}}(\theta, s_i, s_{-i})) - c_i(\sigma_i) \geq \sum_{\theta, s_i, s_{-i}} (u_i(\delta(\tau_i(s_i), s_{-i}), \theta)\alpha_{\sigma_i', \sigma_{-i}}(\theta, s_i, s_{-i})) - c_i(\sigma_i').$$

That is, each i must find it optimal to obey the recommended σ_i and report the realized s_i truthfully.

Although the constraints are concisely summarized by the inequalities above, they are actually rather complicated. First, changing σ_i affects the joint distribution α of (θ, s) and the agent's cost in a non-trivial way since agent i cannot affect agent $-i$'s information structure. Second, an agent may potentially want to make a double deviation, that is, change σ_i and at the same time change his reporting strategy.

Remark 3. *Here, we do not explicitly impose individual rationality constraints. It is not difficult to accommodate these constraints: let us require that any feasible direct mechanism must have an extra message* $m_i^{\emptyset}$ *(a "non-participation" message) so that i's message space is now* $S_i \cup \{m_i^{\emptyset}\}$,

[12]The proof proceeds as follows. First, imagine an auxiliary game where there is no principal, but instead, there is a fictitious player ("player 0") who is indifferent across all decisions in any state. At first, each agent i plays σ_i privately, and then observes the realized signal s_i privately. Then, (without any communication), player 0 chooses $d \in D$. Interpreting this as a baseline extensive-form game, it is easy to see that our current game (with the principal) is the mediated communication game of this auxiliary game in the sense of Forges (1986) (see also Myerson (1986*a*)). Thus, her revelation principle applies.

[13]We focus on a deterministic recommendation of σ, rather than any stochastic recommendation. Accordingly, δ is denoted simply by $\delta(s)$ instead of $\delta(r, s)$. Since first best implementation is achieved with pure recommendations, our focus on them is without loss of generality.

and $\delta(m_i^{\emptyset}, m_{-i})$ is some specific allocation (a "non-participation allocation") for agent i, for any given m_{-i}. When the non-participation message is included into the set of messages for each agent, the individual rationality constraints, both at the ex ante and interim stages, are captured by the above incentive compatibility constraints.[14]

3 MAIN RESULT

Fix any function $d^* : \Theta \to D$, which describes all the economically relevant outcomes in this environment except for the information acquisition costs. If the principal could observe θ, then any d^* is attainable without any information acquisition cost on the agents' side. In this sense, one may interpret this d^* together with zero cost for the agents as the *first-best* outcome.[15]

In this section, for any given d^*, we explicitly construct a mechanism that implements d^* at zero cost for the agents. That is, the first best outcome can be attained *even though the principal cannot directly observe θ.*

Theorem 7. *Fix any $d^* : \Theta \to D$. Under Assumption 6, there exists a mechanism (σ, S, δ) such that (i) $\sum_s \alpha(\theta, s) 1_{\{\delta(s)=d^*(\theta)\}} = \mu_0(\theta)$ for all θ, and (ii) $c_i(\sigma_i) = 0$ for all i.*

Proof. The theorem is proved by construction.

Since Θ is a finite set, we assume without loss of generality that $\Theta = \{1, \ldots, T\}$. Let $K > \max\{I, T\}$ be a prime number. Because is a uniform measure on $X = [0,1]$, we can find a partition of X, denoted by $\{X_{\theta\psi}\}_{(\theta,\psi)\in\{1,\ldots,T\}\times\{1,\ldots,K\}}$, satisfying $\int_0^1 1_{\{x\in X_{\theta\psi}\}} dx = \frac{1}{K}\mu_0(\theta)$ for any θ and ψ. Define a measurable function $\Psi : [0,1] \to \{1,\ldots,K\}$ such that, if $x \in \cup_{\theta\in\Theta} X_{\theta\psi}$, then $\Psi(x) = \psi$. Immediately, Ψ is uniformly distributed on $\{1,\ldots,K\}$ conditional on any realization θ of Θ, hence Ψ is independent of Θ.

Now consider the following information acquisition action profile: for each $i \in \{1,\ldots I\}$, $S_i = \{1,\ldots,K\}$, and $\sigma_i(x) = \Theta(x) + i \cdot \Psi(x) \mod K$ for any $x \in [0,1]$. Note that the residual is calculated as in standard modular arithmetic except when $\Theta(x) + i \cdot \Psi(x)$ is divisible by K, in which case we set $\sigma_i(x) = K$ instead of 0. The following lemma gives the properties of (S, σ) that we need to prove the theorem.

[14] Ex interim individual rationality is guaranteed because agent i can always deviate to $\tau_i(\cdot) \equiv m_i^{\emptyset}$. Ex ante individual rationality is guaranteed because agent i can always deviate to a costless σ_i and then to $\tau_i(\cdot) \equiv m_i^{\emptyset}$.

[15] For example, one may assume that $d^*(\theta)$ is the best decision of the principal given his own preferences in state θ.

Lemma 15. *The above* (S,σ) *satisfies:*

(i) For any $i \in \{1,\dots,I\}$*,* σ_i *is independent of* Θ*.*

(ii) Conditional on any realization of (s_i, s_j) *such that* $i \neq j$*, the joint distribution of* Θ *and* $(\sigma_k)_{k\neq i,j}$ *is degenerate.*

Proof of the lemma. By definition of (S,σ), for each i, we have $s_i = \theta + i \cdot \psi \mod K$, where the random variables Ψ and Θ are independent. Thus the signal profile $s = (s_i)_{i=1}^I$ is defined in the same way as in Zhu (2021).[16] Thus, this lemma is directly implied by Lemma 2 in Zhu (2021). □

The first property says that Θ and σ_i are independent, implying $c_i(\sigma_i) = 0$. The second property says that, given s_i, s_j with $i \neq j$, we can identify the true payoff-relevant state θ and any signal realization s_k without error, that is, there exist $\hat{\theta}(s_i, s_j)$ and $\hat{s}_k(s_i, s_j)$ such that:

$$\Pr\big(\Theta = \hat{\theta}(s_i, s_j)|s_i, s_j\big) = \Pr\big(\sigma_k = \hat{s}_k(s_i, s_j)|s_i, s_j\big) = 1.$$

Let the principal recommend the above σ, and offer the decision rule δ as follows: $\delta(s) = d^*(\theta)$ if (i) for any i, j with $i \neq j$, we have

$$\theta = \hat{\theta}(s_i, s_j);$$

or if (ii) there is i such that, for any j, k where i, j, k are all different, we have

$$\theta = \hat{\theta}(s_j, s_k).$$

In any other case, $\delta(s)$ is arbitrary.

Clearly, if the agents obey the recommendation and report their signals truthfully, then the first best outcome is attained. Therefore, we complete the proof by showing that the proposed mechanism satisfies incentive compatibility. Take any agent i, and suppose that he deviates to any σ_i' and reports $\tau_i(s_i)$ when s_i is realized. First, his cost of information acquisition increases weakly. Second, his reporting decision does not affect the social decision at all, because the principal executes $\delta(s) = d^*(\hat{\theta}(s_j, s_k))$ for an arbitrary pair (j, k) which does not include i. Therefore, the mechanism is incentive compatible. □

[16] In fact, our signal profile s coincides with what Zhu (2021) calls “the IUAR disclosure policy, where IUAR is short for *individually uninformative but aggregately revealing.*

3.1 Impossibility results with two and three agents

One could ask whether a result similar to Theorem 7 obtains with two or three agents. The general answer to this question is *no*. With three agents, although it is possible to determine whether *some* agent has unilaterally deviated or not, it is not possible to identify who the deviator is (and hence not possible to identify the true θ). To see that, consider the following counterexample.

Counterexample 1. *Suppose that there are two payoff-relevant states, i.e.* $\Theta = \{1,2\}$ *and consider the mechanism with* $K = 5$. *Computing* $s_i = \theta + i\psi \mod 5$, *we obtain:*

$\theta = 1$	Agent 1	Agent 2	Agent 3
$\psi = 1$	2	3	4
$\psi = 2$	3	5	2
$\psi = 3$	4	2	5
$\psi = 4$	5	4	3
$\psi = 5$	1	1	1

$\theta = 2$	Agent 1	Agent 2	Agent 3
$\psi = 1$	3	4	5
$\psi = 2$	4	1	3
$\psi = 3$	5	3	1
$\psi = 4$	1	5	4
$\psi = 5$	2	2	2

Suppose the principal observes an out-of-equilibrium signal realization profile $(2,5,4)$. *There are two unilateral deviations that lead to this profile. First, the true profile might be* $(2,3,4)$ *in state* $\theta = 1$ *with agent 2 deviating. Second, the true profile might be* $(1,5,4)$ *in state* $\theta = 2$ *with agent 1 deviating. Hence, the principal cannot identify the deviator, nor can the principal infer the true state.*

Note, however, that if there exists a social decision $d \in D$ that can serve as a severe punishment for all agents for any given θ, then out-of-equilibrium reports can be severely punished by the principal,[17] and a similar first-best implementation result obtains.

With two agents, each agent has much more freedom. The authors work on a separate project with two agents. There, even under Assumption 6 and even when monetary transfers are available to the principal, an extremely positive result similar to Theorem 7 does not generally hold. The optimal mechanism might involve some costly information acquisition.

[17] Consider e.g. environments with monetary transfers, in which the principal can use large fines to punish agents for inconsistent reports.

4.1 Full-surplus extraction in common value auctions

Consider the following common value auction environment. The seller (principal) has a single indivisible good, and there are $I \geq 4$ bidders. The value of the good is common to all the bidders, denoted by $\theta \in \Theta$, where Θ is finite. In fact, the analysis of this section can be straightforwardly extended to the case of "non-pure" common values where each i's valuation is $v_i(\theta)$. Let $\mu_0(\theta)$ denote the probability that θ is the bidders' common value.

Each bidder i's payoff is $\theta q_i - t_i - c_i(\sigma_i)$ if he wins the good with probability q_i, pays t_i to the seller, and spends $c_i(\sigma_i)$ as his information acquisition cost. In case he does not participate in the mechanism, his outside-option payoff is 0. The seller's payoff is revenue, $\sum_{i=1}^{I} t_i$.

The first-best expected surplus of this society is the expected common value:

$$\sum_{\theta \in \Theta} \mu_0(\theta)\theta = \mathbb{E}[\theta].$$

There are several cases where the seller can easily earn $\mathbb{E}[\theta]$. First, if the seller *knows* θ, then he can simply post price θ. Even if the seller does not know θ, if the bidders know θ as their common knowledge (i.e., as *free* information), then again the seller can earn $\mathbb{E}[\theta]$. Conversely, if all the bidders are completely *uninformed* (so that each only knows the common prior μ_0), then again, the seller can post price $\mathbb{E}[\theta]$.

Notice that, with costly information acquisition as considered in our paper, neither of the above ideas would work. First, although it might be possible to make every bidder fully learn θ in some equilibrium, it does not yield $\mathbb{E}[\theta]$ as long as full information is strictly costly. Second, if the seller posts price $\mathbb{E}[\theta]$, then each bidder has a strong incentive of knowing whether the true θ is below $\mathbb{E}[\theta]$ or not: If i finds that $\mathbb{E}[\theta|s_i] < \mathbb{E}[\theta]$ given some signal s_i, he would not buy the good. As long as such information is not too costly, the bidder would be better off by acquiring it.

Therefore, with a general information acquisition cost function, the equilibrium information should be somewhere between full and no information, and it is *a priori* unclear how the seller should find the optimal balance of information and rent extraction. Nevertheless, as long as the cost functions satisfy Assumption 6, Theorem 7 implies that the full-surplus extraction is possible.

Corollary 3. *Under Assumption 6, there is a mechanism which yields* $\mathbb{E}[\theta]$ *as the seller's expected revenue (and each bidder earns 0).*

It is worth emphasizing that the logic here is very different from that of Crémer and McLean (1988*b*). In their paper, the seller exploits an *exogenously given* correlated signal structure, in order to construct a side-bet scheme that extracts the entire surplus. In our case, each bidder can choose any information structure. Indeed, if he prefers, a bidder can choose an information structure such that his information is *independent* from all the other bidders' signals (conditional on the state of the world). The Crémer-McLean lottery scheme, therefore, does not work here. Also, in their auction, each bidder's payoff can be strictly negative *ex post*, while in our case, it is zero *ex post*. Indeed, if the seller offered a negative *ex post* payoff in our auction, bidders would have a strong incentive to get a signal which includes a realization indicative of that event and then abstain from the auction following that realization.

4.2 First-best implementation in collective decision-making

Consider a committee with a designer (principal) and $I \geq 4$ members (agents) deciding whether to hire or not to hire a job market candidate. Formally, $d \in D = \{h, nh\}$. The quality of the candidate is $\theta \in \Theta$, which is unobserved *ex ante*. The designer and all members of the committee hold a common prior belief $\mu_0 \in \Delta(\Theta)$ about the candidate's quality.

The utility that each member obtains from hiring / not hiring the candidate is defined as follows:

$$u_i(d, \theta) = \begin{cases} u_i(\theta), & \text{if } d = h \\ 0, & \text{if } d = nh \end{cases}$$

Without loss of generality, we assume that $u_i(\theta) = k_i \theta$.

Only the committee members can acquire information about the candidate at cost $c_i(\sigma_i)$. The designer aims to maximize the expected sum of all members' gross utilities.[18] That is, ideally, he wants to hire the candidate if and only if $\sum_i k_i \theta \geq 0$. The first best expected surplus of all committee members is given by:

$$\sum_{\theta \mid \sum_i k_i \theta \geq 0} \mu_0(\theta) \sum_i k_i \theta \equiv W^{FB}$$

[18] The result extends to the case where the designer maximizes expected sum of members' net utilities (taking into account the information acquisition costs).

It is useful to note that the existing literature (see e.g. Li (2001) and Gerardi and Yariv (2008)) typically assumes that the committee members have access to information structures whose realized signals are independently distributed across them, conditional on the state of the world. Under these restrictions, the first best outcome cannot be implemented. There are two main forces that prevent the committee from implementing the first best with these restricted information structures: free-riding problem and conflict of interest. First, when committee members have a conflict of interest, they may prefer not to report their own acquired information truthfully. Second, even if all members share a common preference, information could be underprovided relative to the social optimum, because it is essentially a public good used to make a collective decision. For example, Li (2001) suggests that distorting the decision rule away from the *ex post* optimal rule (which is optimal under exogenous information) could help to alleviate the free-riding issue.

In contrast to the previous literature, our results show that with more flexible (even though still costly) information acquisition the designer can implement the first best outcome. On the one hand, having access to a wider range of information acquisition technologies enlarges the set of feasible deviations for the agents. On the other hand, the principal now has more flexibility in designing information structures recommended to the agents. Given these two opposing effects, it is not immediately clear *a priori* whether the first best outcome becomes more or less difficult to attain. It turns out that the second effect dominates: with a larger set of feasible mechanisms, the principal is able to incentivize the agents to acquire and report their information truthfully, no matter what social choice rule the principal is trying to implement. Indeed, Theorem 7 implies that the first best is implementable as long as Assumption 6 about the cost functions holds.

Corollary 4. *Under Assumption 6, there is a mechanism which yields W^{FB} as the total expected surplus (the decision is made under full information with no cost).*

Our construction helps to resolve both of the issues that prevent first best implementation with conditionally independent signals. Recall that it costs nothing for an agent to acquire an *"individually uninformative"* signal which is assigned to him under the optimal mechanism. Therefore, the distorted provision of information is resolved. Moreover, even if committee members have a conflict of interest, under our mechanism, they cannot do better than being truthful since any unilateral deviation can be detected by the designer.

5 CONCLUDING REMARKS

It is quite natural that agents may desire to refine their information in response to a mechanism. This paper proposes one possible framework, based on a class of information acquisition cost functions, such that the cost of information depends on the informativeness of each agent's signal about the state of the world, but not on its informativeness about other agents' signals. We show that such a specification leads to an extremely positive result.

One natural criticism may be that our mechanism induces full information if the signals are aggregated, even though any single signal is completely uninformative: would it be reasonable to assume that such σ_i is costless? Because the answer is necessarily yes under Assumption 6, the question is essentially whether Assumption 6 itself is reasonable. Assumption 6 is satisfied, in particular, in any information acquisition environment, in which the cost of information acquisition does not depend on the correlation structure among signals but only depends on their individual informational content. We argued in Section 1 that there are information acquisition environments, for which this assumption is indeed a reasonable one. In general, however, the correlation structure might affect the cost of information acquisition in various ways. On the one hand, there seem to be cases where more positive correlation is more expensive. For example, fix agent 1's private information, and consider agent 2. If acquiring a positively correlated information necessarily means that agent 2 must *steal* (perhaps a part of) agent 1's information, more positive correlation will be more costly. Strulovici (2021), for instance, considers an environment where hard evidence is scarce in the sense that, if one agent "picks up" a piece of evidence, then it becomes difficult for the others to get the same or similar evidence. On the other hand, there are opposite situations, where *less correlated* signals are more costly. For example, suppose that there are 2 agents and 3 newspapers, and σ_i corresponds to the decision of which newspapers to buy. Suppose further that newspaper 3 is free, hence both agents read it and acquire perfectly correlated signals for free. To acquire less correlated signals, at least one of the agents has to buy additional information (e.g. agent 1 buys newspaper 1 and/or agent 2 buys newspaper 2), hence less correlation can be more costly in this example.

This discussion suggests that we must think more about modelling the microstructure of information acquisition, in order to determine which information structures are more costly. Mechanism design with such more specific information acquisition cost structures would certainly be an interesting future direction, and we hope this article could serve as a first step in that direction.

Bibliography

Bang, S. H. and Kim, J. (2013), 'Price discrimination via information provision', *Information Economics and Policy* **25**(4), 215–224.

Battaglini, M. (2005), 'Long-term contracting with markovian consumers', *American Economic Review* **95**(3), 637–658.

Battaglini, M. and Lamba, R. (2019), 'Optimal dynamic contracting: The first-order approach and beyond', *Theoretical Economics* **14**(4), 1435–1482.

Bergemann, D., Brooks, B. and Morris, S. (2017), Informationally robust optimal auction design. *Working paper*.

Bergemann, D., Heumann, T. and Morris, S. (2022), 'Screening with persuasion', *arXiv preprint arXiv:2212.03360* .

Bergemann, D. and Morris, S. (2005), 'Robust mechanism design', *Econometrica* pp. 1771–1813.

Bergemann, D. and Pesendorfer, M. (2007), 'Information structures in optimal auctions', *Journal of economic theory* **137**(1), 580–609.

Bergemann, D. and Schlag, K. (2011), 'Robust monopoly pricing', *Journal of Economic Theory* **146**(6), 2527–2543.

Bergemann, D. and Välimäki, J. (2002), 'Information acquisition and efficient mechanism design', *Econometrica* **70**(3), 1007–1033.

Bergemann, D. and Wambach, A. (2015), 'Sequential information disclosure in auctions', *Journal of Economic Theory* **159**, 1074–1095.

Bikhchandani, S. (2010), 'Information acquisition and full surplus extraction', *Journal of Economic Theory* **145**, 2282–2308.

Bikhchandani, S. and Obara, I. (2017), 'Mechanism design with acquisition of correlated information', *Economic Theory* **63**, 783–812.

Brooks, B. and Du, S. (2021), 'Optimal auction design with common values: An informationally-robust approach', *Econometrica* **89**(3), 1313–1360.

Caplin, A. and Dean, M. (2015), 'Revealed preference, rational inattention, and costly information acquisition', *American Economic Review* **105**, 2183–2203.

Carroll, G. (2017), 'Robustness and separation in multidimensional screening', *Econometrica* **85**(2), 453–488.

Carroll, G. and Meng, D. (2016), 'Locally robust contracts for moral hazard', *Journal of Mathematical Economics* **62**, 139–166.

Chen, Y.-C. and Li, J. (2018), 'Revisiting the foundations of dominant-strategy mechanisms', *Journal of Economic Theory* **178**, 294–317.

Chung, K. and Ely, J. (2007), 'Foundations of dominant-strategy mechanisms', *The Review of Economic Studies* **74**(2), 447–476.

Condorelli, D. and Szentes, B. (2020), 'Information design in the holdup problem', *Journal of Political Economy* **128**(2), 681–709.

Courty, P. and Li, H. (2000), 'Sequential screening', *The Review of Economic Studies* **67**(4), 697–717.

Crémer, J. and McLean, R. (1988*a*), 'Full extraction of the surplus in bayesian and dominant strategy auctions', *Econometrica* pp. 1247–1257.

Crémer, J. and McLean, R. (1988*b*), 'Full extraction of the surplus in bayesian and dominant strategy auctions', *Econometrica* pp. 1247–1257.

Denti, T. (2020), 'Unrestricted information acquisition', *Working paper* .

Eliaz, K. and Spiegler, R. (2008), 'Consumer optimism and price discrimination', *Theoretical Economics* **3**(4), 459–497.

Eső, P. and Szentes, B. (2007), 'Optimal information disclosure in auctions and the handicap auction', *The Review of Economic Studies* **74**(3), 705–731.

Forges, F. (1986), 'An approach to communication equilibria', *Econometrica* **54**(6), 705–731.

Gentzkow, M. and Kamenica, E. (2017), 'Bayesian persuasion with multiple senders and rich signal spaces', *Games and Economic Behavior* **104**, 411–429.

Gerardi, D. and Yariv, L. (2008), 'Information acquisition in committees', *Games and Economic Behavior* **62**(2), 436–459.

Gershkov, A. and Szentes, B. (2009), 'Optimal voting schemes with costly information acquisition', *Journal of Economic Theory* **144**(1), 36–68.

Gleyze, S. and Pernoud, A. (2020), Infromationally simple incentives. *Working paper.*

Grubb, M. D. (2009), 'Selling to overconfident consumers', *American Economic Review* **99**(5), 1770–1807.

Guo, Y., Li, H. and Shi, X. (2022), Optimal discriminatory disclosure, Technical report, Working paper.

Guo, Y. and Shmaya, E. (2019), 'The interval structure of optimal disclosure', *Econometrica* **87**(2), 653–675.

Jehiel, P., Meyer-ter Vehn, M. and Moldovanu, B. (2012), 'Locally robust implementation and its limits', *Journal of Economic Theory* **147**(6), 2439–2452.

Jehiel, P., ter Vehn, M. M., Moldovanu, B. and Zame, W. R. (2006), 'The limits of ex post implementation', *Econometrica* **74**(3), 585–610.

Johnson, J. P. and Myatt, D. P. (2006), 'On the simple economics of advertising, marketing, and product design', *American Economic Review* **96**(3), 756–784.

Kalai, A. T., Kalai, E., Lehrer, E. and Samet, D. (2010), 'A commitment folk theorem', *Games and Economic Behavior* **69**(1), 127–137.

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

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A. and Li, M. (2017), 'Persuasion of a privately informed receiver', *Econometrica* **85**(6), 1949–1964.

Krähmer, D. (2020), 'Information disclosure and full surplus extraction in mechanism design', *Journal of Economic Theory* **187**, 105020.

Krähmer, D. and Strausz, R. (2015*a*), 'Ex post information rents in sequential screening', *Games and Economic Behavior* **90**, 257–273.

Krähmer, D. and Strausz, R. (2015*b*), 'Optimal sales contracts with withdrawal rights', *The Review of Economic Studies* **82**(2), 762–790.

Krähmer, D. (2020), 'Information disclosure and full surplus extraction in mechanism design', *Journal of Economic Theory* **187**.

Krähmer, D. (2021), 'Information design and strategic communication', *American Economic Review: Insights* **3**(1), 51–66.

Lewis, T. R. and Sappington, D. E. (1994), 'Supplying information to facilitate price discrimination', *International Economic Review* pp. 309–327.

Li, H. (2001), 'A theory of conservatism', *Journal of political Economy* **109**(3), 617–636.

Li, H. and Shi, X. (2017), 'Discriminatory information disclosure', *American Economic Review* **107**(11), 3363–3385.

Liu, Q. (2015), 'Correlation and common priors in games with incomplete information', *Journal of Economic Theory* **148**(3), 49–75.

Lopomo, G., Rigotti, L. and Shannon, C. (2021), 'Uncertainty in mechanism design', *working paper* .

Lopomo, G., Rigotti, L. and Shannon, C. (2022), 'Uncertainty and robustness of surplus extraction', *Journal of Economic Theory* **199**, 105088.

Madarász, K. and Prat, A. (2017), 'Sellers with misspecified models', *The Review of Economic Studies* **84**(2), 790–815.

Matejka, P. and McKay, A. (2015), 'Rational inattention to discrete choices: A new foundation for the multinomial logit', *American Economic Review* **105**, 272–298.

Mensch, J. (2020), 'Screening inattentive agents', *Working paper* .

Morris, S. and Yang, M. (2021), Coordination and continuous stochastic choice. Working paper.

Myerson, R. (1986*a*), ‘Multistage games with communication’, *Econometrica* **54**(2), 323–358.

Myerson, R. B. (1981), ‘Optimal auction design’, *Mathematics of operations research* **6**(1), 58–73.

Myerson, R. B. (1986*b*), ‘Multistage games with communication’, *Econometrica: Journal of the Econometric Society* pp. 323–358.

Ollár, M. and Penta, A. (2017), ‘Full implementation and belief restrictions’, *American Economic Review* **107**(8), 2243–2277.

Pavan, A., Segal, I. and Toikka, J. (2014), ‘Dynamic mechanism design: A myersonian approach’, *Econometrica* **82**(2), 601–653.

Persico, N. (2004), ‘Committee design with endogenous information’, *The Review of Economic Studies* **71**(1), 165–191.

Peters, M. and Troncoso-Valverde, C. (2013), ‘A folk theorem for competing mechanisms’, *Journal of Economic Theory* **148**(3), 953–973.

Pomatto, L., Strack, P. and Tamuz, O. (2020), ‘The cost of information’, *Working paper* .

Ravid, D., Anne-Katrin, R. and Szentes, B. (2020), Learning before trading: On the inefficiency of ignoring free information. *Working paper*.

Renault, J., Renou, L., and Tomala, T. (2014), ‘Secure message transmission on directed networks’, *Games and Economic Behavior* **85**, 1–18.

Renou, L. and Tomala, T. (2012), ‘Mechanism design and communication networks’, *Theoretical Economics* **7**(3), 489–533.

Roesler, A.-K. and Szentes, B. (2017), ‘Buyer-optimal learning and monopoly pricing’, *American Economic Review* **107**(7), 2072–2080.

Shannon, C. (1949), ‘Communication theory of secrecy systems’, *Bell System Technical Journal* **28**(4), 656–715.

Shi, X. (2012), ‘Optimal auctions with information acquisition’, *Games and Economic Behavior* **74**, 666–686.

Sims, C. (2003), ‘Implications of rational inattention’, *Journal of Monetary Economics* **50**, 665–690.

Skreta, V. (2006), 'Mechanism design for arbitrary type spaces', *Economics Letters* **91**(2), 293–299.

Smolin, A. (2023), 'Disclosure and pricing of attributes', *The Rand Journal of Economics* .

Strulovici, B. (2021), 'Can society function without ethical agents? an informational perspective', *Working paper* .

Wei, D. and Green, B. (2023), '(reverse) price discrimination with information design', *American Economic Journal: Microeconomics* .

Wilson, R. (1969), 'Competitive bidding with disparate information', *Management Science* **15**(7), 446–448.

Yamashita, T. (2018), 'Revenue guarantee in auction with a (correlated) common prior and additional information', *Working paper* .

Yamashita, T. and Zhu, S. (2022), 'On the foundations of ex post incentive compatible mechanisms', *American Economic Journal: Microeconomics* **14**(4), 494–514.

Yang, M. (2015), 'Coordination with flexible information acquisition', *Journal of Economic Theory* **158**, 721–738.

Zhao, X. (2016), 'Heterogeneity and unanimity: Optimal committees with information acquisition', *Working paper* .

Zhu, S. (2021), 'Private disclosure with multiple agents', *Working paper* .

Zhu, S. (2023), 'Private disclosure with multiple agents', *Journal of Economic Theory* **212**, 105705.