Generic Bayesian Implementability and Flows*

Zhonghong Kuang[†] Yi Liu[‡] Fan Wu[§]

September 1, 2023

Abstract

We study when an allocation rule is implementable in Bayesian mechanisms. We provide a necessary and sufficient condition on this implementability problem with no restriction on the allocation rule, the joint type distribution, or the belief. Our proof is based on the fundamental duality theorem and network flow theory. Generically, for almost all joint type distribution, any allocation is implementable in Bayesian mechanisms.

JEL Classification: C61, D82

Keywords: Mechanism Design, Bayesian Mechanism, Bayesian Implementation

*For helpful comments and insights, we thank Fedor Sandomirskiy, Omer Tamuz, Luciano Pomatto, Peter Caradonna. All errors are our own.

[†]Zhonghong Kuang: School of Economics, Renmin University of China (kuang@ruc.edu.cn).

[‡]Yi Liu: Institute for Interdisciplinary Information Sciences, Tsinghua University (liu-yi22@mails.tsinghua.edu.cn).

[§]Fan Wu: Division of the Humanities and Social Sciences, California Institute of Technology (fwu@caltech.edu).

1 Introduction

One of the most fundamental problems in mechanism design is whether an allocation rule is implementable. This question appears in various scenarios, including public goods provision, bilateral trade, matching, etc. In the public goods provision problem, a governor needs to decide whether to provide a public good. Its value for each person is unknown to the governor. Given any specific allocation rule, is there a mechanism where everyone truthfully reports his value?

Rochet (1987) characterizes when an allocation rule is implementable in dominant-strategy mechanisms. The implementability condition reduces to the valuation difference satisfying cyclic monotonicity for all players and all report profiles (of the other players). This condition is demanding as the dominant-strategy mechanism is a strong concept. Consequently, many allocation rules are not implementable, which significantly limits the scope of applications. What if we relax the dominant strategy to Bayesian mechanisms? What is the implementability condition for Bayesian mechanisms?

We provide a sufficient and necessary condition for an allocation rule to be implementable in Bayesian mechanisms. For each player, we say that two types are *consistent* if their posteriors about the other players' type are identical. This consistency is an equivalence relation, and we can partition the type space into consistent classes. We show that an allocation rule is implementable in Bayesian mechanisms if and only if the expected valuation difference satisfies cyclic monotonicity within each consistent class for all players (Theorem 1).

We call a belief profile *non-degenerate correlated* if each player's posterior belief (about the other players' type) varies with his own type (an injection). This condition is weak as it is generically true in the space of all belief profiles/type distributions. As a corollary of our characterization, our second result shows that generically, any allocation rule is implementable in Bayesian mechanisms (Theorem 2).

Our results are quite general in several aspects. First, we do not assume independence on the prior distribution. That is, we allow for arbitrary correlation in the players' joint type distribution. Second, we do not need to assume a common prior. We only need a first-order belief, and belief hierarchy is irrelevant. The common prior assumption is widely adopted in the literature. For example, Crémer and McLean (1988)'s full rent extraction relies heavily on the common prior assumption. Third, we do not impose any restrictions on the allocation rule. Fourth, we allow for interdependent preference. The only substantive assumption we need is that the preference is quasi-linear in transfer.

Our technical contribution is a novel proof strategy to derive the implementability condition. Our proof is based on the fundamental duality theorem and network flow theory. Vohra (2011) provides an extensive summary of how mechanism design intimately relates to linear programming. Our proof might be helpful for future research as we directly analyze the polyhedron of incentive-compatible Bayesian mechanisms.

Bergemann et al. (2012) notice a result similar to our Theorem 1 in the setting of efficient auction. Müller et al. (2007) characterize the implementability condition of Bayesian mechanisms assuming independent common prior. Gershkov et al. (2013) show that the set of implementable interim utilities under some conditions is the same under dominant-strategy and Bayesian mechanisms, although the set of implementable allocation rules differs.

The rest of this paper is organized as follows. In Section 2, we set up the model. Section 3 characterizes the implementability condition for Bayesian mechanisms. Section 4 discusses the implications of our main result. In Section 5, we apply our results to several examples in the literature.

2 Setting

There are $n \geq 1$ agents (he) and a principal (she). Each agent i has a type θ_i drawn from a common type space $\Theta = \{\theta^1, \cdots, \theta^m\}$.[1] The state of the game is the *type profile* $\boldsymbol{\theta} = (\theta_1, \cdots, \theta_n) \in \Theta^n$. Let $\mathbb{P}$ denote the common prior. Let $\mathbb{P}_{-i}(\cdot|\theta_i)$ denote agent i's belief conditional on his type θ_i. Let $(\mathbb{P}_{-1}, \cdots, \mathbb{P}_{-n})$ denote a

[1]The type space does not need to be identical across agents.

belief profile. (Our result does not need the common prior assumption. We only need a belief profile.) A mechanism consists of an allocation rule and a transfer rule. The implemenation problem is: when is an allocation rule implementable? By the revelation principle, we focus on direct mechanisms.

Let θ'_i denote the report type of agent i, which is not necessarily θ_i, and $\boldsymbol{\theta'} = (\theta'_1, \cdots, \theta'_n)$ denote the *report profile*. An *allocation rule* g is a function that assigns an outcome to each report profile, i.e., $g : \Theta^n \to X$ where X denotes a set of outcomes. We allow for interdependent values across agents. That is, agents' values do not only depend on their own types but on the types of all agents. The value that agent i assigns to an allocation x is denoted $v_i(x|\boldsymbol{\theta})$. Let $\boldsymbol{v} = (v_1(\cdot), \cdots, v_n(\cdot))$ denote the value profile.

A transfer is a function $\boldsymbol{t} : \Theta^n \to \mathbb{R}^n$. Given a report profile $\boldsymbol{\theta'}$, agent i receives $t_i(\boldsymbol{\theta'})$, where $t_i(\cdot)$ is the *transfer* function to agent i. Let $\boldsymbol{t} = (t_1(\cdot), \cdots, t_n(\cdot))$ denote the transfer profile. When the actual type is θ_i and the report profile is $\boldsymbol{\theta'}$, the payoff of agent i is $u_i(\boldsymbol{\theta'}, \boldsymbol{\theta}) = v_i(g(\boldsymbol{\theta'})|\boldsymbol{\theta}) + t_i(\boldsymbol{\theta'})$.

Note that our setting can nest various models in mechanism design. We present several examples.

Model 1 (Public goods)**.** Clarke (1971) and Groves (1973) study the public good provision problem. The governor decides whether to provide a public good to benefit n agents. Agent i's valuation of the public good is θ_i, which is unknown to the governor. The governor makes her decision based on the report profile $\boldsymbol{\theta'}$. The allocation rule g maps a report profile to an outcome $x \in X = \{0, 1\}$, where 1 indicates approval and 0 indicates disapproval. The benevolent principal wants to provide the public good if the sum of the value of all agents exceeds the production cost C, namely $x = 1$ if and only if $\sum_{i=1}^{n} \theta_i > C$. The principal designs the payment ϕ_i for each agent i to induce truthful reporting. If $g(\boldsymbol{\theta'}) = 0$, all agents pay zero. Otherwise, agents share the production cost. Agents i's payoff is $\mathbf{1}_{x=1}\theta_i - \phi_i(\boldsymbol{\theta'})$.[2]

Model 2 (Bilateral Trade)**.** Myerson and Satterthwaite (1983) introduce the bi-

[2]Rob (1989) and Pesendorfer (1998) study the pollution claim settlements, which is similar to the public good provision problem.

lateral trade model. A seller (agent 1) owns a single indivisible good. There is one potential buyer (agent 2). Agent i's private valuation of the good is θ_i. The mechanism designer tries to induce efficient trade. The trade should occur if and only if $\theta_1 < \theta_2$. The report type is (θ'_1, θ'_2). The allocation rule g maps a report profile to an outcome $x \in X = \{0, 1\}$, where 1 indicates trade and 0 indicates no trade. If the trade takes place, the payment $\phi(\theta'_1, \theta'_2)$ is the price of the good. If the trade does not occur, $\phi(\theta'_1, \theta'_2) = 0$. The seller's payoff is $\theta_1(1 - g(\theta'_1, \theta'_2)) + \phi(\theta'_1, \theta'_2)$ and the buyer's payoff is $\theta_2 g(\theta'_1, \theta'_2) - \phi(\theta'_1, \theta'_2)$.

Model 3 (Assigning Multiple Objects)**.** Moulin (2009) introduces the following objects assignment problem. There are p identical objects and n agents with $1 \leq p < n$. Each agent demand at most one object. Agent i's valuation of an object is θ_i. The principal asks for agents' valuations. The allocation rule g decides which p agents would receive the object given a report profile $\boldsymbol{\theta}'$. The principal designs the payment $\phi_i(\boldsymbol{\theta}')$ for each agent i. Agent i's payoff is $\theta_i \mathbf{1}_{i \text{ receives the good}} - \phi_i(\boldsymbol{\theta}')$.

We now define incentive-compatible mechanisms.

Definition 1 (**DICM**)**.** *A mechanism $(g, \boldsymbol{t})$ is a dominant-strategy incentive-compatible mechanism if truthful reporting by all agents forms a dominant-strategy equilibrium, i.e.,*

$$v_i(g(\theta_i, \boldsymbol{\theta}'_{-i})|\boldsymbol{\theta}) + t_i(\theta_i, \boldsymbol{\theta}'_{-i}) \geq v_i(g(\theta'_i, \boldsymbol{\theta}'_{-i})|\boldsymbol{\theta}) + t_i(\boldsymbol{\theta}'), \forall \boldsymbol{\theta}_{-i} \in \Theta^{n-1}, \boldsymbol{\theta}' \in \Theta^n, i = 1, \cdots, n.$$

If we assume that each agent's value of outcome only depends on his own type, we recover the standard definition.

$$v_i(g(\theta_i, \boldsymbol{\theta}'_{-i})|\theta_i) + t_i(\theta_i, \boldsymbol{\theta}'_{-i}) \geq v_i(g(\theta'_i, \boldsymbol{\theta}'_{-i})|\theta_i) + t_i(\theta'_i, \boldsymbol{\theta}'_{-i}), \forall \boldsymbol{\theta}' \in \Theta^n, i = 1, \cdots, n. \tag{DIC}$$

Definition 2 (**BICM**)**.** *A mechanism $(g, \boldsymbol{t})$ is a Bayesian incentive-compatible* *mechanism if truthful reporting by all agents forms a Bayesian equilibrium, i.e.,*

$$\mathbb{E}_{\boldsymbol{\theta}_{-i}}\left[v_i(g(\theta_i,\boldsymbol{\theta}_{-i})|\boldsymbol{\theta})+t_i(\theta_i,\boldsymbol{\theta}_{-i})|\theta_i\right]\geq\mathbb{E}_{\boldsymbol{\theta}_{-i}}\left[v_i(g(\theta'_i,\boldsymbol{\theta}_{-i})|\boldsymbol{\theta})+t_i(\theta'_i,\boldsymbol{\theta}_{-i})|\theta_i\right],\forall i=1,\cdots,n. \tag{BIC}$$

We study when an allocation rule is implementable in the Bayesian incentive-compatible mechanism, also known as the implementability or implementation problem (Rochet, 1987; Müller et al., 2007). We do not need to impose any individual rationality (IR) constraint. For any reservation utility, we can always translate the transfer up to meet the IR constraint, without affecting the IC constraint. In the same logic, without loss of generality, we assume that $\boldsymbol{t}\geq\boldsymbol{0}$.

Definition 3 (Implementability of Mechanisms). *We say that an allocation g is implementable in dominant-strategy/Bayesian mechanisms if there exists a $\boldsymbol{t}\geq\boldsymbol{0}$ such that $(g,\boldsymbol{t})$ satisfy the (DIC)/(BIC) condition.*

The most classic example is the efficient implementation problem. We say an allocation rule is *efficient* if it maximizes the sum of agents' payoff. A principal with the efficient allocation as her objective is called *benevolent*.

3 Implementatability Conditions

For the following analysis, it is useful to define the *value difference* between θ_i and θ'_i as

$$Dv_i(\theta'_i,\theta_i|\boldsymbol{\theta_{-i}},\boldsymbol{\theta'_{-i}})\triangleq v_i(g(\theta_i,\boldsymbol{\theta'_{-i}})|\boldsymbol{\theta})-v_i(g(\theta'_i,\boldsymbol{\theta'_{-i}})|\boldsymbol{\theta}).$$

Moreover, we define the *expected value difference* between θ_i and θ'_i

$$EDv_i(\theta'_i,\theta_i)\triangleq\mathbb{E}_{\boldsymbol{\theta}_{-i}}\left[v_i(g(\theta_i,\boldsymbol{\theta_{-i}})|\boldsymbol{\theta})-v_i(g(\theta'_i,\boldsymbol{\theta_{-i}})|\boldsymbol{\theta})|\theta_i\right].$$

We review the results on the implementability of the dominant-strategy mechanisms. Rochet (1987) obtains this result abstracting away interdependent values.

For the lemma below, we assume that each agent's value of outcome only depends on his own type.

Lemma 1. *(Rochet, 1987). An allocation rule g is implementable in dominant-strategy mechanisms if and only if value difference Dv satisfies cyclic monotonicity, i.e., for all i, for any finite types $\theta^{(1)}, \cdots, \theta^{(k)} \in \Theta$, for all $\boldsymbol{\theta'_{-i}}$,*

$$Dv_i(\theta^{(1)}, \theta^{(2)}|\boldsymbol{\theta'_{-i}}) + Dv_i(\theta^{(2)}, \theta^{(3)}|\boldsymbol{\theta'_{-i}}) + \cdots + Dv_i(\theta^{(k)}, \theta^{(1)}|\boldsymbol{\theta'_{-i}}) \geq 0. \quad (1)$$

For sufficency, we illustrate Rochet (1987)'s construction of the transfer function. Construct a directed graph: (i) each type is represented by a node; (ii) the length of the arc from node θ to node θ' is $Dv(\theta, \theta')$. For each type θ, we can find the shortest length path among those directed paths starting at θ, denoted by $w(\theta)$, i.e.,

$$w(\theta) \triangleq \min_{\theta^{(n)}, \text{all sequences from } \theta^{(0)}=\theta \text{ to } \theta^{(n)}} \sum_{k=0}^{n-1} Dv(\theta^{(k)}, \theta^{(k+1)}). \quad (2)$$

According to the cyclic monotonicity, all directed cycles in the constructed directed graph have positive length. Thus, w is well-defined.

Then transfer $t(\theta) = w(\theta) - \min_\theta w(\theta)$ is incentive compatible. Note that $w(\theta) + Dv(\theta', \theta)$ represents the length of a path starting from θ', which is necessarily no less than $w(\theta')$—the shortest length path starting from θ'. Hence, $w(\theta) + Dv(\theta', \theta) \geq w(\theta')$, implying $t(\theta) \geq t(\theta') - Dv(\theta', \theta)$.

We next characterize the implementability condition for Bayesian mechanisms. First, we introduce some notations. For an agent i, we say two types θ^j, θ^k are *consistent* if two conditional probability distributions associated with θ^j, θ^k are identical, i.e., $\mathbb{P}_{-i}(\cdot|\theta^j) \equiv \mathbb{P}_{-i}(\cdot|\theta^k)$.[3] Then, for each i, we can partition the type space Θ into several consistent classes. Let $\Theta_i^{(j)}$ denote the j-th consistent class for agent i, i.e., $\Theta = \bigcup_{j=1}^{J_i} \Theta_i^{(j)}$ and $\forall j \neq k, \Theta_i^{(j)} \cap \Theta_i^{(k)} = \varnothing$, where J_i is the number of consistent classes for agent i. We provide an example to illustrate the definition

[3]Namely, $\mathbb{P}_{-i}(\boldsymbol{\theta}_{-i} = \hat{\boldsymbol{\theta}}_{-i}|\theta_i = \theta^j) = \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i} = \hat{\boldsymbol{\theta}}_{-i}|\theta_i = \theta^k)$ for all $\hat{\boldsymbol{\theta}}_{-i} \in \Theta^{n-1}$.

of consistent types.

Example 1. Consider a two-player problem with $\Theta = \{\theta^H, \theta^M, \theta^L\}$. The prior distribution is summarized in the following table. For agent 1, θ^M, θ^L are consistent types, since their conditional probability distributions are identical. Hence, the consistent classes for agent 1 are $\Theta_1^{(1)} = \{\theta^H\}$ and $\Theta_1^{(2)} = \{\theta^M, \theta^L\}$. Meanwhile, there is no consistent type for agent 2. Thus, the consistent classes for agent 2 are $\Theta_2^{(1)} = \{\theta^H\}$, $\Theta_2^{(2)} = \{\theta^M\}$, and $\Theta_2^{(3)} = \{\theta^L\}$. □

$\mathbb{P}(\cdot)$	$\theta_2 = \theta^H$	$\theta_2 = \theta^M$	$\theta_2 = \theta^L$
$\theta_1 = \theta^H$	$\frac{1}{6}$	0	$\frac{1}{9}$
$\theta_1 = \theta^M$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$
$\theta_1 = \theta^L$	$\frac{1}{18}$	$\frac{1}{18}$	$\frac{1}{18}$

Table 1: Prior Distribution

We are ready to provide a sufficient and necessary condition for an allocation rule to be implementable in Bayesian mechanisms. It turns out we only need the expected value difference to satisfy the cyclic monotonicity within each consistent class.

Theorem 1. ***Charaterization***. *An allocation rule is implementable in Bayesian mechanism if and only if for all i, all consistent classes $\Theta_i^{(j)}$, and all finite types $\theta^{(1)}, \cdots, \theta^{(k)} \in \Theta_i^{(j)}$,*

$$EDv_i(\theta^{(1)}, \theta^{(2)}) + EDv_i(\theta^{(2)}, \theta^{(3)}) + \cdots + EDv_i(\theta^{(k)}, \theta^{(1)}) \geq 0. \tag{3}$$

We can compare this implementability condition versus the one under dominant-strategy mechanisms. As agents only care about the expected payoff in a Bayesian mechanism, it is intuitive that the expected value difference supersedes the value difference (conditional on the others' report). Yet, it is surprising that we only need the cyclic monotonicity to hold within each consistent class. We provide a proof sketch below to offer some intuition. The formal proof is relegated to Appendix A.

Proof Sketch of Theorem 1. For necessity, for each agent, for k types, we can write k incentive-compatible inequalities. For all i and $\boldsymbol{\theta}_{-i}$, we have

$$\mathbb{E}_{\boldsymbol{\theta}_{-i}}[t(\theta^{(\kappa)}, \boldsymbol{\theta}_{-i})|\theta^{\kappa}] \geq \mathbb{E}_{\boldsymbol{\theta}_{-i}}[t(\theta^{(\kappa-1)}, \boldsymbol{\theta}_{-i})|\theta^{\kappa}] - EDv(\theta^{(\kappa-1)}|\theta^{(\kappa)}), \forall \kappa \in \{1, 2, \cdots, k\},$$

with $\theta^{(0)} \triangleq \theta^{(k)}$. Telescoping sum yields the condition.

For sufficiency, we use primal-dual technique based on the fundamental duality theorem. We translate the existence problem into a system of inequalities. By Farkas' lemma, the following two statements are equivalent: (i) There exists a vector $\boldsymbol{T}_i$ such that $\boldsymbol{T}_i \geq \mathbf{0}, A_i\boldsymbol{T}_i \geq \boldsymbol{b}_i$. (ii) For any vector $\boldsymbol{y}_i$ such that $\boldsymbol{y}_i \geq \mathbf{0}, \boldsymbol{y}_i^{\top} A_i \leq \mathbf{0}$, we have $\boldsymbol{b}_i^{\top}\boldsymbol{y}_i \leq 0$. From the linear programming perspective, $\boldsymbol{y}_i$ is the Lagrange multiplier of $A_i\boldsymbol{T}_i \geq \boldsymbol{b}_i$.

Back to our problem, we first rewrite the BIC condition, a polyhedron, as the form of statement (i) above. Let $\boldsymbol{P}_i(\theta^r)$ denote the posterior of agent i conditional on his type being θ^r. It is a m^{n-1}-dimensional row vector of $\mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta^r)$ with m^{n-1} different $\boldsymbol{\theta}_{-i}$ arranged in lexicographic order. Each row of matrix A_i consists of the coefficient of BIC condition like $(\mathbf{0}, \boldsymbol{P}_i(\theta^r), \mathbf{0}, -\boldsymbol{P}_i(\theta^r), \mathbf{0})$. And let $\boldsymbol{T}_i$ denote the corresponding vector of transfer and $\boldsymbol{b}_i$ denote the vector that consists of the corresponding minus expected value difference, i.e. $-EDv_i(\theta^s, \theta^r)$. Then the BIC condition can be written as $\{\boldsymbol{T}_i | A_i\boldsymbol{T}_i \geq \boldsymbol{b}_i, \boldsymbol{T}_i \geq \mathbf{0}\}$. Now it suffices to show that the polyhedron $\{\boldsymbol{T}_i | A_i\boldsymbol{T}_i \geq \boldsymbol{b}_i, \boldsymbol{T}_i \geq \mathbf{0}\}$ is non-empty.

By our previous discussion, the polyhedron being non-empty is equivalent to the existence of the vectors y_i. Each vector y_i has $m(m-1)$ entries, corresponding to pairs of different types. The entries are indexed by $r \circ s$ in a lexicographic order: $1\circ 2, 1\circ 3, \cdots, 1\circ m, \cdots, m\circ 1, m\circ 2, \cdots, m\circ(m-1)$. By the fundamental duality theorem, the $r \circ s$ entry of y_i, $y_i^{r\circ s} \geq 0$ is the Lagrange multiplier of

$$\sum_{\boldsymbol{\theta}_{-i}} \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta^r) t_i(\theta^r, \boldsymbol{\theta}_{-i}) \geq \sum_{\boldsymbol{\theta}_{-i}} \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta^r) t_i(\theta^s, \boldsymbol{\theta}_{-i}) - EDv_i(\theta^s, \theta^r)$$

With some algebraic manipulation, we can get $\boldsymbol{y}_i^{\top} A_i = \mathbf{0}$. Expand it, we have for any r,

$$\sum_{s,s\neq r} y_i^{ros}\boldsymbol{P}_i(\theta^r) - \sum_{s,s\neq r} y_i^{sor}\boldsymbol{P}_i(\theta^s) = \mathbf{0}.$$

Since all posterior $\boldsymbol{P}_i$ has L1-norm 1, we get $\sum_{s,s\neq r} y_i^{ros} - \sum_{s,s\neq r} y_i^{sor} = 0$. Then the above equation has a strong geometric interpretation that any posterior $\boldsymbol{P}_i(\theta^r)$ is in the convex hull of other posterior vectors. In Figure 1, consider the vector that has the largest L2-norm, i.e. the red vector. It cannot be in the convex hull of other different vectors. So for any two inconsistent types θ^r, θ^s, we have $y^{ros} = y^{sor} = 0$.

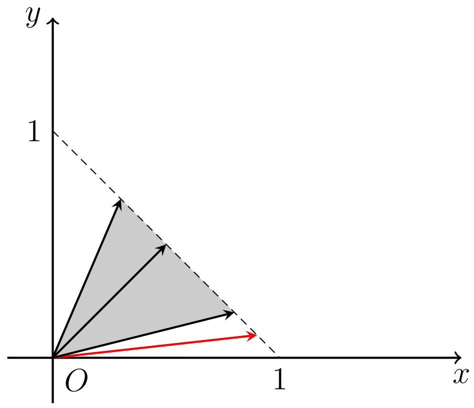

Figure 1: Geometric intuition

Then the Lagrange multiplier of any two inconsistent types θ^r, θ^s is always zero. By complementary slackness, these incentive-compatible inequalities do not affect the existence of the transfer payment. So we only need to consider the inequality of those consistent types, which reduces to Lemma 1.

□

4 Further Discussions

In the literature, there is a common assumption that types are independently distributed across players. Yet, this assumption is quite strong. In reality, types are typically correlated. For example, in a procurement process (Anton and Yao, 1992; Jarman and Meisner, 2017), a principal decides to purchase a good from multiple potential suppliers. (The principal is benevolent and wants to purchase a good from the lowest-cost supplier.) Each supplier's cost is his private information. The costs are commonly interdependent, as there are industrial-level shocks that induce a positive correlation, such as changes in the price of raw materials or exchange rates. When one supplier has a higher cost, his belief about others' costs is generally higher.

Motivated by such considerations, we introduce a concept that only involves minimal correlation. We say that a belief profile is *non-degenerate correlated* if for any agent i, his belief about the others' type distribution conditional on his own type θ_i varies with θ_i, i.e., for all $\theta^r \neq \theta^s$ and $i = 1, \cdots, n$, $\exists \boldsymbol{\theta}_{-i} \in \Theta^{n-1}, \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta^r) \neq \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta^s)$.[4] We also say that a common prior is *non-degenerate correlated* if such condition holds. Generically, almost all joint type distribution is non-degenerate correlated.[5] Moreover, our condition is weaker than other related conditions in the literature.

One similar assumption is the Crémer–McLean conditions (Crémer and McLean, 1988). They derive conditions under which full rent extraction is feasible in dominant-strategy auctions and Bayesian auctions. We can write agent i's belief about the other agents' types conditional on agent i's type θ_i as a vector with m^{n-1} entries. Agent i's conditional belief is described by the set of vectors of this form, one for each type of agent i. The first Crémer–McLean condition requires that these vectors are linearly independent for every agent. The second requires that for each agent, any vector cannot be a convex combination of the others. By contrast, a non-degenerate correlated belief profile only requires that these vectors are pairwise distinct for every agent. Hence, both Crémer-McLean conditions are more restrictive.[6]

Another related concept is the stochastic relevance condition in Miller et al. (2005). We say that type θ_i is *stochastically relevant* for type θ_j if the distribution

[4]The non-degenerate correlated condition is the same as condition (A4) in Kandori and Matsushima (1998).

[5]The set of joint distributions that fail to meet the non-degenerate correlated condition has Lebesgue measure zero.

[6]When $n \geq 3$, every joint distribution satisfying the Crémer–McLean condition is non-degenerate correlated, but not vice versa. When $n = 2$, the non-degenerate correlated condition is equivalent to the Crémer–McLean condition.

of θ_j conditional on θ_i varies with θ_i. We say that the *stochastic relevance condition* holds if for all $i \neq j$, θ_i is stochastically relevant for θ_j. Every joint distribution satisfying the stochastic relevance condition is non-degenerate correlated, but not vice versa.[7]

Given a non-degenerate correlated distribution, each consistent class $\Theta_i^{(m)}$ is a singleton, leading to the following result.

Theorem 2. *If the belief profile is non-degenerate correlated, any allocation is implementable in a Bayesian mechanism.*

This result is very powerful. Recall that non-degenerate correlated condition is weak and holds generically in the set of all joint type distributions. Hence, generically, any allocation rule is implementable.

Yet, our result is silent on the magnitude of the transfer. In our problem, the principal's only objective is to implement an allocation rule. She has an infinite amount of funds to achieve this objective. We do not impose any restraint on the transfer like budget balance. Consequently, we get a stronger implementability result in return.

5 Applications

In this section, we apply our results to several examples in the literature.

5.1 Communication and Collusion

Kandori and Matsushima (1998) analyze whether communication can sustain cooperation in a market with secret price cutting. In the market, each firm's sale level is its private information. Firms cannot directly observe price cutting from competing firms. Other than private price cutting, demand shocks can also affect sales levels. Hence, each firm's own sales can only imperfectly reflect whether

[7]When $n \geq 3$, the stochastic relevance condition is more restrictive than the non-degenerate correlated condition. When $n = 2$, the non-degenerate correlated condition is equivalent to the stochastic relevance condition.

opponents cut their prices. Kandori and Matsushima (1998) study whether agents would tell the truth when communicating with other agents to form long-term cooperation.

To maintain high prices, firms need to punish those deviators. However, since products are generally differentiated, the firms typically receive different levels of sales and may end up having diverse beliefs about what might have happened. At the end of each period, players can communicate with each other about what they have privately observed. Since communication is cheap, firms can freely provide any false information if it suits their interest.

Kandori and Matsushima (1998) model the above problem as an infinite repeated game with discount factor $\delta \to 1$. In each period t, each firm i chooses an action (i.e., price) $a_i^t \in A_i$. Then a signal (i.e., sales profile) $\boldsymbol{w}^t$ randomly realizes according to the action profile $\boldsymbol{a}^t$ from the distribution $p(\boldsymbol{w}|\boldsymbol{a})$. The likelihood function $p(\boldsymbol{w}|\boldsymbol{a})$ is uniform across different stages. Each firm i only observes its own signal (i.e., sales) $w_i^t \in W_i$. They assume the non-degenerate correlated condition, i.e. for all agent i, for all a, for all $w_i \neq w_i'$, the posterior distribution $\tilde{p}_{-i}(w_{-i}|a, w_i) \neq \tilde{p}_{-i}(w_{-i}|a, w_i')$ for some w_{-i}.

At the end of each period, each firm communicates with other firms by sending a message $m_i^t \in M_i$ (reporting its sales). In each period, firm i's strategy $s_i^t = (\alpha_i^t, \eta_i^t)$ consists of a strategy of action $\alpha_i^t : A_i^{t-1} \times W_i^{t-1} \times M^{t-1} \to \Delta(A_i)$ and a strategy to send message $\eta_i^t : A_i^t \times W_i^t \times M^{t-1} \to \Delta(M_i)$. In period t, firm i obtains a stage payoff $\phi_i(\alpha_i^t, w_i^t)$.

Kandori and Matsushima (1998) reduce the infinite repeated game to a T-period repeated game (as $T \to \infty$) with a transfer $\boldsymbol{t}$ paid in the last stage. Namely, the stage game repeats itself for T times, and afterward, each firm i receives $t_i(\boldsymbol{m}^T)$ where $\boldsymbol{m}^T = (m_1^T, \cdots, m_n^T)$. The payoff of firm i is thus

$$\frac{1-\delta}{1-\delta^T} \sum_{t=1}^{T} \phi_i(a_i^t, w_i^t)\delta^{t-1} + t_i(\boldsymbol{m}^T).$$

They restrict attention to stationary action, where each firm adopts the same strategy in every period, $\alpha_i^t = \alpha_i$ for all t. The payoff of firm i is translated into

$$\mathbb{E}_{\boldsymbol{w}\sim p(\cdot|\boldsymbol{\alpha})}\phi_i(\alpha_i, w_i) + \mathbb{E}\{t_i(\boldsymbol{m})|\boldsymbol{\alpha}, \boldsymbol{w}\},$$

where $\boldsymbol{\alpha} = (\alpha_1, \cdots, \alpha_n)$ denotes the profile of stationary strategies. The former term, $\mathbb{E}_{\boldsymbol{w}\sim p(\cdot|\boldsymbol{\alpha})}u_i(\alpha_i, w_i)$, is solely determined by the stationary strategy profile $\boldsymbol{\alpha}$. They argue that we can treat an action profile $\boldsymbol{\alpha}$ as exogenously given, and players have no incentive to deviate if all firms truthfully report their messages. Moreover, the incentive for truth-telling is strict.[8]

We consider the auxiliary problem with

$$v_i((g(\boldsymbol{m}))|\boldsymbol{w}) = \begin{cases} \mathbb{E}_{\boldsymbol{w}\sim p(\cdot|\boldsymbol{\alpha})}\phi_i(\alpha_i, w_i) - 1 & m_i = w_i, \\ \mathbb{E}_{\boldsymbol{w}\sim p(\cdot|\boldsymbol{\alpha})}\phi_i(\alpha_i, w_i) & m_i \neq w_i. \end{cases}$$

According to our Theorem 2, there exists a transfer $\hat{\boldsymbol{t}}$ that implements the Bayesian-Nash equilibrium in the auxiliary problem. Therefore, this transfer $\hat{\boldsymbol{t}}$ induces a strict Bayesian-Nash equilibrium in the original problem.

Corollary 1. *There exists a transfer that induces truthful report in a strict Bayesian-Nash equilibrium.*

This is the Theorem 2 of Kandori and Matsushima (1998).

5.2 Information Elicitation

Miller et al. (2005) consider an information elicitation problem, also known as the crowdsourcing problem. Many decision processes depend on eliciting truthful evaluations from n agents, such as online recommender systems and academic reviewing. The mechanism designer hopes to design a reward system to induce honest reports.

Consider the online recommender systems. The quality of the product, ω, is a random state chosen from the set Ω. Each agent i privately observes a signal

[8] See their Section 4 for a detailed discussion.

θ_i about the product quality, where θ_i is drawn from the set Θ. The conditional distribution of θ_i given ω is denoted by $f_i(\theta|\omega) = \Pr(\theta_i = \theta|\omega)$. Miller et al. (2005) assume that $\theta_1, \cdots, \theta_n$ are conditionally independently and identically distributed: $f_i(\theta|\omega) \equiv f(\theta|\omega)$. They assume the stochastic relevance condition we mentioned earlier. It is stronger than our non-degenerate correlated condition. Here, we only need to assume the non-degenerate correlated condition.

The principal asks the agents to report their private signals. She designs a reward system that assigns to agent i a reward $t_i(\boldsymbol{\theta}')$ based on the report profile $\boldsymbol{\theta}'$. The payoff of each agent is $u_i(\boldsymbol{\theta}', \theta_i) = t_i(\boldsymbol{\theta}')$. Given this payoff structure, even if there is no transfer at all $\boldsymbol{t} \equiv \mathbf{0}$, truthful reporting by all agents forms a Bayesian-Nash equilibrium. Their research question is: is it possible to design a transfer $\boldsymbol{t}$ such that truthful reporting forms a strict Bayesian-Nash equilibrium? Truthtelling is the unique best response for one agent in a *strict* Bayesian-Nash equilibrium, provided that other agents tell the truth

$$\mathbb{E}_{\boldsymbol{\theta}_{-i}}\left[t_i(\theta_i, \boldsymbol{\theta}_{-i})|\theta_i\right] > \mathbb{E}_{\boldsymbol{\theta}_{-i}}\left[t_i(\theta_i', \boldsymbol{\theta}_{-i})|\theta_i\right], \forall i = 1, \cdots, n.$$

To apply our results, consider an auxiliary problem with an allocation rule such that $v_i(g(\boldsymbol{\theta}')|\boldsymbol{\theta}) = \phi_i(\theta_i'|\theta_i)$. Namely, the agent i's value gained from the allocation rule g only depends on agent i's type θ_i and his report θ_i'. The payoff of each agent is

$$u_i(\boldsymbol{\theta}', \theta_i) = \phi(\theta_i'|\theta_i) + t_i(\boldsymbol{\theta}').$$

(In the original problem, $\phi_i(\theta_i'|\theta_i) \equiv 0$ for all i, θ_i and θ_i'.) Let

$$\hat{\phi}_i(\theta_i'|\theta_i) = \begin{cases} -1 & \theta_i' = \theta_i, \\ 0 & \text{otherwise.} \end{cases}$$

According to Theorem 2, there exists a transfer $\hat{\boldsymbol{t}}$ that implements the Bayesian-Nash equilibrium in the auxiliary problem. Therefore, this transfer $\hat{\boldsymbol{t}}$ induces a strict Bayesian-Nash equilibrium in the original problem, which coincides with Proposition 1 in Miller et al. (2005).

Then we consider scenarios with a reporting cost. Miller et al. (2005) assume a fixed and identity-independent cost $c > 0$ of report, namely $\phi_i(\theta_i'|\theta_i) \equiv -c$ for all i, θ_i and θ_i'. We can consider an auxiliary problem such that

$$\tilde{\phi}_i(\theta_i'|\theta_i) = \begin{cases} -c - 1 & \theta_i' = \theta_i, \\ -c & \text{otherwise.} \end{cases}$$

According to Theorem 2, there exists a transfer $\tilde{\boldsymbol{t}}$ that implements the Bayesian-Nash equilibrium in the auxiliary problem with reporting cost. Therefore, this transfer $\tilde{\boldsymbol{t}}$ induces a strict Bayesian-Nash equilibrium in the original problem with cost, which coincides with Proposition 2 in Miller et al. (2005).

However, assuming a fixed and identity-independent cost is restrictive in reality. More importantly, the reporting cost generates another economic incentive to misreport beyond the model in Miller et al. (2005). For example, when assessing a research paper, submitting a report for revise-and-resubmit is more costly than suggesting a rejection. This may create an incentive to reject the paper. This moral hazard problem becomes severe as the paper length grows. Similar issues can be found when collecting data in fieldwork or lab experiment.

Even with general cost structures, we can still address the information elicitation problem. Let $\phi_i(\theta_i'|\theta_i)$ denotes the cost of reporting θ_i' when agent i's true type is θ_i. We can define the auxiliary problem by assuming

$$\bar{\phi}_i(\theta_i'|\theta_i) = \begin{cases} \phi_i(\theta_i'|\theta_i) - 1 & \theta_i' = \theta_i, \\ \phi_i(\theta_i'|\theta_i) & \text{otherwise.} \end{cases}$$

By Theorem 2, there always exists a transfer rule $\bar{\boldsymbol{t}}$ in this auxiliary problem. Therefore, this transfer $\bar{\boldsymbol{t}}$ induces a strict Bayesian-Nash equilibrium with general cost structure. We thereby get a stronger result with a weaker assumption (non-degenerate correlated condition) than Miller et al. (2005).

Corollary 2. *Under non-degenerate correlated distribution, for any reporting cost ϕ, there exists a transfer $\boldsymbol{t}$ such that truthful reports forms a strict Bayesian-Nash* *equilibrium.*

5.3 Implementability with Independent Common Prior

An extremely simple case for the implementability problem is what happens if we have independent common prior. Müller et al. (2007) directly extend Rochet (1987)'s original proof and obtains the following result.

Corollary 3. *If the common prior distribution is independent, an allocation g is implementable in Bayesian mechanism if and only if the expected value difference EDv satisfies cyclic monotonicity, i.e, for all i and for any finite types $\theta^{(1)}, \cdots, \theta^{(k)} \in \Theta$,*

$$EDv_i(\theta^{(1)}, \theta^{(2)}) + EDv_i(\theta^{(2)}, \theta^{(3)}) + \cdots + EDv_i(\theta^{(k)}, \theta^{(1)}) \geq 0. \tag{4}$$

Yet, this is also a corollary/degenerate case of our Theorem 1. When the prior distribution is independent, the entire type space Θ is a consistent class for each agent, leading to this result.

6 Concluding Remarks

Our paper assumes the type space to be finite. One direction for future research is to generalize our results to infinite type space. Moreover, as we do not impose (ex-ante) budget balance, IR constraints are irrelevant in our model. Alternatively, we can achieve budget balance by shifting the transfer function if we do not impose any IR constraints. One potential extension is to impose both IR constraints and the budget balance. The IR constraints would naturally depend on the context of interests, such as players' outside options.

A Omitted Proofs

Lemma 2. ***Network Flow****. Given $r, s, k \in \{1, \cdots, m\}$ where m is an integer, $m(m-1)$ non-negative real $y_{s \circ r}$, $m(m-1)$ real numbers $Dv_{s \circ r}$ where $s \neq r$, if*

1. $\sum_{s, s\neq k} y_{s \circ k} = \sum_{r, r \neq k} y_{k \circ r}, \quad \forall k$
2. *for all finite sequence $s_1, \cdots, s_\kappa \in \{1, 2, \cdots, m\}$,*

$$Dv_{s_1 \circ s_2} + \cdots + Dv_{s_\kappa \circ s_1} \geq 0,$$

then $\sum_{(s,r), s \neq r} Dv_{s \circ r} y_{s \circ r} \geq 0$.

Proof. Construct a directed graph $G(V, E)$ where $V = \{1, 2, ..., m\}$ and each directed edge (s, r) has flow $y_{s \circ r}$. Then the first condition of the lemma is that the sum of each node's in-degree flow equals the sum of its out-degree flow.

Consider the following updating process. If there exists a cycle $(s_1, s_2, \cdots, s_k)$ such that $y_{s_1 \circ s_2} \times \cdots \times y_{s_\kappa \circ s_1} \neq 0$, i.e., all flow are non-zero. Then let $y = \min\{y_{s_1 \circ s_2}, \cdots, y_{s_\kappa \circ s_1}\}$, and update the values of $y_{s_1 \circ s_2}, y_{s_2 \circ s_3}, \cdots, y_{s_\kappa \circ s_1}$ as the following:

$$\begin{aligned} y_{s_1 \circ s_2} &\leftarrow y_{s_1 \circ s_2} - y \\ &\cdots \\ y_{s_\kappa \circ s_1} &\leftarrow y_{s_\kappa \circ s_1} - y \end{aligned}$$

After the updating process, $\sum_{s, s \neq k} y_{s \circ k} = \sum_{r, r \neq k} y_{k \circ r}, \forall k$ remains, because a cycle cross node k must have the same in-degree and out-degree then both right hand side and left hand side minus a constant.

After the updating process, the flow of each edge remains non-negative. The left hand side of the target inequality, $\sum_{(s,r), s \neq r} Dv_{s \circ r} y_{s \circ r}$, is subtracted by $y(Dv_{s_1 \circ s_2} + \cdots + Dv_{s_\kappa \circ s_1})$. This implies that the left hand side of the target inequality, $\sum_{(s,r), s \neq r} Dv_{s \circ r} y_{s \circ r}$, weakly decreases after the update.

Repeat the above updating process until there exists not cycle $(s_1, s_2, \cdots, s_k)$ such that $y_{s_1 \circ s_2} \times \cdots \times y_{s_\kappa \circ s_1} \neq 0$. Then, each directed cycle of the remaining graph must have a edge with zero flow. Now we delete the edge with zero flow, the remaining graph has no directed cycle.

If the remaining graph has no edge, then all $y_{sor} = 0$. The target also becomes 0. Thus, $\sum_{(s,r),s\neq r} Dv_{sor}y_{sor}$ weakly decreases to 0 and hence $\sum_{(s,r),s\neq r} Dv_{sor}y_{sor} \geq 0$ in the original problem. Otherwise, there must exist a node in remaining graph which has no in-degree and has at least one out-degree. But this node has different sum of in-degree and out-degree, which brings a contradiction.

□

Proof of Theorem 1. Necessity. Suppose a mechanism $(g, \boldsymbol{t})$ is **BIC**. Rewrite the (BIC) condition with conditional probabilities, we have

$$\sum_{\boldsymbol{\theta}_{-i}} \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta_i)t_i(\theta_i, \boldsymbol{\theta}_{-i}) \geq \sum_{\boldsymbol{\theta}_{-i}} \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta_i)t_i(\theta_i', \boldsymbol{\theta}_{-i}) - EDv_i(\theta_i', \theta_i) \qquad \text{(BIC')}$$

For any finite sequence $\theta^{(0)}, \theta^{(1)}, \cdots, \theta^{(k)} \in \Theta_i^{(j)}$ with $\theta^{(0)} = \theta^{(k)}$, we have $\forall \kappa \in \{1, 2, \cdots, k\}$,

$$\sum_{\boldsymbol{\theta}_{-i}} \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta^{(\kappa)})t_i(\theta^{(\kappa)}, \boldsymbol{\theta}_{-i}) \geq \sum_{\boldsymbol{\theta}_{-i}} \mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta^{(\kappa)})t_i(\theta^{(\kappa-1)}, \boldsymbol{\theta}_{-i}) - EDv_i(\theta^{(\kappa-1)}, \theta^{(\kappa)})$$

Since the types are chosen in the same consistent class, the conditional distributions are the same. Summing over the inequalities above, we obtain

$$EDv_i(\theta^{(1)}, \theta^{(2)}) + EDv_i(\theta^{(2)}, \theta^{(3)}) + \cdots + EDv_i(\theta^{(k)}, \theta^{(1)}) \geq 0.$$

Sufficiency. This proof proceeds in four steps. Let $\boldsymbol{P}_i(\theta^r)$ denote the posterior of agent i conditional on his type being θ^r. It is a m^{n-1}-dimensional row vector of $\mathbb{P}_{-i}(\boldsymbol{\theta}_{-i}|\theta^r)$ with m^{n-1} different $\boldsymbol{\theta}_{-i}$ arranged in lexicographic order.

$$\boldsymbol{P}_i(\theta^r) = (\mathbb{P}_{-i}(\theta^1, \cdots, \theta^1|\theta^r), \cdots, \mathbb{P}_{-i}(\theta^m, \cdots, \theta^m|\theta^r)).$$

Similarly, let $\boldsymbol{T}_i$ denote the m^n dimension column vector that consists of $t_i(\boldsymbol{\theta})$ with $\boldsymbol{\theta}$ in lexicographic order.

Step 1: Convert the BIC condition to a linear programming feasibility problem.

In a Bayesian equilibrium, for each $\theta^r, \theta^s \in \Theta$, for each agent i, we have

$$\boldsymbol{P}_i(\theta^r)\boldsymbol{T}_i^r \geq \boldsymbol{P}_i(\theta^r)\boldsymbol{T}_i^s - EDv_i(\theta^s, \theta^r).$$

Rearrange the terms,

$$\boldsymbol{P}_i(\theta^r)\boldsymbol{T}_i^r - \boldsymbol{P}_i(\theta^r)\boldsymbol{T}_i^s \geq -EDv_i(\theta^s, \theta^r)$$

where the left hand side can be further expressed as the inner product of two vectors. One is the column vector $\boldsymbol{T}_i$, and the other is a row vector composed of $\boldsymbol{P}_i(\theta^r), -\boldsymbol{P}_i(\theta^r)$ and $\boldsymbol{0}$ in some order. Let $\boldsymbol{P}_i^{r\circ s}$ denote this m^n dimensional row vector.

Now, we construct matrix A_i in the linear programming problem. Each row in A_i is $\boldsymbol{P}_i^{r\circ s}$. We sort them according to row index $r \circ s$ with $r \neq s$ in a lexicographic order: $1 \circ 2, 1 \circ 3, \cdots, 1 \circ m, \cdots, m \circ 1, m \circ 2, \cdots, m \circ (m-1)$. The matrix A_i is

$$\begin{pmatrix} \boldsymbol{P}_i(\theta^1) & -\boldsymbol{P}_i(\theta^1) & \boldsymbol{0} & \boldsymbol{0} & \cdots \\ \boldsymbol{P}_i(\theta^1) & \boldsymbol{0} & -\boldsymbol{P}_i(\theta^1) & \boldsymbol{0} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \boldsymbol{P}_i(\theta^1) & \boldsymbol{0} & \boldsymbol{0} & \cdots & -\boldsymbol{P}_i(\theta^1) \\ -\boldsymbol{P}_i(\theta^2) & \boldsymbol{P}_i(\theta^2) & \boldsymbol{0} & \boldsymbol{0} & \cdots \\ \boldsymbol{0} & \boldsymbol{P}_i(\theta^2) & -\boldsymbol{P}_i(\theta^2) & \boldsymbol{0} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{pmatrix}$$

Then, we construct a $m(m-1)$ dimension column vector $\boldsymbol{b}_i$. Let $b_i^{r\circ s} = -EDv_i(\theta^s, \theta^r)$ where $s \neq r$. $\boldsymbol{b}_i$ consists of all values of $b_i^{r\circ s}$ with $r \circ s$ in lexicographic order: $1 \circ 2, 1 \circ 3, \cdots, 1 \circ m, \cdots, m \circ 1, m \circ 2, \cdots, m \circ (m-1)$. For example, when $m = 3$, the column vector $\boldsymbol{b}_i$ is

$$(-EDv_i(\theta^2, \theta^1), -EDv_i(\theta^3, \theta^1), -EDv_i(\theta^1, \theta^2), -EDv_i(\theta^3, \theta^2), -EDv_i(\theta^1, \theta^3), -EDv_i(\theta^2, \theta^3))^\top$$

Then the (BIC) conditions of transfer function could be written as the constraints in the following linear programming,

$$\min \mathbf{0}^\top \boldsymbol{T}_i, A_i \boldsymbol{T}_i \geq \boldsymbol{b}_i, \boldsymbol{T}_i \geq \mathbf{0}$$

We only need to prove that the above linear programming problem is feasible.

Step 2: Construct the dual linear programming problem.

Now, we construct the dual linear programming of the above linear programming problem. Let y_i^{ros} where $r \neq s$ denote $m(m-1)$ decision variables and $\boldsymbol{y}_i$ collect all $m(m-1)$ values in a lexicographic order: $1 \circ 2, 1 \circ 3, \cdots, 1 \circ m, \cdots, m \circ 1, m \circ 2, \cdots, m \circ (m-1)$.

$$\max \boldsymbol{b}_i^\top \boldsymbol{y}_i, \boldsymbol{y}_i^\top A_i \leq \mathbf{0}, \boldsymbol{y}_i \geq \mathbf{0}$$

By the strong duality theorem in linear programming, the original problem is feasible if and only if the optimum in the duality problem is zero. Clearly, zero is achievable by $\boldsymbol{y}_i = \mathbf{0}$. In the following, we prove $\boldsymbol{b}_i^\top \boldsymbol{y}_i \leq 0$.

By $\boldsymbol{y}_i^\top A_i \leq \mathbf{0}$, we have

$$\sum_{s, s \neq r} y_i^{ros} \boldsymbol{P}_i(\theta^r) - \sum_{s, s \neq r} y_i^{sor} \boldsymbol{P}_i(\theta^s) \leq \mathbf{0}, \quad \forall r \in \{1, 2, \cdots, m\}.$$

Sum over r, we have an equality

$$\sum_r \left(\sum_{s, s \neq r} y_i^{ros} \boldsymbol{P}_i(\theta^r) - \sum_{s, s \neq r} y_i^{sor} \boldsymbol{P}_i(\theta^s) \right) = \mathbf{0}$$

Then all the inequalities of $\boldsymbol{y}_i^\top A_i$ must be binding. That is

$$\sum_{s, s \neq r} y_i^{ros} \boldsymbol{P}_i(\theta^r) - \sum_{s, s \neq r} y_i^{sor} \boldsymbol{P}_i(\theta^s) = \mathbf{0}, \quad \forall r \in \{1, 2, \cdots, m\}.$$

Since the L_1 norm of each $\boldsymbol{P}_i(\theta^r)$ is 1, take the L_1 norm to get

$$\sum_{s,s\neq r} y_i^{ros} - \sum_{s,s\neq r} y_i^{sor} = 0, \quad \forall r \in \{1,2,\cdots,m\}.$$

Step 3: We show that if θ^r, θ^s are not in the same consistent class, then $y_i^{ros} = y_i^{sor} = 0$.

Order the consistent classes of agent i such that the L_2 norm of $\boldsymbol{P}_i(\theta)$ where $\theta \in \Theta_i^{(m)}$ is decreasing in m. Now consider the first consistent class of agent i, i.e., $\Theta_i^{(1)}$. Take any type in this set $\theta^r \in \Theta_i^{(1)}$, we have

$$\sum_{s,s\neq r} y_i^{ros} \boldsymbol{P}_i(\theta^r) - \sum_{s,s\neq r} y_i^{sor} \boldsymbol{P}_i(\theta^s) = \boldsymbol{0},$$

then

$$\begin{aligned}\sum_{s,s\neq r} y_i^{ros} \|\boldsymbol{P}_i(\theta^r)\|_2 &= \left\|\sum_{s,s\neq r} y_i^{sor} \boldsymbol{P}_i(\theta^s)\right\|_2 \leq \sum_{s,s\neq r} \|y_i^{sor} \boldsymbol{P}_i(\theta^s)\|_2 \\ &\leq \sum_{s,s\neq r} y_i^{sor} \|\boldsymbol{P}_i(\theta^r)\|_2 = \|\boldsymbol{P}_i(\theta^r)\|_2 \sum_{s,s\neq r} y_i^{ros}\end{aligned}$$

Then all the inequalities above must be binding. Consider the triangle inequality. For all s such that $s \neq r$, either we have $y_i^{sor} = 0$ or $\boldsymbol{P}_i(\theta^s)$ is collinear with $\boldsymbol{P}_i(\theta^r)$. Since they all have unit L_1 norm, collinearity implies $\boldsymbol{P}_i(\theta^s) = \boldsymbol{P}_i(\theta^r)$, i.e., $\theta^s \in \Theta_i^{(1)}$. Then for all $\theta^s \notin \Theta_i^{(1)}$ we have $y_i^{sor} = 0$. So by step 2,

$$\sum_{s,s\neq r} y_i^{ros} = \sum_{s,s\neq r,\theta^s \in \Theta_i^{(1)}} y_i^{sor}.$$

Sum over all types in $\Theta_i^{(1)}$,

$$\sum_{\theta^r \in \Theta_i^{(1)}} \sum_{s,s\neq r} y_i^{ros} = \sum_{\theta^r \in \Theta_i^{(1)}} \sum_{s,s\neq r,\theta^s \in \Theta_i^{(1)}} y_i^{sor}.$$

That is

$$\sum_{\theta^r \in \Theta_i^{(1)}, \theta^s \notin \Theta_i^{(1)}} y_i^{ros} = 0$$

Since all $\boldsymbol{y}_i \geq \mathbf{0}$, then $y_i^{ros} = 0$ for $\theta^s \notin \Theta_i^{(1)}$ and $\theta^r \in \Theta_i^{(1)}$. So far we have proven that for $\theta^s \notin \Theta_i^{(1)}$ and $\theta^r \in \Theta_i^{(1)}$, $y_i^{ros} = y_i^{sor} = 0$.

Similarly, we obtain that for all θ^s, θ^r in distinct consistent classes, $y_i^{sor} = y_i^{ros} = 0$. And for any $\theta^r \in \Theta_i^{(j)}$ we also have

$$\sum_{\theta^s \in \Theta_i^{(j)}, s \neq r} y_i^{ros} = \sum_{\theta^s \in \Theta_i^{(j)}, s \neq r} y_i^{sor}.$$

Step 4: Apply the Network Flow Lemma 2 to every consistent class $\Theta_i^{(j)}$.

In every consistent class $\Theta_i^{(j)}$, the equation above and cyclic monotonicity of EDv satisfy the assumption of the Network Flow Lemma 2. Thus,

$$-\boldsymbol{b}_i^\top \boldsymbol{y}_i = \sum_{(s,r), r \neq s} (-b_i^{ros}) y_i^{ros} = \sum_j \sum_{\theta^r, \theta^s \in \Theta_i^{(j)}, \theta^r \neq \theta^s} (-b_i^{ros}) y_i^{ros} \geq 0.$$

□

References

James J. Anton and Dennis A. Yao. Coordination in split award auctions. *Quarterly Journal of Economics*, 107(2):681–707, 1992.

Dirk Bergemann, Stephen Morris, and Satoru Takahashi. Efficient auctions and interdependent types. *American Economic Review*, 102(3):319–324, 2012.

Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

Jacques Crémer and Richard P McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.

Alex Gershkov, Jacob K Goeree, Alexey Kushnir, Benny Moldovanu, and Xianwen Shi. On the equivalence of bayesian and dominant strategy implementation. *Econometrica*, 81(1):197–220, 2013.

Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

Felix Jarman and Vincent Meisner. Ex-post optimal knapsack procurement. *Journal of Economic Theory*, 171:35–63, 2017.

Michihiro Kandori and Hitoshi Matsushima. Private observation, communication and collusion. *Econometrica*, 66(3):627–652, 1998.

Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.

Hervé Moulin. Almost budget-balanced vcg mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.

Rudolf Müller, Andrés Perea, and Sascha Wolf. Weak monotonicity and bayes–nash incentive compatibility. *Games and Economic Behavior*, 61(2):344–358, 2007.

Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

Martin Pesendorfer. Pollution claim settlements under correlated information. *Journal of Economic Theory*, 79(1):72–105, 1998.

Rafeal Rob. Pollution claim settlements under private information. *Journal of Economic Theory*, 47(2):307–333, 1989.

Jean-Charles Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16(2):191–200, 1987.

Rakesh V. Vohra. *Mechanism Design: A Linear Programming Approach*. Cambridge University Press, 2011.