Selling Information in Competitive Environments

by

Amir Nouripour

B.Sc. in Electrical Engineering, Sharif University, Tehran (2019)
B.Sc. in Mathematics and Applications, Sharif University, Tehran (2019)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2021

Author . . . .
Department of Electrical Engineering and Computer Science
August 27, 2021

Certified by . . . .
Munther A. Dahleh
William A. Coolidge Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by . . . .
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Sciences
Chair, Department Committee on Graduate Students

Selling Information in Competitive Environments

by

Amir Nouripour

Submitted to the Department of Electrical Engineering and Computer Science
on August 27, 2021, in partial fulfillment of the
requirements for the degree of
Master of Science in Electrical Engineering and Computer Science

Abstract

We consider a setting where data buyers compete in a game of incomplete information, about which a data seller owns some payoff relevant information. We formulate the problem facing the seller as a joint information and mechanism design problem: deciding which information to sell, while at the same time eliciting the private value types of the buyers and collecting payments. We derive the welfare- and revenue-optimal mechanisms for a class of binary games. Our results reveal some important features of selling information in competitive environments: (i) the negative externalities arising from competition among buyers increase the profitability of selling exclusive information to one of the buyers; (ii) in order for the buyers to follow the seller's action recommendations, the extent of exclusive sales must be limited; (iii) these same equilibrium constraints also limit the distortions in the allocation of information distortion that can be introduced by a monopolist data seller; (iv) the fiercer the competition across buyers the stronger the previous two limitations, and the weaker the impact of market power on the equilibrium allocation of information.

Thesis Supervisor: Munther A. Dahleh
Title: William A. Coolidge Professor of Electrical Engineering and Computer Science

Acknowledgments

First, I want to thank my advisor Prof. Munther A. Dahleh for providing me with the opportunity to find and think on a problem I like and helping me through it by his illuminating guidance.

To Thibaut Horel: Thanks for your patience, knowledge, and all I learnt from you.

To Prof. Alessandro Bonatti: Thanks for the new insights and intuitions that you brought to this project.

To my friends: Thanks for keeping the life fun.

Finally, to my family: Thanks for all your support and devotion. Thanks for keeping me motivated and your guidance through my whole life. Thanks for everything.

Contents

Chapter 1

Introduction

In online and offline markets alike, most sales of information take place in competitive environments: credit bureaus sell consumer scores to competing lenders; and large digital platforms sell targeted audiences to competing retailers, e.g., through keyword advertising slots or sponsored listings.

Settings such as these share two key properties: (i) the marginal profitability of acquiring information for any buyer is unknown to the seller, and (ii) selling information to a buyer can have a (negative) impact on other buyers' profits.

In this work, we study the design of mechanisms for selling information in the presence of allocative externalities. Our work is directly inspired by [1], in which these externalities were modeled as *intrinsic* to the buyers: the negative marginal effect on a given buyer of another buyer acquiring information was assumed to be part of their privately observed type. In contrast, we model these externalities as *explicitly stemming from the competition* in which the buyers engage.

Our goal is to propose a tractable formulation of this situation and to answer a range of questions about the allocative and welfare properties of such information selling mechanisms, including:

- What is the structure of the optimal (revenue maximizing) mechanism? How does it depend on the form and on the intensity of the downstream competition?
- How does the information being sold affect the equilibrium strategies and payoffs in

the downstream game?

- How does the provision of information to competing players translate into revenue for the seller, and into social welfare? In other words, can we relate the value of data for the seller to in terms of its added value in the downstream competition?

We formulate the seller's choice of the optimal mechanism for the sale of information as an *information design* problem *with elicitation* [5]. Specifically, we consider two data buyers and a data seller. The data buyers compete in a simultaneous-moves, finite game of incomplete information (the "downstream game"). Loosely, the more information a buyer has, the better her decisions are in the downstream game, which results in *lower* payoffs for the other player.

The seller is informed about a payoff relevant state variable: in our simplest example, the state identifies which action is dominant for each player in the downstream game. The seller maximizes revenues by selling this information. However, each buyer also has a private value type (e.g., a technology parameter) that scales her payoffs in the downstream game linearly. Thus, each buyer has a private willingness to pay for any additional information the seller offers. In turn, the seller must first elicit the buyers' private payoff types, and then send them informative signals. This is a joint mechanism and information design problem, wherein the seller's choice is subject to the buyers' obedience and truthful reporting constraints.

Our leading example considers the interaction between a digital platform (information seller) and two merchants (information buyers) who wish to leverage the platform's information to offer (e.g., to advertise) one of their products to a single consumer. We first describe why this economically relevant setting is indeed one of information trading. We then describe our mechanism design problem in the context of this example.

Motivating Example Large digital platforms (e.g., Amazon, Facebook, and Google in the US, Alibaba and JD in China) collect an ever-increasing amount of information on their users' online behavior (e.g., browsing, shopping, social media interactions), which allows them to precisely estimate individual consumers' tastes for various products.

These platforms typically monetize their proprietary information by selling targeted advertising campaigns (Facebook, Google) or sponsored marketplace listings (Amazon) to advertisers and retailers. Such practices amount to *indirect sales of information* (see [2] for a thorough discussion): while the platform does not trade its consumers' data for payment (*direct sales*), it is nonetheless able to create value for merchants by allowing them to condition their strategies (in particular, which product to offer) on the consumers' information (e.g., their browsing or shopping history, and third-party cookies). In the context of our static model, direct transfers of information and indirect sales of information are, in fact, equivalent.[1]

For an individual merchant, the expected volume of sales on a platform is driven by two factors: (i) the degree of targeting, or precision of its campaign, as measured by the ability to show the right product to each consumer; (ii) the exclusivity of its campaign, i.e., the mismatch between its competitors' offers and the consumer's tastes. Merchants are willing to pay more for better-targeted campaigns, and even more for exclusive access to targeted campaigns.

However, the merchants' willingness to pay for an advertising campaign also depends on the profitability of making each sale, i.e., on their cost structure. As the latter is privately known to the merchant, the platform must elicit it through its choice of mechanism. Abstracting from the details and dynamics of online advertising auctions, the platform's problem reduces to designing a menu of (information structure, payment) pairs. Each information structure corresponds to an advertising campaign with a certain degree of precision and exclusivity.

An information structure maps the true state (i.e., the consumer's tastes) and the competing merchants' choices into a distribution over informative signals. Thus, each buyer chooses an information structure from the platform's menu and pays the platform upfront, but the joint distribution of states and signals depends on all buyers' choices.

These signals can be viewed as action recommendations, e.g., what product to show the consumer. As such, the platform's problem is subject to two types of constraints simultaneously: first, the merchant must find it worthwhile to follow the platform's recommendation (obedience constraints); second, the merchant must choose an option from the platform's

[1] In a more realistic, dynamic setting, there are several reasons why digital platforms prefer indirect sales, e.g., maintaining a reputation for privacy with their consumers, and using click-through rates to certify the quality of their information.

menu that correctly reveals its own preferences (truthful reporting constraints).

Summary of the Results We begin by setting up the model as a joint information and mechanism design problem, we restrict attention to a linear setting with independent types. For a binary example in this class, we characterize the welfare- and the revenue-maximizing mechanisms.

Our results highlight two defining features of selling information to competing players, and show how information and competition interact in shaping the optimal mechanism. Both features distinguish the sale of information to competing firms from the sale of a physical good, and from the case of a single buyer.

First, any buyer can always ignore (or indeed, reverse) the seller's recommendation. The resulting *obedience* constraint limits the social planner's ability to reveal information to one player *exclusively.* Likewise, obedience disciplines the monopolist seller's ability to distort the allocation of information to maximize revenue at the expense of social welfare. Intuitively, the seller wants to distort the allocation of any buyer type with a negative Myersonian virtual value, so as to minimize her payoff and reduce the information rents of higher types.

In our setting, this distortion corresponds to recommending the wrong action in every state. However, the buyer would not follow such a recommendation in any mechanism that does so too often. Therefore, the seller can do no better than to reveal "zero net information" to a low-value buyer, i.e., to probabilistically send the right and the wrong recommendation in a way that leaves the buyer indifferent over any course of action.

There are, of course, many such information structures (including entirely uninformative ones). However, the seller is not indifferent among them. Indeed, she prefers to reveal the correct state to both buyers when their types are sufficiently low. This relaxes a low type's obedience constraint and allows the seller to issue the correct recommendation to one player *exclusively* (and the wrong recommendation to the other player) when their two types are sufficiently different.

Second, providing information to one firm naturally imposes a negative externality on its competitors. In our setting, these negative externalities expand the profitability of selling information. In particular, each buyer is willing to pay a positive price as long as either (a) she is strictly better off following the seller's recommendation, or (b) her opponent does not receive the correct recommendation with probability one. As a result, the seller can charge a strictly positive price to some types with negative Myersonian virtual values, including some types whose obedience constraint binds, i.e., those who do not receive any valuable information themselves.

Related Literature [3] study the sale of data to a single buyer with private information about her beliefs over a payoff relevant state. Their problem is similar to us insofar as the optimal mechanism can be represented through a menu of options (i.e., distributions over action recommendations) and associated prices. However, because there is a single buyer, the externalities that are central to our analysis do not arise.

[6] study a very similar setting, but only consider mechanisms with a single option: selling the true state distorted by Gaussian noise. Their problem then consists of finding the optimal covariance matrix of the noise and the associated prices. In particular, the covariance matrix is not designed as a function of the buyers' private types.

As mentioned, our work is closely related to the information design literature, e.g., [4], [11], and the references therein. Most papers in this literature view the the problem as a pure information design question as opposed to a mechanism design question in a quasilinear setting. Consequently, these papers do not derive the implications of optimal information design for the designer's revenue.

[12] study competing information sellers who offer selling exogenous, imperfect information structures about a binary state to buyers with known types who compete in a downstream game with binary actions.

Finally, [8] and [7] study specific instances of selling information in competitive markets. The former consider the sale of cost information to a large number of perfectly competitive firms, each one facing a privately informed manager. The latter sell information about a consumer's location along a Hotelling line to two firms with fixed locations who can use this data to set personalized prices.

Chapter 2

Model

We consider n buyers who compete in a game of incomplete information (the "downstream game"), and a monopolist data seller who observes a payoff relevant state variable. Loosely speaking, the seller elicits the buyers' private payoff types, and sells informative signals to the data buyers. A complete specification of our model thus comprises (i) the payoff functions of the data buyers in the downstream game, (ii) the information structure describing which payoff relevant parameters is observed by each agent, and (iii) the information exchange and transfers between the data buyers and the data seller occurring prior to the downstream game. We now describe each component in turn after introducing some notation.

Notation For a tuple of sets $(\mathcal{S}_i)_{i\in[n]}$, we write $\mathcal{S} = \prod_{i=1}^{n} \mathcal{S}_i$ and $\mathcal{S}_{-i} = \prod_{j\neq i} \mathcal{S}_j$. Similarly for $s \in \mathcal{S}$, s_i (resp. s_{-i}) denotes the projection of s on $\mathcal{S}_i$ (resp. $\mathcal{S}_{-i}$). Finally $\Delta(\mathcal{S})$ denotes the set of probability distributions over $\mathcal{S}$.

Downstream Game We consider a game of incomplete information between n players, depending on an unknown parameter θ (the *state of the world*) and denote by Θ the set of all possible states.

Each player $i \in [n]$ is described by a set of *types* $\mathcal{V}_i$, a set of actions $\mathcal{A}_i$, and a utility function $u_i : \mathcal{A} \times \Theta \times \mathcal{V} \to \mathbb{R}$.

Information Structures The monopolist information seller chooses a set of signals $\mathcal{S}$, and a message (bid) space $\mathcal{B}$ to design a communication rule $\sigma : \Theta \times \mathcal{B} \to \Delta(\mathcal{S})$ and payment function $p : \Theta \times \mathcal{B} \to \mathbb{R}^n_{\geq 0}$. Given a vector of bids $b \in \mathcal{B}$ and state $\theta \in \Theta$, we write $\sigma(\,\cdot\,; \theta, b) : \mathcal{S} \to [0, 1]$ for the corresponding probability distribution over $\mathcal{S}$.

The players' utility functions $(u_i)_{i \in [n]}$, the mechanism (σ, p), as well as the joint distribution of the random variables $(\theta, V) \in \Theta \times \mathcal{V}$ are commonly known at the onset of the game.

Timing The interaction between the information seller and the players, and among the players in the downstream game, takes place as follows:

1. Each player $i \in [n]$ observes their type V_i and the information seller observes the state θ.

2. Each player reports a bid B_i to the information seller, where B_i is a V_i-measurable random variable in $\mathcal{B}_i$.

3. The information seller generates signals $S \in \mathcal{S}$ distributed as $\sigma(\theta, B)$ and reveals S_i to each player $i \in [n]$ in exchange for payment $p_i(\theta, B)$.

4. Each player i chooses an action $A_i \in \mathcal{A}_i$ which is (V_i, S_i)-measurable and obtains utility $u_i(A; \theta, V) - p_i(\theta, B)$.

Linear Payoffs Beginning with Section 3.2, we shall restrict attention to a specific instance of our problem. We assume that the random variables $(\theta, V_1, \ldots, V_n)$ are mutually independent and that the utility of buyer i only depends on V through V_i only, and it does so in a linear manner:

$$u_i(a; \theta, v) = v_i \cdot \pi_i(a; \theta)$$

for some function $\pi_i : \mathcal{A} \times \Theta \to \mathbb{R}$. In other words, the buyers have private-value, *payoff types* that capture their marginal valuations, and reveal nothing about the state of the world.

This setup allows us to describe settings that capture our motivating scenario (sponsored links or listing) above, such as the following examples.

Example 1 (Competing Locations)**.** *Consider a variant of Hotelling's product-location game. A unit mass of consumers are located along the real line according to a Gaussian distribution* $\mathcal{N}(\theta, 1)$*. As customary in discrete-choice models of product differentiation, a consumer's location equivalently represents her ideal product variety. Two competing firms* $i = 1, 2$ *(who do not observe the mean of the consumers' distribution* θ*), choose a single location (i.e., offer a single variety from a continuum of products)* $a_i \in \mathbb{R}$*. For a given action profile* $(a_1, a_2) \in \mathbb{R}^2$*, each firm captures the mass of consumers closest to it. Thus, if* $a_1 < a_2$*, the quantity sold by firm* 1 *is given by*

$$\pi_1(a_1, a_2; \theta) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{a_1+a_2}{2}} e^{-\frac{1}{2}(x-\theta)^2} dx.$$

If the two firms were homogeneous, this would describe a zero-sum game. However, each firm is privately informed about its unit profit margin, i.e., firm i's payoff in the downstream game is given by $v_i \cdot \pi_i$*. The asymmetry in the two firms' profitability levels affects the seller's incentives to design mechanism that reveal more or less precise information about* θ*.*

The following example describes an even simpler setup that is meant to capture the reduced-form competition between two predictors.

Example 2 (Competing Predictions)**.** *Suppose the state of the world follows a known Gaussian distribution* $\theta \sim \mathcal{N}(\mu, 1)$*. Each player wishes to minimize a quadratic loss function* $L_i(a_i; \theta) = -(a_i - \theta)^2$ *and externalities are linear, in the sense that* $\pi_i(a; \theta) = -L_i(a_i; \theta) - \lambda \cdot L_j(a_j; \theta)$*, where* $\mathcal{A}_1 = \mathcal{A}_2 = \mathbb{R}$ *and* $\lambda \in (0, 1)$*. Each player is also privately informed about the marginal value of scaling down her losses, i.e., her payoffs in the downstream game are again given by* $v_i \cdot \pi_i$*.*

Unlike in Example 1, where a player's ideal location choice depends both on the state and on her opponent's location, Example 2 does not involve strategic interaction among the players, because each one has a dominant strategy to choose an action that matches the (unknown) state θ. In Chapter 4, Sections 4.2 and 4.3, we shall fully solve for the welfare- and revenue-maximizing mechanisms for the downstream game described in Example 3, which is a binary state and action version of the strategic setting in Example 2.

Chapter 3

Incentive Compatibility and Solution Concept

3.1 Definitions

In the game described in Chapter 2, each data buyer makes two strategic decisions: (i) report a bid $B_i \in \mathcal{B}_i$ after observing their private type V_i, and (ii) take an action A_i in the downstream game after receiving signal $S_i \in \mathcal{S}_i$ from the information seller.

By the revelation principle for dynamic games, see e.g. Section 6.3 in [10], it is without loss of generality to assume that the seller's set of signals $\mathcal{S}_i$, is equal to the set of actions $\mathcal{A}_i$, and that the buyers' reports lie in their own type space, $\mathcal{V}_i$, instead of a general message space $\mathcal{B}_i$, as long as we consider incentive compatible mechanisms.

For such mechanisms, a signal from the seller can thus be thought of as recommending to the buyer which action to take in the downstream game. Henceforth, we therefore denote the seller's recommendation by A_i, and then buyer's choice of action by a_i. Incentive compatibility (below) requires each buyer to both report her true type and to follow the seller's recommendation.

Definition 1 (Incentive Compatibility)**.** *A mechanism* (σ, p) *is* incentive compatible *if, for* *each* $(v_i, v_i') \in \mathcal{V}_i^2$ *and for each deviation function* $\delta : \mathcal{A}_i \to \mathcal{A}_i$,

$$\mathbb{E}\big[u_i(A_i, A_{-i}; \theta, V) - p_i(\theta, V) \,|\, V_i = v_i, B_i = v_i\big] \geq \mathbb{E}\big[u_i(\delta(A_i), A_{-i}; \theta, V) - p_i(\theta, v_i', V_{-i}) \,|\, V_i = v_i, B_i = v_i'\big],$$

where A *is distributed as* $\sigma(\theta, B_i, V_{-i})$*. The input into the deviation function* δ *is the seller's recommended action, and its output is the action the buyer takes after receiving the recommendation.*

This definition of incentive compatibility is closely related to the one of [5, Section 3.1] in the context of information design with elicitation. Unlike in our setting, this previous definition did not consider any payment to the information designer.

Definition 1 requires the mechanism to be robust to *double deviations* in which the data buyer both misreports their private type and deviate from the seller's recommendation. This implies in particular, when considering single deviations, that the mechanism is both *truthful* and *obedient* as defined next.

Definition 2 (Obedience)**.** *A mechanism* (σ, p) *is* obedient *if players have no incentive to deviate from the action recommendation of the information seller assuming everyone reports their type truthfully. Formally, for each* $\delta : \mathcal{A}_i \to \mathcal{A}_i$ *and* $v_i \in \mathcal{V}_i$,

$$\mathbb{E}\big[u_i(A_i, A_{-i}; \theta, V) - p_i(\theta, V) \,|\, V_i = v_i\big] \geq \mathbb{E}\big[u_i(\delta(A_i), A_{-i}; \theta, V) - p_i(\theta, V) \,|\, V_i = v_i\big]$$

where A *is distributed as* $\sigma(\theta, V)$*—in particular, data buyer* i*'s report is truthful.*

Equivalently one can write obedience as: for each $(a_i, a_i') \in \mathcal{A}_i^2$ *and* $v_i \in \mathcal{V}_i$,

$$\mathbb{E}\big[u_i(a_i, A_{-i}; \theta, V) \,|\, V_i = v_i, A_i = a_i\big] \geq \mathbb{E}\big[u_i(a_i', A_{-i}; \theta, V) \,|\, V_i = v_i, A_i = a_i\big]\,,$$

where A *is distributed as* $\sigma(\theta, V)$*.*

The first expression shows data buyer i's strategic behavior before receiving the action recommendation when she intends to report her type in the first stage of the game. At this stage, the buyer's strategy specifies a course of action following any action recommendation from the seller. Obedience requires that no deviations $\delta : \mathcal{A}_i \to \mathcal{A}_i$ are more profitable than obedience, i.e., the identity mapping $id : \mathcal{A}_i \to \mathcal{A}_i$.

The second expression shows data buyer i's strategic behavior after receiving the action recommendation at the second stage, and expresses that no other action results in a better expected utility. As mentioned before, these two are equivalent.

The second expression (which assumes every player reports her type truthfully) shows that obedience is only a property of the downstream game and of the recommendation rule σ, which thus correlates the actions of the data buyers. The distribution of actions resulting from an obedient recommendation rule in a game of incomplete information is a *Bayes correlated equilibrium* as defined and studied in [4, 5].

Definition 3 (Truthfulness)**.** *A mechanism is* truthful *if players have no incentive to misreport their type, assuming that everyone follows the seller's recommendations in the downstream game. Formally, for each* $(v_i, v_i') \in \mathcal{V}_i^2$*,*

$$\mathbb{E}\big[u_i(A;\theta,V) - p_i(\theta,V_i',V_{-i}) \,|\, V_i = v_i, B_i = v_i\big] \geq \mathbb{E}\big[u_i(A;\theta,V) - p_i(\theta,V_i',V_{-i}) \,|\, V_i = v_i, B_i = v_i'\big]$$

where A *is distributed as* $\sigma(\theta, B_i, V_{-i})$*.*

As already mentioned, incentive compatibility implies both obedience and truthfulness, but the converse is not true in general. However, as we will see in Chapter 4, when private types are independent linear valuation coefficients, the converse holds and incentive compatibility is thus equivalent to having obedience and truthfulness.

The externalities across the players in the downstream game imply that the recommendations given by the data seller affects the players' utilities even when they choose not to participate in the mechanism. This is because players engage in downstream competition even when they acquire no information from the seller, and possibly face fiercer competition when other players receive additional information instead.

Therefore, specifying the players' individual rationality constraints in our setting requires characterizing the *outside option* of the buyers, which results from the information structure chosen by the seller when a player chooses not to participate in the mechanism. Regardless of her objective function, it is optimal for the seller to choose the worst possible information structure for a non-participating player. We then have the following definition.

Definition 4 (Individual Rationality)**.** *For each $v_i \in \mathcal{V}_i$,*

$$\mathbb{E}\big[u_i(A;\theta,V) - p_i(\theta,V)\,|\,V_i = v_i\big] \geq \min_{\sigma_{-i}} \max_{a_i \in \mathcal{A}_i} \mathbb{E}\big[u_i(a_i, A_{-i};\theta,V)\,|\,V_i = v_i\big]$$

where the minimization on the right-hand side is over all communication rules $\sigma_{-i} : \Theta \times \mathcal{V}_{-i} \to \Delta(\mathcal{A}_{-i})$ and where A_{-i} is distributed as $\sigma_{-i}(\theta, V_{-i})$.

3.2 Characterization of Incentive Compatible and Truthful Mechanisms

For the rest of the work, we will focus on a specialized version of the model in Chapter 2 which covers settings such as those of Examples 1 and 2. In particular, we make the following assumption.

Assumption 1 (Linear Payoffs)**.** *The random variables $(\theta, V_1, \ldots, V_n)$ are mutually independent and for each $i \in [n]$, the utility of buyer i is given by*

$$u_i(a;\theta,v) = v_i \cdot \pi_i(a;\theta)$$

for some function $\pi_i : \mathcal{A} \times \Theta \to \mathbb{R}$.

With this formulation, the downstream game is interpreted as an imperfectly competitive market in which the outcome $(\pi_1, \ldots, \pi_n)$ determines the quantity sold by each data buyer, and their types represent their marginal value of making a sale (i.e., their unit profit margins). Because the downstream game determines market shares, we maintain the following assumption on the nature of competition.

Assumption 2 (Competitive Market)**.** *Players in the downstream game impose negative externalities on each other, that is,*

$$\pi_i(a_i, a_{-i};\theta) \geq \pi_i(a_i', a_{-i};\theta) \Longrightarrow \forall j \neq i \quad \pi_j(a_i, a_{-i};\theta) \leq \pi_j(a_i', a_{-i};\theta)\,,$$

for each $a_i, a_i' \in \mathcal{A}_i$, $a_{-i} \in \mathcal{A}_{-i}$ and state $\theta \in \Theta$.

We begin the analysis with a characterization of incentive compatibility (Definition 1). For the linear payoff model with independent types, the next proposition shows that incentive compatibility reduces to requiring truthfulness and obedience separately. In other words double deviations are not beneficial to the data buyers whenever a mechanism is immune to single deviations.

Proposition 1 (IC Characterization)**.** *Under Assumption 1, a mechanism is incentive compatible whenever it is truthful and obedient.*

Proof. Indeed we can write, using the notations of Definition 1, for each $(v_i, v_i') \in \mathcal{V}_i$:

$$\begin{aligned}
\mathbb{E}\big[u_i(A;\theta,V) - p_i(\theta,V)|V_i = v_i, B_i = v_i] &= \mathbb{E}\big[v_i \cdot \pi_i(A,\theta) - p_i(\theta,V)|V_i = v_i, B_i = v_i] \\
&\geq \mathbb{E}\big[v_i' \cdot \pi_i(A,\theta) - p_i(\theta,v_i',V_{-i})|V_i = v_i, B_i = v_i'] \\
&= \mathbb{E}\big[v_i' \cdot \pi_i(A,\theta) - p_i(\theta,v_i',V_{-i})|B_i = v_i'] \\
&= \mathbb{E}\big[u_i(A;\theta,B_i,V_{-i}) - p_i(\theta,B_i,V_{-i})|B_i = v_i'],
\end{aligned}$$

where the first equality is by Assumption 1, the inequality is by truthfulness (Definition 3), the third equality is because A is independent of V_i conditioned on B_i by Assumption 1, and the last equality uses the form of the payoffs from Assumption 1.

Moreover, note that since A is distributed as $\sigma(\theta, B_i, V_{-i})$, the last expectation in the previous derivation is exactly of the same form as the left-hand side in the first equation of Definition 2 for a vector of types (B_i, V_{-i}). Thus, by obedience for all $\delta : \mathcal{A}_i \to \mathcal{A}_i$,

$$\mathbb{E}\big[u_i(A;\theta,B_i,V_{-i}) - p_i(\theta,B_i,V_{-i})|B_i = v_i'] \geq \mathbb{E}\big[u_i(\delta(A_i),A_{-i},\theta) - p_i(\theta,B_i,V_{-i})|B_i = v_i'].$$

The previous two inequalities together imply incentive compatibility, which concludes the proof. □

In order to characterize truthful mechanisms, we follow the classical result of [9], which we restate in Proposition 2 below using our notation. In particular, recall that, in the general formulation of our model, the payment function p is allowed to depend on the realized state and on every player's reported type. However, note that we can restrict attention without loss of generality to interim-stage payments, i.e., where players' payments are a function of their own report only.

Thus, let (σ, p) be a mechanism and define for player $i \in [n]$, the interim share $\tilde{\pi}_i(V_i) \coloneqq \mathbb{E}[\pi_i(A;\theta) \mid V_i]$ and interim payment $\tilde{p}_i(V_i) \coloneqq \mathbb{E}[p_i(\theta, V) \mid V_i]$. We then have the following familiar characterization result. For completeness, a proof is provided in Appendix A.

Proposition 2 (Truthfulness Characterization)**.** *The mechanism (σ, p) is truthful if and only if for each player i:*

1. *The interim share $\tilde{\pi}_i$ is non-decreasing.*

2. *The interim payment $\tilde{p}_i$ is given for $v_i \in \mathcal{V}_i$ by*

$$\tilde{p}_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \underline{v} \cdot \tilde{\pi}_i(\underline{v}) + \tilde{p}_i(\underline{v}) - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds \,. \tag{3.1}$$

3.3 Characterization of Obedient Mechanisms

For the rest of the work, we will further restrict attention to the following example in order to characterize obedient mechanisms in this section and optimal mechanisms in Chapter 4 (for both the social planner and the monopolist seller).

Example 3 (Main Example)**.** *Consider two states of the world, $\Theta = \{0, 1\}$, two players, and two actions for each player in the downstream game, $\mathcal{A}_1 = \mathcal{A}_2 = \{0, 1\}$. The payoffs π_i in the downstream game are as follows:*

	0	*1*
0	$1-\alpha, 1-\alpha$	$1, -\alpha$
1	$-\alpha, 1$	$0, \quad 0$

$\theta = 0$

	0	*1*
0	$0, \quad 0$	$-\alpha, 1$
1	$1, -\alpha$	$1-\alpha, 1-\alpha$

$\theta = 1$

In each state of the world, it is a dominant strategy for each player to play the action matching that state. Furthermore, each player benefits from the other player choosing the wrong action. Formally, for $\{i, j\} = \{1, 2\}$,

$$\pi_i(a; \theta) = \mathbf{1}\{a_i = \theta\} - \alpha \mathbf{1}\{a_j = \theta\}. \tag{3.2}$$

A special case of this game occurs when $\alpha = 1$*. In this case we have a zero-sum game which is fully competitive. It can be viewed as a simplified version of the Hotelling model of Example 1, where every consumer is located in one of two locations* $\theta \in \{0, 1\}$*, which is unknown to the players. If both players locate themselves in a same place they will serve the consumer with probability* $1/2$*. Otherwise, the player in the correct location wins the consumer with probability* 1*. By varying* $\alpha \in [0, 1]$*, one can change the intensity of the competition between the firms.*

In addition to having a dominant strategy in each state of the world, a crucial property of the game in Example 3 is that the gain in utility when going from the suboptimal to the optimal action does not depend on the state nor on the action of the other player. This allows us to characterize obedient recommendation rules simply as those that recommend the optimal action to each player with sufficiently high probability.

In particular, the dominant strategy for each player in the absence of any signal about θ is to play the action corresponding to the most likely state under the common prior. For such a strategy,

$$\mathbb{P}[A_i = \theta \mid V_i] = \max\{\mathbb{P}[\theta = 0], \mathbb{P}[\theta = 1]\}.$$

Hence, the characterization of obedience in Proposition 3 below states that following their recommended action makes the players more likely to be correct than if they were simply basing their action on the common prior.

Proposition 3 (Obedience Characterization)**.** *Under Assumption 1 and for the downstream game given by Example 3, a recommendation rule* σ *is obedient if and only if for each player* $i \in \{1, 2\}$*, it holds almost surely that*

$$\mathbb{P}[A_i = \theta \mid V_i] \geq \max\{\mathbb{P}[\theta = 0], \mathbb{P}[\theta = 1]\}.$$

Proof. Since there are only two actions, $a_i \in \{0, 1\}$, we can rewrite obedience as

$$\mathbb{E}[\pi_i(a_i, A_j; \theta) - \pi_i(1 - a_i, A_j; \theta) \,|\, A_i = a_i, V_i] \geq 0, \tag{3.3}$$

for $\{i, j\} = \{1, 2\}$ and $a_i \in \{0, 1\}$. Using (3.2), we have that almost surely

$$\begin{aligned}\pi_i(a_i, A_j; \theta) - \pi_i(1 - a_i, A_j; \theta) &= \mathbf{1}\{a_i = \theta\} - \alpha\mathbf{1}\{A_j = \theta\} - \big(\mathbf{1}\{1 - a_i = \theta\} - \alpha\mathbf{1}\{A_j = \theta\}\big) \\ &= \mathbf{1}\{a_i = \theta\} - \mathbf{1}\{a_i = 1 - \theta\}\,.\end{aligned}$$

where the last expression crucially does not depend on A_j. Hence, obedience is equivalent to

$$\mathbb{P}[\theta = a_i \,|\, A_i = a_i, V_i] - \mathbb{P}[\theta = 1 - a_i \,|\, A_i = a_i, V_i] \geq 0\,,$$

for $i \in \{1, 2\}$ and $a_i \in \{0, 1\}$.

By Bayes' rule and independence of θ and V_i, this is equivalent to

$$\mathbb{P}[\theta = a_i] \cdot \mathbb{P}[A_i = a_i \,|\, \theta = a_i, V_i] \geq \mathbb{P}[\theta = 1 - a_i] \cdot \mathbb{P}[A_i = a_i \,|\, \theta = 1 - a_i, V_i]\,,$$

which is in turn equivalent to $\mathbb{P}[A_i = \theta \,|\, V_i] \geq \mathbb{P}[\theta = 1 - a_i]$. Since the left-hand side no longer depends on a_i, we can combine the two constraints when $a_i = 0$ and when $a_i = 1$ and obtain the statement of the lemma. □

Chapter 4

Welfare and Revenue Optimal Mechanisms

We now turn to the seller's objective and consider both social welfare and revenue maximization. We show below that, for the downstream game of Example 3, both objectives can be written as a weighted sum of the probabilities that the mechanism recommends the dominant strategy to each player (see Eq. (4.1) below). Hence, we first describe in Section 4.1 an optimal mechanism for a general class of objective functions of this form, which we then instantiate in Section 4.2 and Section 4.3 to derive mechanisms maximizing social welfare and revenue, respectively.

4.1 Optimal Mechanisms

We consider a general objective function of the form

$$\mathbb{E}\big[w_1(V)\mathbb{P}[A_1 = \theta \,|\, V] + w_2(V)\mathbb{P}[A_2 = \theta \,|\, V]\big] \tag{4.1}$$

for functions $w_1, w_2 : \mathcal{V} \to \mathbb{R}$. The form of this expression as well as the characterization of obedience obtained in Proposition 3 suggest that a convenient parametrization of the auctioneer's problem is in terms of the functions $h_i : V \mapsto \mathbb{P}[A_i = \theta \,|\, V]$ from $\mathcal{V}$ to $[0, 1]$ for $i \in \{1, 2\}$. These functions can easily be expressed in terms of the recommendation rule σ.

Indeed, we have almost surely

$$\begin{aligned}\mathbb{P}[A_1 = \theta \,|\, V] &= \mathbb{E}[\mathbf{1}\{A_1 = \theta\} \,|\, V] \\ &= \mathbb{E}[\mathbf{1}\{A_1 = \theta, A_2 = 1-\theta\} + \mathbf{1}\{A_1 = \theta, A_2 = \theta\} \,|\, V] \\ &= \mathbb{E}[\sigma(\theta, 1-\theta; \theta, V_1, V_2) \,|\, V] + \mathbb{E}[\sigma(\theta, \theta; \theta, V_1, V_2) \,|\, V]\end{aligned}$$

and similarly for $\mathbb{P}[A_2 = \theta \,|\, V]$. Conversely the following lemma shows how to construct a recommendation rule which has h_1 and h_2 as its marginals.

Lemma 1 (Recommendation Rule from Marginals)**.** *Let h_1 and h_2 be two measurable functions from $\mathcal{V}$ to $[0,1]$, then there exists a recommendation rule $\sigma : \Theta \times \mathcal{V} \to \Delta(\mathcal{A})$ such that, almost surely for $i \in \{1,2\}$,*

$$\mathbb{P}[A_i = \theta \,|\, V] = h_i(V)\,.$$

Proof. Given two functions h_1 and h_2 satisfying the lemma's assumptions, one can choose for example σ such that for all $(x, v_1, v_2) \in \Theta \times \mathcal{V}_1 \times \mathcal{V}_2$, the distribution $\sigma(x, v_1, v_2) \in \Delta(\mathcal{A})$ has independent coordinates with marginals given by h_1 and h_2 respectively. Formally, we have

$$\sigma(a_1, a_2; x, v_1, v_2) = \begin{cases} h_1(v_1, v_2)h_2(v_1, v_2) & \text{if } (a_1, a_2) = (x, x) \\ h_1(v_1, v_2)\big(1 - h_2(v_1, v_2)\big) & \text{if } (a_1, a_2) = (x, 1-x) \\ \big(1 - h_1(v_1, v_2)\big)h_2(v_1, v_2) & \text{if } (a_1, a_2) = (1-x, x) \\ \big(1 - h_1(v_1, v_2)\big)\big(1 - h_2(v_1, v_2)\big) & \text{if } (a_1, a_2) = (1-x, 1-x), \end{cases}$$

which ends the proof. □

In other words, any choice of the marginal functions $\mathbb{P}[A_i = \theta \,|\, V]$ can be "realized" by a recommendation rule. Hence, as long as the criterion being optimized by the auctioneer and the constraints on the recommendation rule can be expressed in terms of $\mathbb{P}[A_i = \theta \,|\, V]$, we will directly optimize over these quantities. An optimal information structure σ in this class can then be obtained using Lemma 1.

We now describe a general recommendation rule that optimizes criteria of the form (4.1) (including welfare and revenue maximization) subject to obedience constraints. Our first result (Lemma 2 below) describes and solves the "core" of this optimization problem.

Lemma 2 (Variational Lemma)**.** *Let* $f : \mathbb{R} \to \mathbb{R}_{\geq 0}$ *be a probability density function and let* $F : \mathbb{R} \to [0,1]$ *be its associated cumulative density function. Let* $g : \mathbb{R} \to \mathbb{R}$ *be a non-increasing function such that* $g \cdot f$ *is absolutely integrable and define* $t_g \coloneqq \sup\{t \in \mathbb{R} \mid g(t) \geq 0\}$*. Let* $\mathcal{F}$ *be the set of measurable functions* $\mathbb{R} \to [0,1]$*, and* $c \in [0,1]$ *be a constant. Then, a solution to*

$$\max_{h \in \mathcal{F}} \; \mathcal{L}(h) \coloneqq \int_{\mathbb{R}} h(v)g(v)f(v)dv$$
$$s.t. \; \int_{\mathbb{R}} h(v)f(v)dv \geq c$$

is given by $h^\star : v \mapsto \mathbf{1}\{v \leq v^\star\}$ *where* $v^\star = \max\{F^{-1}(c), t_g\}$*.*

Proof. We note that under the assumptions of the lemma F is non-decreasing and absolutely continuous hence it admits a right inverse and $q_c \coloneqq F^{-1}(c)$ is well-defined.

We first show that $h^\star$ is feasible. Indeed, since $v^\star \geq q_c$ and F is non-decreasing

$$\int_{\mathbb{R}} h^\star(v)f(v)dv = \int_{-\infty}^{v^\star} f(v)d(v) = F(v^\star) \geq F(q_c) = c\,.$$

We conclude by showing that $\mathcal{L}(h^\star) - \mathcal{L}(h) \geq 0$ for all feasible h by distinguishing two cases depending on the relative position of q_c and t_g.

1. If $q_c \leq t_g$, equivalently $v^\star = t_g$, then

$$\mathcal{L}(h^\star) - \mathcal{L}(h) = \int_{t_g}^{\infty} (-h)(v)g(v)f(v)dv + \int_{-\infty}^{t_g} (1-h)(v)g(v)f(v)dv \geq 0\,.$$

   Indeed, since g is non-increasing, $g(v) \geq 0$ for $v < t_g$ and $g(v) \leq 0$ for $v > t_g$. Since $0 \leq h(v) \leq 1$ for all $v \in \mathbb{R}$, this implies that both integrals are non-negative.

2. If $q_c > t_g$, equivalently $v^\star = q_c$. Then

$$\begin{aligned}\mathcal{L}(h^\star) - \mathcal{L}(h) &= \int_{q_c}^{\infty} (-h)(v)g(v)f(v)dv + \int_{-\infty}^{q_c} (1-h)(v)g(v)f(v)dv \\ &\geq g(q_c)\int_{q_c}^{\infty} (-h)(v)f(v)dv + g(q_c)\int_{-\infty}^{q_c} (1-h)(v)f(v)dv \\ &= g(q_c)F(q_c) - g(q_c)\int_{\mathbb{R}} h(v)f(v)dv \\ &\geq g(q_c)F(q_c) - g(q_c)c = 0,\end{aligned}$$

where the first inequality uses that g is non-increasing and $0 \leq h \leq 1$, and the last inequality uses that h is feasible and $g(q_c) < 0$.

□

We are now ready to describe our optimal mechanism.

Proposition 4 (Optimal Mechanism)**.** *Assume that V_1 and V_2 are identically and independently distributed with absolutely continuous cumulative distribution function F and write $v^\star := F^{-1}\big(\max\{\mathbb{P}[\theta = 0], \mathbb{P}[\theta = 1]\}\big)$. Consider the objective*

$$W := \mathbb{E}\big[w_1(V)\mathbf{1}\{A_1 = \theta\} + w_2(V)\mathbf{1}\{A_2 = \theta\}\big]$$

where $w_1, w_2 : \mathcal{V} \to \mathbb{R}$ are functions such that for all $(v_1, v_2) \in \mathcal{V}$, $w_1(v_1, \cdot)$ and $w_2(\cdot, v_2)$ are non-increasing functions. Define $v_1^\star, v_2^\star : \mathcal{V} \to \mathbb{R}$ by

$$v_2^\star(v_1) = \sup\{v_2 \in \mathbb{R} \,|\, w_1(v_1, v_2) \geq 0\} \quad \textit{and} \quad v_1^\star(v_2) = \sup\{v_1 \in \mathbb{R} \,|\, w_2(v_1, v_2) \geq 0\}$$

with the usual convention $\sup \emptyset = -\infty$. Then the recommendation rule characterized[1] *by*

$$\mathbb{P}[A_i = \theta \,|\, V] = \mathbf{1}\big\{V_j \leq \max\{v_j^\star(V_i), v^\star\}\big\}$$

for all $\{i, j\} = \{1, 2\}$, maximizes W subject to obedience.

[1] Note that this recommendation rule is defined in terms of the functions $V \mapsto \mathbb{P}[A_i = \theta \,|\, V]$ for $i \in \{1, 2\}$. Recall that a recommendation rule compatible with these functions can be constructed using Lemma 1. Furthermore since the functions in fact take values in $\{0, 1\}$, it is easy to see that such a rule is in fact uniquely defined.

Proof. Define $\mu \coloneqq \max\{\mathbb{P}[\theta=0], \mathbb{P}[\theta=1]\}$ so that $v^\star = F^{-1}(\mu)$ and define $h_i : \mathcal{V} \to [0,1]$ by $h_i(V) = \mathbb{P}[A_i = \theta \mid V]$. By Proposition 3 and using the law of total expectation, obedience can be written as, for $i \in \{1,2\}$ and almost surely

$$\mathbb{P}[A_i = \theta \mid V_i] = \mathbb{E}[\mathbf{1}\{A_i = \theta\} \mid V_i] = \mathbb{E}\big[\mathbb{E}[\mathbf{1}\{A_i = \theta\} \mid V] \mid V_i\big] = \mathbb{E}[h_i(V) \mid V_i] \geq \mu\,.$$

Similarly we can rewrite the objective in terms of the functions h_i

$$W = \mathbb{E}\big[w_1(V)\mathbb{P}[A_1 = \theta \mid V] + w_2(V)\mathbb{P}[A_2 = \theta \mid V]\big] = \mathbb{E}\big[w_1(V)h_1(V) + w_2(V)h_2(V)\big]\,.$$

Hence the optimization problem we need to solve is

$$\begin{aligned}\max\ & \mathbb{E}\big[w_1(V)h_1(V) + w_2(V)h_2(V)\big]\\ \text{s.t. } & \mathbb{E}[h_i(V) \mid V_i] \geq \mu,\ \text{ for } i \in \{1,2\}.\end{aligned}$$

This problem decomposes as the sum of two optimization problems, one for each h_i, $i \in \{1,2\}$. We focus on the one for h_1, the other one being identical after swapping 1 and 2. Writing the expectations in terms of the density f associated with F, the problem is

$$\begin{aligned}\max\ & \int_{\mathbb{R}} \left(\int_{\mathbb{R}} w_1(v)h_1(v)f(v_2)dv_2\right) f(v_1)dv_1\\ \text{s.t. } & \int_{\mathbb{R}} h_1(v_1,v_2)f(v_2)dv_2 \geq \mu\,, \text{for all } v_1 \in \mathbb{R}\,.\end{aligned}$$

Since the constraint on h_1 is an integral constraint on the partial function $h_1(v_1,\cdot)$ for each v_1, an optimal solution is obtained by choosing for each $v_1 \in \mathbb{R}$, $h_1(v_1,\cdot)$ maximizing the inner integral $\int_{\mathbb{R}} w_1(v_1,v_2)h_1(v_1,v_2)f(v_2)dv_2$ subject to this integral constraint. But since $w_1(v_1,\cdot)$ is non-increasing by assumption, this optimization problem over $h_1(v_1,\cdot)$ has exactly the form solved in Lemma 2, from which we get the solution

$$h_1(v_1,v_2) = \mathbf{1}\big\{v_2 \leq \max\{v_2^\star(v_1), v^\star\}\big\}\,,$$

and similarly for h_2, which concludes the proof by definition of h_1 and h_2. $\square$

4.2 Welfare Maximization

Using (3.2) we can write the expected social welfare for the game of Example 3 as

$$W = \mathbb{E}\big[(V_1 - \alpha V_2)\mathbf{1}\{A_1 = \theta\} + (V_2 - \alpha V_1)\mathbf{1}\{A_2 = \theta\}\big] ,$$

which is of the form solved by Proposition 4. We then obtain the following characterization of the welfare-maximizing (second best) mechanism.

Proposition 5 (Welfare Optimal Mechanism)**.** *Assume that V_1 and V_2 are identically distributed with absolutely continuous c.d.f. F and write $v^\star := F^{-1}\big(\max\{\mathbb{P}[\theta = 0], \mathbb{P}[\theta = 1]\}\big)$. A recommendation rule maximizing social welfare subject to obedience is given by*

$$(A_1, A_2) = \begin{cases} (\theta, \theta) & \text{if } V_1 \le v^\star \text{ and } V_2 \le v^\star \\ (\theta, 1-\theta) & \text{if } V_1 \ge \max\{v^\star, V_2/\alpha\} \\ (1-\theta, \theta) & \text{if } V_2 \ge \max\{v^\star, V_1/\alpha\} . \end{cases}$$

Proof. We apply Proposition 4 with $w_1(v_1, v_2) = (v_1 - \alpha v_2)$ and $w_2(v_1, v_2) = (v_2 - \alpha v_1)$ for which $v_2^\star(v_1) = v_1/\alpha$ and $v_1^\star(v_2) = v_2/\alpha$. The optimal mechanism is thus characterized by

$$\mathbb{P}[A_i = \theta \,|\, V] = \mathbf{1}\big\{V_j \le \max\{v^\star, V_i/\alpha\}\big\} .$$

Since these probabilities take values in $\{0, 1\}$, we obtain a partition of the type space $\mathcal{V}$ into regions in which the recommended action profile (A_1, A_2) is constant given (θ, V_1, V_2). □

We remark on several noteworthy properties of the welfare-optimal (second best) mechanism. First, the recommended action profile is deterministic given both players' types and the state θ. This appears at first glance to be in contradiction with the obedience constraint, since a player could infer the state from their recommendation and thus deviate when being recommended the suboptimal action. This apparent contradiction disappears when one remembers that the players only observe their own type and have to reason in expectation over the other player's type. Hence, the seller is able to exploit a player's uncertainty about

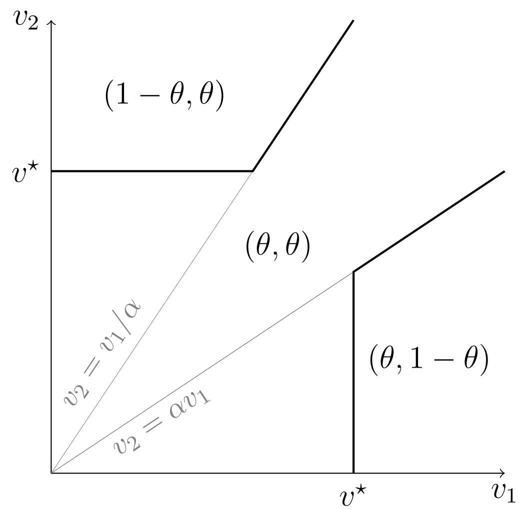

Figure 4-1: Representation of the second-best recommendation rule from Proposition 5 for $\alpha = 2/3$, showing the recommended action profile as a function of both players' types.

the other player's type to "hide" the true state of the world and recommend the suboptimal action with positive probability (across realizations of the other player's type) without violating obedience constraints.

Second, a simple computation shows that

$$\mathbb{P}[A_i = \theta \,|\, V_i] = \max\{F(v^\star), F(V_i/\alpha)\},$$

which indicates that the obedience constraint is binding for player i whenever $V_i \leq \alpha v^\star$. The socially efficient (first best) allocation would instead set $(A_1, A_2) = (\theta, 1-\theta)$ when $V_1 \geq V_2/\alpha$ and $(A_1, A_2) = (1-\theta, \theta)$ when $V_2 \geq V_1/\alpha$, and finally $(A_1, A_2) = (\theta, \theta)$ when $\alpha V_1 \leq V_2 \leq V_1/\alpha$. Therefore, the area in Figure 4-1 where $V_i \leq v^\star$ for $i \in \{1, 2\}$ and $\max_i V_i > \min_i V_i/\alpha$ captures the loss in efficiency due to obedience constraints.

Finally, it is easy to verify that the second best mechanism is implementable, i.e. satisfies the players' truth-telling constraints. Indeed, by Proposition 2 it suffices to verify that the interim share of the data buyers are non-decreasing.

Proposition 6 (Truthfulness of Second Best)**.** *For the mechanism of Proposition 5, the interim share* $\tilde{\pi}_i(V_i) = \mathbb{E}\big[\pi_i(A, \theta) \,|\, V_i\big]$ *of player* $i \in \{1, 2\}$ *is non-decreasing. Hence the second best mechanism is implementable.*

Proof. We have

$$
\begin{aligned}
\tilde{\pi}_i(V_i) &= \mathbb{E}\big[\mathbb{P}[A_i = \theta \,|\, V] - \alpha \mathbb{P}[A_j = \theta \,|\, V] \,|\, V_i\big] \\
&= \mathbb{E}\big[\mathbf{1}\big\{V_j \le \max\{v^\star, V_i/\alpha\}\big\} \,|\, V_i\big] - \alpha \mathbb{E}\big[\mathbf{1}\big\{V_i \le \max\{v^\star, V_j/\alpha\}\big\} \,|\, V_i\big] \\
&= \max\{F(v^*), F(V_i/\alpha)\} - \alpha\big(1 - \mathbb{P}[V_i > v^* \wedge V_j < \alpha V_i \,|\, V_i]\big) \\
&= \max\{F(v^*), F(V_i/\alpha)\} + \alpha \mathbf{1}\{V_i > v^*\} \cdot F(\alpha V_i) - \alpha\,,
\end{aligned}
$$

where the first equality uses (3.2) and the law of total expectation, the second equality uses the characterization of the second best mechanism from Proposition 5 and the remaining two equalities are basic algebra. The final expression is clearly non-decreasing in V_i as sum and product of non-negative non-decreasing functions. ◻

Intuitively, a higher type is revealed the correct state more often by the social planner, which makes it possible to find transfers that would induce truthful reporting of the players' types. Of course these transfers do not correspond to a monopolist data seller's optimal choice. In the next section, we will see how a monopolist data seller modifies the second best mechanism to maximize the associated payments.

4.3 Revenue Maximization

Throughout this section, we assume that V_1 and V_2 are identically distributed with absolutely continuous cumulative distribution function F. Denoting by f the associated probability density function, we define the virtual value $\phi : \mathcal{V} \to \mathbb{R}$ as

$$\phi(v) := v - \frac{1 - F(v)}{f(v)},$$

and call a distribution F *regular* if the associated virtual values ϕ are non-decreasing.

Recall that the outside option determining data buyer i's participation constraint (Definition 4) is given by the expected quantity $\tilde{\pi}_i$ sold in the worst-case scenario—where all her competitors learn the state with probability one—scaled by the private type v_i. Despite this dependency on the private types, we show in Lemma 3 below that the seller's expected revenue $R \coloneqq \mathbb{E}\big[\sum_i p_i(V)\big]$ can be expressed in terms of the virtual valuations, as in the standard Myersonian auction [9]. The proof is given in Appendix A.

Lemma 3 (Reduction to Virtual Surplus)**.** *Assuming truthfulness and individual rationality, optimal allocation for maximizing virtual social surplus is also the solution to the seller's revenue maximization problem. Moreover,*

$$\tilde{p}_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - K \cdot \underline{v} - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s)ds\,, \tag{4.2}$$

where $K \coloneqq \min_{\sigma_j \in \Delta(\mathcal{A}_j)} \max_{a_i \in \mathcal{A}_i} \mathbb{E}\big[\pi_i(a_i, A_j; \theta)\big]$ *is the expected market share in a player's outside option (recall we assume that the game is symmetric, hence* K *does not depend on* i*).*

Hence, we now focus on maximizing the virtual surplus which can be written using (3.2) in terms of the recommendation to each player,

$$R^{\dagger} = \mathbb{E}\big[\big(\phi(V_1) - \alpha\phi(V_2)\big)\mathbf{1}\{A_1 = \theta\} + \big(\phi(V_2) - \alpha\phi(V_1)\big)\mathbf{1}\{A_2 = \theta\}\big].$$

We can then apply Proposition 4 to characterize the optimal information structure.

Proposition 7 (Revenue Optimal Mechanism)**.** *For* $\alpha \in [0,1]$*, if the distributions of* V_1 *and* V_2 *are identical and regular, the recommendation rule maximizing virtual surplus subject to obedience is given by*

$$\mathbb{P}[A_i = \theta \,|\, V] = \mathbf{1}\big\{V_j \le \max\{v^\star, \phi^{-1}(\phi(V_i)/\alpha)\}\big\}\,.$$

Proof. This follows from Proposition 4 applied to $w_1(v_1, v_2) = \phi(v_1) - \alpha\phi(v_2)$ and $w_2(v_1, v_2) = \phi(v_2) - \alpha\phi(v_1)$. By regularity, $w_1(v_1, \cdot)$ and $w_2(\cdot, v_2)$ are non-increasing for all $(v_1, v_2) \in \mathcal{V}$ and $v_i^\star(v_j) = \phi^{-1}(\phi(v_j)/\alpha)$ for $\{i, j\} = \{1, 2\}$. □

In contrast to the welfare-maximization case, the recommendation rule described in Proposition 7 can potentially partition the type space into more than three regions when $\alpha < 1$. To understand this, let us first ignore the obedience constraint, in which case the recommendation rule maximizing virtual surplus is determined by

$$\mathbb{P}[A_i = \theta \mid V] = \mathbf{1}\big\{V_j \leq \phi^{-1}(\phi(V_i)/\alpha)\big\}\,.$$

Since the functions $v \mapsto \phi^{-1}(\phi(v)/\alpha)$ and $v \mapsto \phi^{-1}(\alpha\phi(v))$ intersect at $v_0 := \phi^{-1}(0)$, this recommendation rule partitions the type space into four regions. Crucially, there is now a new region, defined by $\phi^{-1}(\phi(v_1)/\alpha) \leq v_2 \leq \phi^{-1}(\alpha\phi(v_1))$ in which both players receive the wrong action recommendation.

The presence of this region is intuitive, because the seller finds it profitable to distort the allocation (i.e., to lower the quality of her recommendations) to players with low types, so to reduce information rents and increase the payments of higher types. However, unlike settings without externalities, merely having a negative virtual value does not imply a player receives the wrong information. Even absent obedience constraints, the seller knows that distorting one player's recommendations increases the surplus of the other player. Therefore, both players receive the wrong recommendation only if both their virtual values are negative *and* they are sufficiently similar. Conversely, if both $\phi(V_i) < 0$ but V_1 is sufficiently larger than V_2, then the seller prefers issuing the correct recommendation to player 1. Indeed, distorting the recommendation to player 1 would increase player 2's payoff, which has an even stronger negative impact on the seller's profits.

After introducing the obedience constraints, however, one needs to further intersect these regions with the lines $v_2 = v^\star$ and $v_1 = v^\star$ and enforce that $A_i = \theta$ when $V_j \leq v^\star$. Depending on the type distribution, it might be that $v^\star \geq v_0$, in which case the region over which both players receive the wrong recommendation is empty.[2]

Therefore, when $v^\star \geq v_0$, the square where $v_1, v_2 \leq v^\star$, in which obedience requires recommending the optimal action to both players, fully contains the aforementioned region, and the situation is qualitatively the same as in Figure 4-1. However, when $v^\star \leq v_0$, the type space is now partitioned into five regions as in Figure 4-2. Notice that obedience binds for all $v_i < \tilde{v}$ as for those types $\mathbb{P}\big[A_i = \theta \mid V_i\big] = F(v^*)$.

To summarize, the expected quantity sold $\tilde{\pi}_i$ of data buyer i in the revenue-optimal

[2]For example, when the types are uniformly distributed over $[0, 1]$, we have by definition $P(v^\star) = v^\star = \max\{\mathbb{P}[\theta = 0], \mathbb{P}[\theta = 1]\} \geq 1/2$, but $\phi : v \mapsto 2v - 1$ and $v_0 = 1/2$.

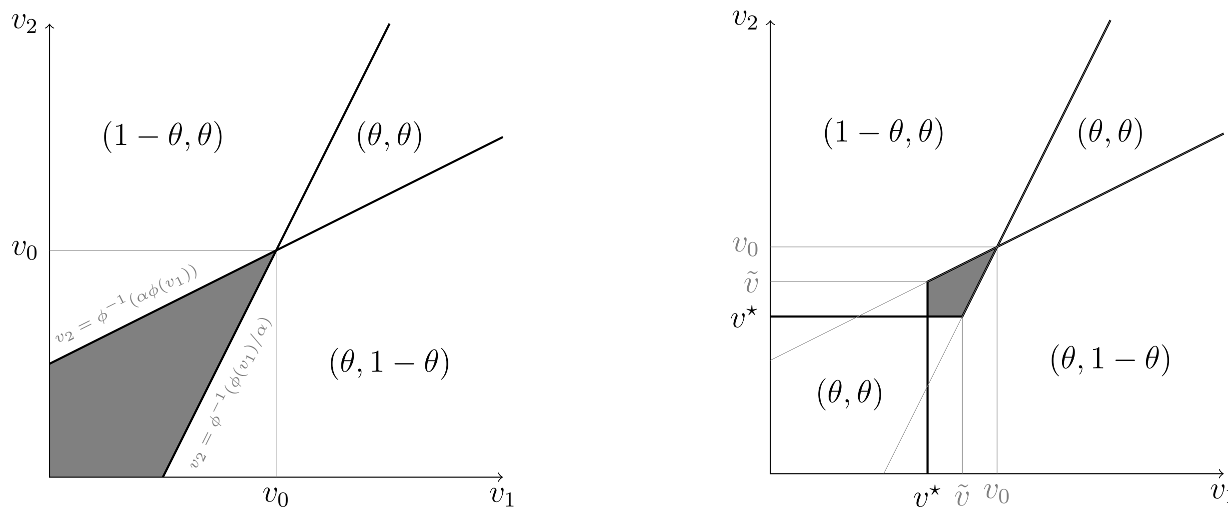

Figure 4-2: Virtual surplus-maximizing recommendation rule from Proposition 7 with $\alpha = 1/2$ and equally likely states. The types are distributed exponentially with parameter 1, so that $\phi(v) = v - 1$ and $v_0 = 1$. On the left, the surplus-maximizing rule ignoring the obedience constraint; on the right with the obedience constraint ($v^\star = \ln 2$). In both panels, the gray area is the region where the suboptimal action $1-\theta$ is recommended to both players.

mechanism is given by the following function. Define $\tilde{v} := \phi^{-1}(\alpha\phi(v^*))$, then

- If $v_0 < v^*$:

$$\tilde{\pi}_i(v_i) = \begin{cases} F(v^*) - \alpha & v_i < \tilde{v} \\ F\big(\phi^{-1}(\frac{\phi(v_i)}{\alpha})\big) - \alpha & \tilde{v} < v_i < v^* \\ F\big(\phi^{-1}(\frac{\phi(v_i)}{\alpha})\big) - \alpha + \alpha F\big(\phi^{-1}(\alpha\phi(v_i))\big) & v^* < v_i. \end{cases}$$

- If $v^* < v_0$:

$$\tilde{\pi}_i(v_i) = \begin{cases} F(v^*) - \alpha & v_i < v^* \\ F(v^*) - \alpha + \alpha F\big(\phi^{-1}(\alpha\phi(v_i))\big) & v^* < v_i < \tilde{v} \\ F(\phi^{-1}(\frac{\phi(v_i)}{\alpha})) - \alpha + \alpha F\big(\phi^{-1}(\alpha\phi(v_i))\big) & \tilde{v} < v_i. \end{cases}$$

We now show this share is increasing, i.e., the obedient mechanism above is also *truthful (implementable)*, and therefore we have the optimal solution.

Proposition 8 (Truthfulness of Optimal Mechanism)**.** *For the mechanism of Proposition 7 and under the same assumptions, the interim share* $\tilde{\pi}_i(V_i) = \mathbb{E}\big[\pi_i(A, \theta) \,|\, V_i\big]$ *of each player* $i \in \{1, 2\}$ *is non-decreasing. Hence the revenue-maximizing mechanism is implementable and the payments are given by*

$$\tilde{p}_i(v_i) = v_i \tilde{\pi}_i(v_i) - \underline{v}\tilde{\pi}_i(\underline{v}) - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds \,. \tag{4.3}$$

Also by comparing eq. (4.3) *to eq.* (4.2)*, we find out that* $\tilde{\pi}_i(\underline{v}) = K$ *and* $\tilde{p}_i(\underline{v}) = 0$*.*

Proof. The proof is almost identical to the one of Proposition 6. Again we write,

$$\begin{aligned}
\tilde{\pi}_i(V_i) &= \mathbb{E}\big[\mathbb{P}[A_i = \theta \,|\, V] - \alpha \mathbb{P}[A_j = \theta \,|\, V] \,|\, V_i\big] \\
&= \mathbb{E}\big[\mathbf{1}\big\{V_j \leq \max\{v^\star, \phi^{-1}(\phi(V_i)/\alpha)\}\big\} \,|\, V_i\big] - \alpha \mathbb{E}\big[\mathbf{1}\big\{V_i \leq \max\{v^\star, \phi^{-1}(\phi(V_j)/\alpha)\}\big\} \,|\, V_i\big] \\
&= \max\{F(v^*), F(\phi^{-1}(\phi(V_i)/\alpha)\} - \alpha\big(1 - \mathbb{P}[V_i > v^* \wedge V_j \leq \phi^{-1}(\alpha\phi(V_i)) \,|\, V_i]\big) \\
&= \max\{F(v^*), F(\phi^{-1}(\phi(V_i)/\alpha)\} + \alpha \mathbf{1}\{V_i > v^*\} \cdot F\big(\phi^{-1}(\alpha\phi(V_i))\big) - \alpha \,,
\end{aligned}$$

where this last expression is non-decreasing in V_i as sum and product of non-decreasing functions, since ϕ is non-decreasing by regularity.

For the payment, by (3.1)

$$\tilde{p}_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \underline{v} \cdot \tilde{\pi}_i(\underline{v}) + \tilde{p}_i(\underline{v}) - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds$$

where $\tilde{p}_i(\underline{v}) = 0$ as in eq. (4.2) we have $\tilde{p}_i(\underline{v}) = \underline{v}(\tilde{\pi}_i(\underline{v}) - K)$. But we can see that $K = \tilde{\pi}_i(\underline{v}) = F(v^*) - \alpha$.

$$\begin{aligned}
K &= \max_{a_i \in \mathcal{A}_i} \mathbb{E}\big[\pi_i(a_i, \theta; \theta)\big] = \max_{a_i \in \mathcal{A}_i} \mathbb{E}\big[\mathbf{1}\{a_i = \Theta\}\big] - \alpha \\
&= \max\big\{\mathbb{P}\{\theta = 0\}, \mathbb{P}\{\theta = 1\}\big\} - \alpha = F(v^*) - \alpha
\end{aligned}$$

Which is equal to $\tilde{\pi}_i(\underline{v})$ as it is computed above the Proposition 8. ∎

We remark on two striking properties of the optimal payments, which apply whenever $v^* < v_0$. First, some types of player i with a negative virtual valuation, $\tilde{v} < v_i < v_0$, are nonetheless charged a positive payment. This occurs because these types are sufficiently high that their opponent j has an even lower type v_j with a significant probability, $F(v_i)$. In other words, the seller finds it optimal to reveal the correct state to player i with probability, $F(\phi^{-1}(\phi(v_i)/\alpha)) > F(v^*)$. Player i then has a strict incentive to follow the seller's recommendation—her obedience constraint is slack.

Second, even types of player i such that $v^* < v_i < \tilde{v}$, and whose obedience constraint binds, pay a strictly positive price. Because their obedience constraint is binding, these types derive no net utility from following the seller's recommendation. However, unlike types in $[0, v^*]$ where the other data buyer always receives the right recommendation, these types' opponent is revealed the correct state with probability $1 - F(\phi^{-1}(\alpha\phi(v_i)))$. Therefore, these types are strictly better off participating, and they can be charged a positive payment. Put differently, the presence of negative externalities augments the profitability of selling information, as the seller charge positive payments in exchange for limiting the information available to each buyer's competitors.

Impact of the competitiveness of the game Finally, we investigate the role of downstream competition intensity on the revenue-optimal mechanism. Figure 4-3 compares two settings, where competition is fiercer in the left panel ($\alpha = 1/2$) than in the right panel ($\alpha = 1/4$).

Reducing the intensity of competition reduces the value of exclusive sales of information (i.e., recommending the right action to one player only) in the first best: at one extreme, if players imposed no externalities on each other, the seller would recommend the right action to any player with a positive virtual value. In particular, in an unconstrained revenue problem, the seller would recommend the wrong action to both players more often when competition is weaker. However, this recommendation profile would violate obedience, which requires the seller to recommend the right action to both players in the square where their types is smaller than $v^\star$. As $v^\star$ is independent of α, the right panel shows how the seller resorts to exclusive sales as a second-best policy under obedience constraints, and does so more often

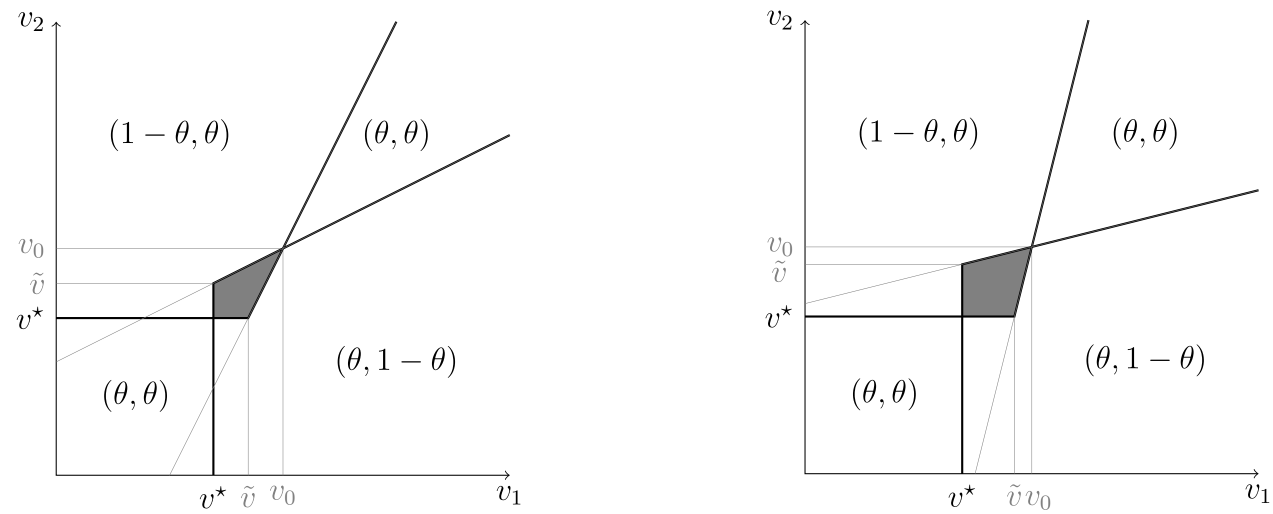

(a) $\alpha = 1/2$ (b) $\alpha = 1/4$

Figure 4-3: Comparison of the revenue-maximizing recommendation rules from Proposition 7 for two different values of α. As in Figure 4-2 the types are exponentially distributed and the states are equally likely. Recall that the larger the value of α, the more competitive the downstream game.

as the competition weakens.

Chapter 5

Future Work and Conclusions

We have begun to explore the implications of selling information to competing players in a mechanism design framework. We have shown that the nature of information disciplines the optimal selling mechanisms for data products, and distinguishes them from canonical (e.g., physical) goods.

In particular, the players' actions in the downstream game introduce a moral hazard problem for the designer's choice of mechanism. The resulting obedience constraints on the information structure prevent a social planner from implementing the efficient level of information exclusivity—the second best involves symmetric information more often than optimal. At the same time, obedience also severely limits the distortions that a monopolist information seller is able to impose on the allocation—the revenue-optimal mechanism provides the correct information to the players more often than the monopolist would like.

In the present work, we characterized optimal mechanisms in the context of a linear model with binary states and actions. In that respect, considerable work remains to be done to extend the applicability of this framework. Natural next steps include removing dominant strategies, introducing strategic complements and substitutes, more than two players, actions, and states. We pursue these avenues in the future work.

In the following, an overview of the ongoing work on strategic complements and substitutes is given.

5.1 Substitute and Complement Competitions

In this section, everything is similar to Example 3 (the main example) except the payoff matrix of the down stream market. The dominant strategy (the right action) for each data buyer is matching his action with the state of the world, and he prefers the other data buyer to take the wrong action. However, the externality is not constant and changes based on the action of the other data buyer. If the rival takes the right action, data buyers payoff increases by α when he switches from wrong action to the right action. If the rival takes the wrong action, the data buyer's payoff increases by β when he takes right action instead of the wrong one. If $\alpha > \beta$, we have a complement competition, and it is substitute if $\alpha < \beta$. The payoff matrix for the down stream market is given below.

$\theta = 0$

	0	1
0	α, α	$1, 0$
1	$0, 1$	$1-\beta, 1-\beta$

$\theta = 1$

	0	1
0	$1-\beta, 1-\beta$	$0, 1$
1	$1, 0$	α, α

In the payoff matrix above, we assume $0 \leq \alpha\,, \beta \leq 1$. Therefore, a data buyer's dominant strategy is taking the action which matches the state of the world. He also prefers the other data buyer to take the wrong action. We also assume $\alpha \geq 1 - \beta$ so it is better for both of them to take the right action rather than both taking the wrong action.

Remark. *When $\alpha = \beta$, the new payoff matrix is a shifted and scaled version of the payoff matrix in Example 3, and we have*

$$\pi_i(a;\theta) = \alpha\Big(\mathbf{1}\{a_i = \theta\} - \frac{1-\alpha}{\alpha}\mathbf{1}\{a_j = \theta\}\Big) + 1 - \alpha\,.$$

Comparing it to (3.2), we find out the results will be the same as the results in the previous chapters but we only need to substitute α by $\frac{1-\alpha}{\alpha}$. In other words,

$$\alpha_{old} = \frac{1-\alpha_{new}}{\alpha_{new}}\,.$$

Moreover, we have $0 \leq \alpha_{old} \leq 1$ *iff* $\frac{1}{2} \leq \alpha_{new} \leq 1$.

In the following, simulation results for revenue-optimal communication rule, σ for different values of α and β from Matlab are given. Notice that payments still can be computed based on the expected market share, $\tilde{\pi}_i$, using Proposition 2, and $\tilde{\pi}_i$ can be computed when we have the σ.

Valuations are distributed exponentially over the finite interval $[2, 5]$. In the blue regions, both data buyers receive the right action, (r, r). In the black region, both receive wrong action, (w, w). In the red region, the first data buyer receives the right action and the second one receives the wrong action, (r, w), and it is the opposite in the green region, (w, r).

As we can see the structure of the solution is different for $\alpha > \beta$, $\alpha = \beta$, and $\alpha < \beta$.

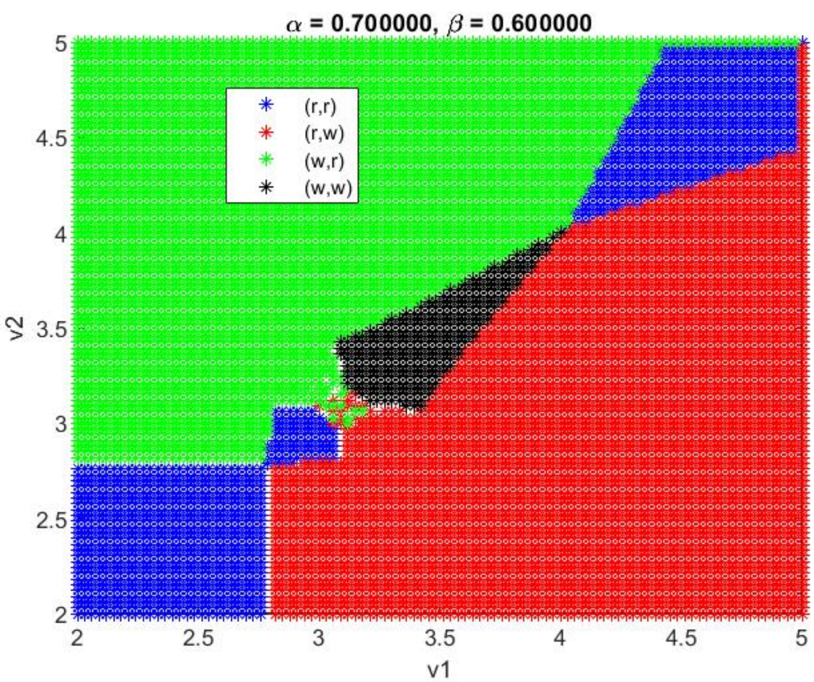

Figure 5-1: $\alpha = 0.7$, $\beta = 0.6$

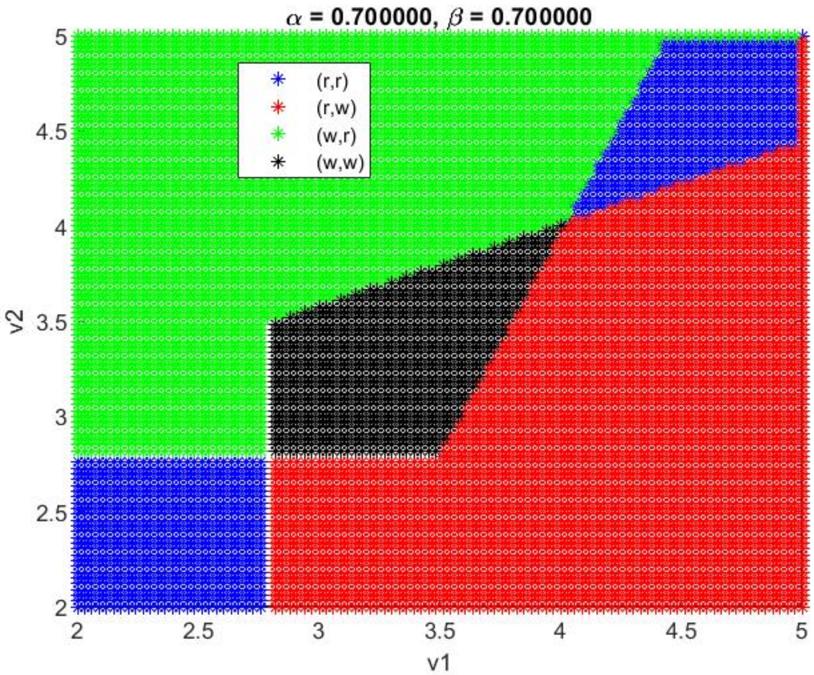

Figure 5-2: $\alpha = 0.7$, $\beta = 0.7$

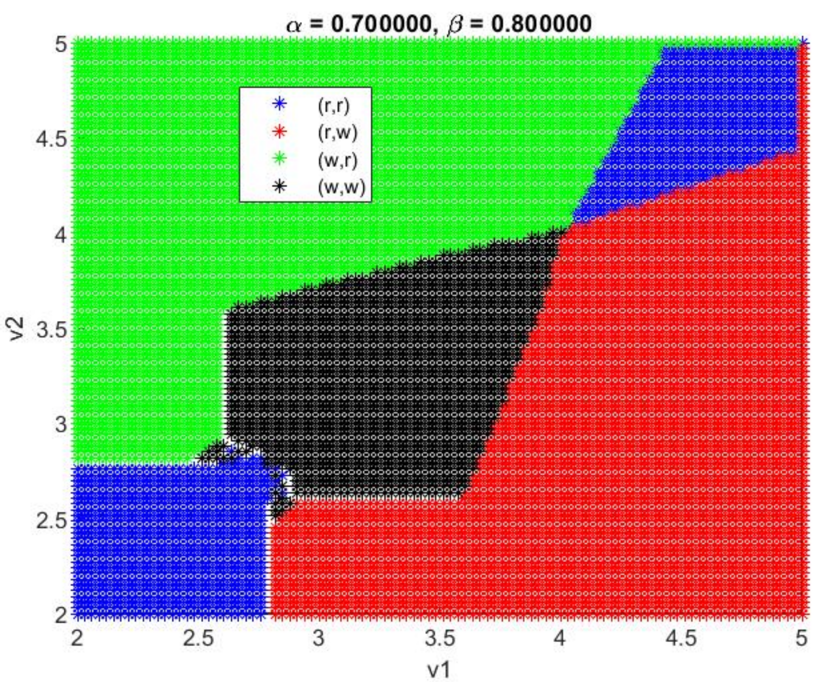

Figure 5-3: $\alpha = 0.7$, $\beta = 0.8$

Appendix A

Missing proofs

Proposition 9 (Proposition 2 restated)**.** *The mechanism (σ, p) is truthful if and only if for each player i:*

1. *The interim share $\tilde{\pi}_i$ is non-decreasing.*
2. *The interim payment $\tilde{p}_i$ is given for $v_i \in \mathcal{V}_i$ by*

$$\tilde{p}_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \underline{v} \cdot \tilde{\pi}_i(\underline{v}) + \tilde{p}_i(\underline{v}) - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds \,. \tag{A.1}$$

Proof. Define $\tilde{u}_i : v_i \mapsto v_i \cdot \tilde{\pi}_i(v_i) - \tilde{p}_i(v_i)$. Then truthfulness is equivalent to

$$\tilde{u}_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \tilde{p}_i(v_i) \geq v_i \cdot \tilde{\pi}_i(v_i') - \tilde{p}_i(v_i') = \tilde{u}_i(v_i') + (v_i - v_i') \cdot \tilde{\pi}_i(v_i') \,,$$

for all $(v_i, v_i') \in \mathcal{V}_i^2$. This is equivalent to saying that $\tilde{\pi}_i(v_i) \in \partial \tilde{u}_i(v_i)$ for all $v_i \in \mathcal{V}_i$ where $\partial \tilde{u}_i(v_i) \subset \mathbb{R}$ denotes the subdifferential of $\tilde{u}_i$ at v_i. By a well-known characterization of convexity, this in turn equivalent to saying that $\tilde{\pi}_i$ is non-decreasing and that

$$\tilde{u}_i(v_i) = \tilde{u}_i(\underline{v}) + \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds \,.$$

This concludes the proof since this last expression is equivalent to (3.1). □

Lemma 4 (Lemma 3 restated)**.** *Assuming truthfulness and individual rationality, optimal allocation for maximizing virtual social surplus is also the solution to the seller's revenue* *maximization problem. Moreover,*

$$\tilde{p}_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - K \cdot \underline{v} - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s)ds\,, \tag{A.2}$$

where $K \coloneqq \min_{\sigma_j \in \Delta(\mathcal{A}_j)} \max_{a_i \in \mathcal{A}_i} \mathbb{E}\big[\pi_i(a_i, A_j; \theta)\big]$ *is the expected market share in a player's outside option (recall we assume that the game is symmetric, hence* K *does not depend on* i*).*

Proof. Note that we can write the seller's revenue as

$$R = \mathbb{E}\Big[\sum_i p_i(\theta, V)\Big] = \mathbb{E}\Big[\sum_i \mathbb{E}[p_i(\theta, V) \mid V_i]\Big] = \mathbb{E}\Big[\sum_i \tilde{p}_i(V_i)\Big].$$

By Proposition 2, we have

$$\begin{aligned} R &= \mathbb{E}\Big[\sum_i V_i \cdot \tilde{\pi}_i(V_i) - \underline{v} \cdot \tilde{\pi}_i(\underline{v}) + \tilde{p}_i(\underline{v}) - \int_{\underline{v}}^{V_i} \tilde{\pi}_i(s)ds\Big] \\ &= \sum_i \mathbb{E}\Big[V_i \cdot \tilde{\pi}_i(V_i) - \int_{\underline{v}}^{V_i} \tilde{\pi}_i(s)ds\Big] + \sum_i \big(- \underline{v} \cdot \tilde{\pi}_i(\underline{v}) + \tilde{p}_i(\underline{v})\big). \end{aligned}$$

We are now going to treat each summand separately. First, we will show that the first one is equal to the virtual surplus and then we will compute $\tilde{p}_i(\underline{v})$ in the second one for the revenue-maximizing mechanism.

We compute the expectation in the first summand as follows:

$$\begin{aligned} \mathbb{E}\Big[V_i \cdot \tilde{\pi}_i(V_i) - \int_{\underline{v}}^{V_i} \tilde{\pi}_i(s)ds\Big] &= \int_{\underline{v}}^{\overline{v}} v_i \cdot \tilde{\pi}_i(v_i) f(v_i) dv_i - \int_{\underline{v}}^{\overline{v}} \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) f(v_i) ds dv_i \\ &= \int_{\underline{v}}^{\overline{v}} v_i \cdot \tilde{\pi}_i(v_i) f(v_i) dv_i - \int_{\underline{v}}^{\overline{v}} \int_{s}^{\overline{v}} \tilde{\pi}_i(s) f(v_i) dv_i ds \\ &= \int_{\underline{v}}^{\overline{v}} \left(v_i - \frac{1 - F(v_i)}{f(v_i)}\right) \cdot \tilde{\pi}_i(v_i) f(v_i) dv_i \\ &= \mathbb{E}\big[\phi(V_i)\tilde{\pi}_i(V_i)\big]\,, \end{aligned}$$

where the second equality uses Fubini's theorem and the last equality is the definition of the virtual value function. Summing the last expression over $i \in [n]$ yields the virtual surplus.

It then only remains to determine the value of $\tilde{p}_i(\underline{v})$ which maximizes revenue. For this, note that individual rationality is the only constraint on $\tilde{p}_i(\underline{v})$ since the payments do not appear in the incentive obedience constraint, and since truthfulness does not constrain the payment at the lowest type by Proposition 2.

We express individual rationality (Definition 4) using the interim market share $\tilde{\pi}_i$:

$$\forall v_i \in \mathcal{V}_i \quad v_i \tilde{\pi}_i(v_i) - \tilde{p}_i(v_i) \geq v_i \cdot K\,.$$

Substituting $\tilde{p}_i$ in the previous equation using (3.1), we have

$$\forall v_i \in \mathcal{V}_i \quad \underline{v}(\tilde{\pi}_i(\underline{v}) - K) + \int_{\underline{v}}^{v_i} (\tilde{\pi}_i(s) - K) ds \geq \tilde{p}_i(\underline{v}).$$

Thus, we can find the maximum value for $\tilde{p}_i(\underline{v})$ by minimizing the left-hand side over v_i:

$$\max \tilde{p}_i(\underline{v}) = \min_{v_i} \left\{ \underline{v}(\tilde{\pi}_i(\underline{v}) - K) + \int_{\underline{v}}^{v_i} (\tilde{\pi}_i(s) - K) ds \right\} = \underline{v}(\tilde{\pi}_i(\underline{v}) - K).$$

The second equality holds because $\tilde{\pi}_i \geq K$ since

$$\begin{aligned}
\tilde{\pi}_i(V_i) &= \mathbb{E}\big[\tilde{\pi}_i(A_i, A_j; \theta) \,|\, V_i\big] \\
&\geq \mathbb{E}\big[\tilde{\pi}_i(a^o, A_j; \theta) \,|\, V_i\big] \\
&\geq \mathbb{E}\big[\tilde{\pi}_i(a^o, \theta; \theta) \,|\, V_i\big] = K,
\end{aligned}$$

where $a^o \coloneqq \arg\max_{a_i \in \mathcal{A}_i} \mathbb{E}\big[\pi_i(a_i, \theta; \theta)\big]$. The first inequality holds by choosing the deviation function δ in Definition 2 to be the constant function a^o, the second inequality holds by Assumption 2, and the last equality is the definition of K. □

Bibliography

[1] Anish Agarwal, Munther Dahleh, Thibaut Horel, and Maryann Rui. Towards data auctions with externalities. *arXiv preprint arXiv:2003.08345*, 2020.

[2] Dirk Bergemann and Alessandro Bonatti. Markets for information: An introduction. *Annual Review of Economics*, 11:85–107, 2019.

[3] Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American Economic Review*, 108(1):1–48, January 2018.

[4] Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.

[5] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

[6] Kostas Bimpikis, Davide Crapis, and Alireza Tahbaz-Salehi. Information sale and competition. *Management Science*, 65(6):2646–2664, 2019.

[7] David Bounie, Antoine Dubus, and Patrick Waelbroeck. Selling strategic information in digital competitive markets. *RAND Journal of Economics*, 52(2):283–313, 2021.

[8] Jakub Kastl, Marco Pagnozzi, and Salvatore Piccolo. Selling information to competitive firms. *RAND Journal of Economics*, 49(1):254–282, 2018.

[9] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[10] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, 1991.

[11] Ina Taneva. Information design. *American Economic Journal: Microeconomics*, 11(4):151–85, 2019.

[12] Yi Xiang and Miklos Sarvary. Buying and selling information under competition. *Quantitative Marketing and Economics*, 11(3):321–351, 2013.