

A Note on Optimal Allocation with Costly Verification.*

Albin Erlanson Andreas Kleiner

30th August 2016

Abstract

We revisit the problem of a principal allocating an indivisible good with costly verification, as it was formulated and analyzed by Ben-Porath et al. (2014). We establish, in this setting, a general equivalence between Bayesian and dominant-strategy incentive compatible mechanisms. We also provide a simple proof showing that the optimal mechanism is a threshold mechanism.

Keywords: Optimal mechanisms; Costly verification; BIC and DIC equivalence

JEL classification: D82

*We are grateful to Daniel Krähmer, Bart Lipman and Benny Moldovanu for detailed comments and discussions. Albin Erlanson gratefully acknowledges financial support from the European Research Council and the Jan Wallander and Tom Hedelius Foundation. Erlanson: Department of Economics at Stockholm School of Economics, albin.erlanson@hhs.se; Kleiner: Department of Economics, W.P. Carey School of Business at Arizona State University, andreas.kleiner@asu.edu

1 Introduction

A crucial part of designing mechanisms is to elicit private information. It is often assumed that private information cannot be verified in any way. However, there are many real-life situations when information indeed is verifiable as it may be based on hard information. In a recent paper, Ben-Porath et al. (2014) (henceforth called BDL) analyzed costly verification in a model where a principal allocates an indivisible good to privately informed agents. They showed that the optimal Bayesian incentive compatible mechanism for the principal is in the class of “threshold mechanisms.”¹ In a threshold mechanism, provided that some report is above the threshold, the agent with the highest reported value is verified and gets the object if she was truthful, and the object is randomly allocated according to a given probability distribution otherwise. If an agent is caught lying she does not receive the object.

Somewhat surprisingly, the optimal mechanism is dominant-strategy incentive compatible. Thus, the optimal mechanism does not use the extra flexibility that a Bayesian mechanism offers. We explain this observation by establishing a more general equivalence: for any Bayesian incentive compatible mechanism there exists an “equivalent” mechanism that is implementable in dominant strategies and induces the same expected verification costs. Similar equivalence results exist in the standard one-dimensional mechanism design setting with single-crossing and quasi-linear utilities (Manelli and Vincent, 2010; Gershkov et al., 2013).

In the second part of this note we provide an alternative proof for the optimality of threshold mechanisms² by using insights from the literature on reduced form auctions.³ To prove the optimality of threshold mechanisms we observe that the relevant incentive constraints are formulated in terms of reduced forms. Thus, we can restate the optimization problem using only reduced forms and optimize over them directly, which is significantly easier. A characterization of feasible reduced form rules is readily available due to Border (1991), and we can show that threshold mechanisms are optimal. Our approach of using reduced forms to solve for optimal mechanisms is one example among several recent papers, for other examples see Mierendorff (2016), Mylovanov and Zapechelnyuk (2015) and Pai and Vohra (2014).

The rest of the note is organized as follows. In Section 2 we introduce the model. In Section 3 we formalize and prove the equivalence between Bayesian and dominant-strategy incentive compatible mechanisms. In Section 4 we provide our alternative proof of the optimality of threshold mechanisms.

¹BDL further showed that a “favored-agent” mechanism is the optimal threshold mechanism. Thus, a favored-agent mechanism is optimal among all incentive compatible mechanisms.

²Another proof of the optimality of threshold mechanisms can be found in Lipman (2015).

³A *reduced form* maps the type of an agent into the expected probability of being allocated the object. The set of feasible reduced forms has an explicit description (Border, 1991) and a nice combinatorial structure (Che et al., 2013).

2 Model and incentive constraints

The principal wants to allocate one indivisible object among agents in $\mathcal{I} \equiv \{1, \dots, I\}$. Agents are privately informed about their types $t_i \in \mathcal{T}_i \equiv [\underline{t}_i, \bar{t}_i]$, where $-\infty < \underline{t}_i < \bar{t}_i < \infty$. The principal receives value t_i when the object is allocated to an agent with type t_i . Monetary transfers are not possible, and all agents strictly prefer to receive the object. Thus, the payoff of an agent is simply the probability of receiving the object.⁴ Types are independently distributed with distribution function F_i and corresponding density f_i . A profile of types is denoted by $t \in \mathcal{T} \equiv \prod_{i \in \mathcal{I}} \mathcal{T}_i$. The principal can verify agent i at a given cost of c_i , in which case the type of agent i is perfectly revealed. The goal of the principal is to maximize the expected value of allocating the good less the expected cost of verification.

By invoking a revelation principle, it is without loss of generality to consider only direct and incentive compatible mechanisms. Denote by $p_i : \mathcal{T} \rightarrow [0, 1]$ the probability that agent i is assigned the good, and by $q_i : \mathcal{T} \rightarrow [0, 1]$ the probability that agent i is verified and assigned the good given that he reported truthfully. A *mechanism* is a tuple $(p, q) = (p_i, q_i)_{i \in \mathcal{I}}$. A mechanism (p, q) is *feasible* if, for all $t \in \mathcal{T}$, $\sum p_i(t) \leq 1$ and $q_i(t) \leq p_i(t)$ for each $i \in \mathcal{I}$. Without loss of generality, we can assume that if an agent is verified and is lying then he will not be assigned the object.

A mechanism is *incentive compatible*, if truthtelling is an equilibrium in the game induced by the mechanism. We consider two notions of incentive compatibility: Bayesian and dominant-strategy incentive compatibility,⁵ which are characterized in the following two lemmas.

Lemma 1. *A mechanism (p, q) is Bayesian incentive compatible (BIC) if and only if, for all $i \in \mathcal{I}$ and all $t_i \in \mathcal{T}_i$,*

$$\inf_{t'_i \in \mathcal{T}_i} \mathbb{E}_{t_{-i}} [p_i(t'_i, t_{-i})] \geq \mathbb{E}_{t_{-i}} [p_i(t_i, t_{-i})] - \mathbb{E}_{t_{-i}} [q_i(t_i, t_{-i})]. \quad (1)$$

Lemma 2. *A mechanism (p, q) is dominant-strategy incentive compatible (DIC) if and only if, for all $i \in \mathcal{I}$, all $t_i \in \mathcal{T}_i$, and all $t_{-i} \in \mathcal{T}_{-i}$*

$$\inf_{t'_i \in \mathcal{T}_i} p_i(t'_i, t_{-i}) \geq p_i(t_i, t_{-i}) - q_i(t_i, t_{-i}). \quad (2)$$

Consider an agent with type t'_i . By reporting his type t'_i truthfully he receives the good with probability $\mathbb{E}_{t_{-i}} [p_i(t'_i, t_{-i})]$. Now, suppose instead that he lies about his type and reports $t_i \neq t'_i$. Then he receives the good with probability $\mathbb{E}_{t_{-i}} [p_i(t_i, t_{-i})]$ conditional upon not being verified, but if he is verified then he is not assigned the good. Thus,

⁴Agents' cardinal preferences can depend on their types, but intensities of the agents' preferences do not play a role in the analysis.

⁵A mechanism is dominant-strategy incentive compatible, if truthtelling is a best response for every type to any possible strategy used by the other agents. This is the weak notion of dominant strategies, which is commonly used in mechanism design.

the probability of receiving the good when lying is $\mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})] - \mathbb{E}_{t_{-i}}[q_i(t_i, t_{-i})]$. This is the right-hand side of the BIC constraints in Lemma 1 above. Note that the payoff from lying is independent of the agent's true type t_i , and the inequality must hold for all types. In particular, the worst payoff from telling the truth must be better than lying. This is the infimum on the left hand side in Lemma 1 above. The only difference for the characterization of DIC mechanisms is that this argument must hold pointwise instead of averaging over the other agents' types.

In this note, instead of using the ex-post allocation rule p , we will sometimes use the corresponding *interim allocation rule* (also called reduced form) \hat{p} , where $\hat{p}_i(t_i) = \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$. A relevant question is whether an interim allocation rule \hat{p} is *feasible* in the sense that there exists a feasible ex-post allocation rule that induces \hat{p} . This question has been answered by Border (1991) and Mierendorff (2011), who characterized the set of feasible reduced forms: a monotone interim allocation rule is feasible if and only if, for all $(\alpha_1, \dots, \alpha_n) \in \mathcal{T}$,

$$\sum_i \int_{\alpha_i}^{\bar{t}_i} \hat{p}_i(t_i) dF_i(t_i) \leq 1 - \prod_i F_i(\alpha_i). \quad (\text{Border})$$

This condition is necessary for a reduced form to be feasible: the left-hand side, denoting the probability that an agent i with type above α_i wins the object, must clearly be lower than the probability that there is an agent i with type above α_i , which is written on the right-hand side.⁶ The content of Border's theorem is to show that the above condition is also sufficient for a nondecreasing reduced form to be feasible.

3 BIC-DIC equivalence

In this section we will establish an equivalence between BIC and DIC mechanisms. We will first formalize the equivalence notion, and then state and prove the equivalence theorem.

BDL showed that the optimal BIC mechanism satisfies the stronger dominant-strategy incentive constraints. At first sight it might be surprising that the optimal mechanism does not use the extra degrees of freedom that Bayesian incentive constraints offer in order to save on verification costs. We show that there is a deeper connection underlying this observation: for any BIC mechanism (p, q) there exists an "equivalent" DIC mechanism (\tilde{p}, \tilde{q}) that induces the same expected verification costs and the same interim allocation rules.

To compare Bayesian with dominant-strategy incentive constraints, we average the pointwise constraints in (2) by taking expectations over reports of the other agents. Consequently, any DIC mechanism must satisfy

$$\mathbb{E}_{t_{-i}}[\inf_{t'_i \in \mathcal{T}_i} p_i(t'_i, t_{-i})] \geq \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i}) - q_i(t_i, t_{-i})]. \quad (3)$$

⁶We use this direction to construct an upper bound on the objective function in Section 4.

Note that the left-hand side of (3) is in general strictly smaller than the left-hand side of the Bayesian constraints in (1). Therefore, in general it is strictly more costly to implement an allocation rule in dominant strategies; only if the expectation operator commutes with the infimum operator, can a rule be implemented at the same verification costs in dominant strategies. Below we show that for any allocation rule, there exists an equivalent rule for which the expectation operator commutes with infimum operator, thereby explaining why it is without loss of optimality to consider only dominant-strategy incentive compatible mechanisms.

Definition 1. Two mechanisms (p, q) and (\tilde{p}, \tilde{q}) are *equivalent* if $\mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})] = \mathbb{E}_{t_{-i}}[\tilde{p}_i(t_i, t_{-i})]$ and $\mathbb{E}_{t_{-i}}[q_i(t_i, t_{-i})] = \mathbb{E}_{t_{-i}}[\tilde{q}_i(t_i, t_{-i})]$ for all i and almost every t_i .

Intuitively, we say that two mechanisms are equivalent if they essentially induce the same interim expected allocation and verification probabilities. Now we can state the equivalence result between BIC and DIC mechanism.

Theorem 1. *For any BIC mechanism (p, q) there exists an equivalent DIC mechanism (\tilde{p}, \tilde{q}) .*

Proof. Let (p, q) be a BIC mechanism. We will construct an equivalent DIC mechanism (\tilde{p}, \tilde{q}) .

Step 1: Constructing allocation rules \tilde{p}

We will define a new type space such that the interim allocation rules are nondecreasing, and then we can construct pointwise nondecreasing allocation rules p' .

Let $\sigma_i : \mathcal{T}_i \rightarrow \mathbb{R}$ be defined as $\sigma_i(t_i) := \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$. Let the new type space be $\tilde{\mathcal{T}}_i \equiv \{x \in \mathbb{R} \mid \sigma_i(t_i) = x \text{ for some } t_i \in \mathcal{T}_i\}$, i.e., the new type space $\tilde{\mathcal{T}}_i$ is the image of σ_i . Denote a type in $\tilde{\mathcal{T}}_i$ by \tilde{t}_i , and let $\tilde{\mathcal{T}} \equiv \prod_{i \in \mathcal{I}} \tilde{\mathcal{T}}_i$. Let G_i denote the distribution function on the type space $\tilde{\mathcal{T}}_i$ such that

$$G_i(\tilde{t}_i) := \int_{t_i: \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})] \leq \tilde{t}_i} dF_i(t_i) \quad \text{for all } i \in \mathcal{I}.$$

Let the allocation rule on the new type space $p'_i : \tilde{\mathcal{T}} \rightarrow [0, 1]$ be defined such that $\mathbb{E}_{\tilde{t}_{-i}}[p'_i(\tilde{t}_i, \tilde{t}_{-i})] = \tilde{t}_i$ for all $\tilde{t}_i \in \tilde{\mathcal{T}}_i$. The existence of such an allocation rule is guaranteed by Border's characterization of reduced forms (Border, 1991).⁷ Thus, the interim allocation rules of p'_i are nondecreasing and Theorem 1 in Gershkov et al. (2013) implies that there

⁷Indeed, denoting by μ the measure induced by the cdf. F , we get for all $\alpha_i \in \tilde{\mathcal{T}}_i$ (setting $A_i := \{\tilde{t}_i \in \tilde{\mathcal{T}}_i \mid \tilde{t}_i \geq \alpha_i\}$):

$$\sum_i \int_{A_i} \tilde{t}_i dG_i(\tilde{t}_i) = \sum_i \int_{\sigma^{-1}(A_i)} \mathbb{E}_{t_{-i}} p_i(t_i, t_{-i}) dF_i(t_i) \leq 1 - \prod_i [1 - \mu(\sigma^{-1}(A_i))] = 1 - \prod_i [1 - G(\alpha_i)].$$

The equalities hold by construction and the inequality holds because p is a feasible reduced form.

exists another feasible allocation rule $p_i''(\cdot, \tilde{t}_{-i})$ that is pointwise nondecreasing and has the same interim allocation rules as $p_i'(\cdot, \tilde{t}_{-i})$.⁸ Now we can define the allocation rules,

$$\tilde{p}_i(t_i, t_{-i}) := p_i''(\sigma_i(t_i), \sigma_{-i}(t_{-i})) \text{ for all } i \in \mathcal{I}. \quad (4)$$

Note that $\mathbb{E}_{t_{-i}}[p_i''(\sigma_i(t_i), \sigma_{-i}(t_{-i}))] = \mathbb{E}_{\tilde{t}_{-i}}[p_i''(\tilde{t}_i, \tilde{t}_{-i})] = \mathbb{E}_{\tilde{t}_{-i}}[p_i'(\tilde{t}_i, \tilde{t}_{-i})] = \tilde{t}_i = \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$. Thus, $\mathbb{E}_{t_{-i}}[\tilde{p}_i(t_i, t_{-i})] = \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$, as desired.

Moreover,

$$\begin{aligned} \mathbb{E}_{t_{-i}}[\inf_{t_i} \tilde{p}_i(t_i, t_{-i})] &= \mathbb{E}_{t_{-i}}[\inf_{t_i} p_i''(\sigma_i(t_i), \sigma_{-i}(t_{-i}))] \\ &= \inf_{t_i} [\mathbb{E}_{t_{-i}} p_i''(\sigma_i(t_i), \sigma_{-i}(t_{-i}))] = \inf_{t_i} [\mathbb{E}_{t_{-i}} \tilde{p}_i(t_i, t_{-i})], \end{aligned}$$

as p_i'' is pointwise nondecreasing.

Step 2: Constructing verification rules \tilde{q}

Now we will construct verification rules \tilde{q}_i such that the mechanism (\tilde{p}, \tilde{q}) is DIC, i.e., satisfies (2), and induces the same expected verification costs. The verification rules are defined as,

$$\tilde{q}_i(t_i, t_{-i}) := \tilde{p}_i(t_i, t_{-i}) - \inf_{t'_i \in \mathcal{T}_i} \tilde{p}_i(t'_i, t_{-i}) \text{ for all } i \in \mathcal{I}. \quad (5)$$

By construction, the incentive constraints in (2) hold as equalities, and it remains to show that the expected verification probabilities are the same.

$$\begin{aligned} \mathbb{E}_{t_{-i}}[\tilde{q}_i(t_i, t_{-i})] &= \mathbb{E}_{t_{-i}}[\tilde{p}_i(t_i, t_{-i})] - \inf_{t'_i \in \mathcal{T}_i} \mathbb{E}_{t_{-i}}[\tilde{p}_i(t'_i, t_{-i})] \\ &= \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})] - \inf_{t'_i \in \mathcal{T}_i} \mathbb{E}_{t_{-i}}[p_i(t'_i, t_{-i})] \\ &\leq \mathbb{E}_{t_{-i}}[q_i(t_i, t_{-i})] \end{aligned}$$

The inequality follows because (p, q) is BIC. Thus, by possibly adding verification we can ensure that $\mathbb{E}_{t_{-i}}[\tilde{q}_i(t)] = \mathbb{E}_{t_{-i}}[q_i(t)]$. \square

4 Threshold mechanisms are optimal

In this section, we follow BDL and analyze the optimal mechanism for the principal. We provide a simple proof showing that it is optimal to use a threshold mechanism.

As a first step in finding the optimal mechanism, BDL construct a relaxation of the principal's optimization problem.⁹ In BDL the objective is $\sum_i \mathbb{E}_t[p_i(t)(t_i - c_i)]$. We consider instead the net-types $t_i - c_i$. This gives us an equivalent formulation of the

⁸Gershkov et al. (2013) show that there is an allocation rule providing the same interim expected utilities (U-equivalence). Since agents in our model only care about the probability with which they receive the good, this implies that the new allocation rule is the same interim allocation rule (P-equivalence).

⁹For a detailed derivation of the relaxed problem see Section V in (Ben-Porath et al., 2014).

optimization problem that is notationally easier to work with. Henceforth, with slight abuse of notation, the type t_i is the net-type $t_i - c_i$.

Define, for each $i \in \mathcal{I}$, $\varphi_i \equiv \inf_{t'_i \in \mathcal{T}_i} \mathbb{E}_{t_{-i}} [p_i(t'_i, t_{-i})]$ and observing that the incentive constraints must be binding in an optimal solution, we get $\mathbb{E}_{t_{-i}} [q_i(t_i, t_{-i})] = \mathbb{E}_{t_{-i}} [p_i(t_i, t_{-i})] - \varphi_i$. Plugging this into the principal's objective function and treating $\varphi \equiv (\varphi_i)_{i \in \mathcal{I}}$ as a parameter, we get the following optimization problem:

$$\begin{aligned} \max_{\{p_i\}_{i \in \mathcal{I}}} \quad & \mathbb{E}_t [\sum_i p_i(t) t_i] & (R) \\ \text{s.t.} \quad & \mathbb{E}_{t_{-i}} [p_i(t_i, t_{-i})] \geq \varphi_i & \forall t_i \in \mathcal{T}_i, i \in \mathcal{I} \\ & \sum_i p_i(t) \leq 1 & \forall t \in \mathcal{T} \\ & 0 \leq p_i(t) & \forall t \in \mathcal{T}, i \in \mathcal{I} \end{aligned}$$

Hence, the problem is to find feasible ex-post allocation rules $p_i : \mathcal{T} \rightarrow [0, 1]$ maximizing the expected value to the principal, subject to an interim incentive constraint that restricts the interim expected value of p_i .

We can simplify the statement of the problem using interim allocation rules:

$$\begin{aligned} \max_{\{\hat{p}_i\}_{i \in \mathcal{I}}} \quad & \sum_i \mathbb{E}_{t_i} [\hat{p}_i(t_i) t_i] \\ \text{s.t.} \quad & \hat{p}_i(t_i) \geq \varphi_i & \forall t_i \in \mathcal{T}_i, i \in \mathcal{I}. \\ & (\text{Border}) \end{aligned}$$

Our key insight to solve this optimization problem is that the constraint set is the infinite-dimensional analogue of a generalized polymatroid. Linear programs with generalized polymatroids as constraints are simple optimization problems that can be solved using a greedy algorithm, that is, an algorithm that chooses the locally optimal solution without taking other constraints into account. Our arguments below generalizes this approach by constructing an analogous solution for the infinite-dimensional problem and showing that it is indeed the optimal solution.

BDL showed that threshold mechanisms solves the optimization problem (R). Following BDL, we define a *threshold mechanism* with threshold α to be a mechanism p with the following reduced form: $\hat{p}_i(t_i) = \prod_{j \neq i} F_j(t_i)$ for $t_i > \alpha$ and $\hat{p}_i(t_i) = \varphi_i$ otherwise. This reduced form can be implemented as follows (where i^* denotes the agent with the highest report). If $t_{i^*} \leq \alpha$, that is, all agents report below the threshold, then each agent i is allocated the object with probability $\frac{\varphi_i}{\prod_{j \neq i} F_j(\alpha)}$. If $t_{i^*} > \alpha$, then agent i^* is verified and conditional that no lie was detected, he receives the object with probability 1. Let

$$\alpha^* = \inf \{ \alpha \in \mathbb{R}_+ \mid \sum_i \varphi_i F_i(\alpha) \leq \prod_i F_i(\alpha) \text{ and } F_i(\alpha) > 0 \text{ for all } i \} \quad (6)$$

and denote by p^* the threshold mechanism with threshold α^* .¹⁰ Given the parameters φ_i , the designer would like to choose the threshold as low as possible: if all reports are below

¹⁰The constraint set is nonempty, and hence α^* is well-defined, because given $\sum_i \varphi_i \leq 1$, $\alpha = \max_i \bar{t}_i$ satisfies all constraints. Note also that, by continuity of F_i and the definition of α^* , $\sum_i \varphi_i F_i(\alpha^*) = \prod_i F_i(\alpha^*)$.

the threshold, he has to randomize the allocation, whereas he can allocate to the highest type if at least one report is above the threshold. However, feasibility requires that the threshold cannot be too low. We prove below that, basically, α^* is the lowest threshold that is still feasible.

The following theorem from BDL is the main step in deriving the optimal mechanism.

Theorem 2. [Theorem 4 in Ben-Porath et al. (2014)] *The threshold mechanism p^* is the essentially unique solution to problem (R), that is, every solution to this problem equals p^* almost everywhere.*

The proof of Theorem 2 consists of three steps. In Step 1, we show that \hat{p}^* is a feasible solution by constructing an ex-post allocation rule inducing the required interim allocation rule and arguing that it satisfies the incentive constraints. In Step 2, we show that \hat{p}^* is indeed an optimal solution to the optimization problem. To do so, we first derive an upper bound on the objective function using that any feasible solution must satisfy Border's constraints. We then show that \hat{p}^* achieves this bound. In Step 3, we show that every optimal reduced form must in fact be equal to \hat{p}^* almost everywhere.

Proof.

Step 1: Feasibility

We will first construct a feasible ex-post rule inducing the reduced form rule \hat{p}^* and then show that \hat{p}^* satisfies the incentive constraints.

Consider the following ex-post rule p^* . It allocates the object to the agent with the highest type whenever $t_j > \alpha^*$ for some j , and whenever $t_j \leq \alpha^*$ for all j it is defined by $p_i^*(t) = \frac{\varphi_i}{\prod_{j \neq i} F_j(\alpha^*)}$.¹¹ This rule induces the interim rule \hat{p}^* . Moreover, it is clearly feasible if $t_j > \alpha^*$ for some j . Assuming $t_j \leq \alpha^*$ for all j and summing over all agents, we have that $\sum_i p_i^*(t) = \sum_i \frac{\varphi_i}{\prod_{j \neq i} F_j(\alpha^*)}$. By definition of α^* and continuity of the $F_j(\cdot)$'s, $\sum_i \frac{\varphi_i}{\prod_{j \neq i} F_j(\alpha^*)} \leq 1$.¹² Thus, p^* is a feasible ex-post rule.

Regarding the incentive constraints, $\hat{p}_i^*(t_i) = \varphi_i$ for all $t_i \leq \alpha^*$. Suppose instead that $t_i > \alpha^*$. By definition of α^* , $F_i(t_i) > 0$ and we obtain $\hat{p}_i^*(t_i) = \frac{\prod_j F_j(t_i)}{F_i(t_i)}$. We will show below that $\sum_j \varphi_j F_j(t_i) - \prod_j F_j(t_i)$ is non-increasing for all $t_i \geq \alpha^*$ and hence $\prod_j F_j(t_i) \geq \sum_j \varphi_j F_j(t_i)$ for all $t_i \geq \alpha^*$. Thus,

$$\hat{p}_i^*(t_i) = \frac{\prod_j F_j(t_i)}{F_i(t_i)} \geq \frac{\sum_j \varphi_j F_j(t_i)}{F_i(t_i)} \geq \varphi_i.$$

Hence, \hat{p}^* is a feasible solution to (R). To finalize the argument we now show that the function $h(x) := \sum_i \varphi_i F_i(x) - \prod_i F_i(x)$ is non-increasing for all $x > \alpha^*$. Note first that $h(\alpha^*) = 0$; if this was strictly negative, $F_i(\alpha^*) > 0$ had to hold for all i and we could decrease α^* without violating any of the constraints in (6). Differentiating h , $h'(x) = \sum_i f_i(x)[\varphi_i - \prod_{j \neq i} F_j(x)] \leq \sum_i f_i(x)[\varphi_i - \prod_{j \neq i} F_j(\alpha^*)] \leq \sum_i \frac{f_i(x)}{F_i(\alpha^*)} [\sum_i \varphi_i F_i(\alpha^*) -$

¹¹If $F_j(\alpha^*) = 0$ for some j , we define the ex-post rule to always allocate to the agent with the highest type.

¹²For all $\alpha > \alpha^*$, we have $\sum_i \varphi_i F_i(\alpha) \leq \prod_i F_i(\alpha)$. Rearranging gives $\sum_i \frac{\varphi_i}{\prod_{j \neq i} F_j(\alpha)} \leq 1$. By continuity of F_i , this inequality holds for α^* as well.

$\prod_i F_i(\alpha^*) = 0$, since $h(\alpha^*) = 0$. Thus, $h'(x) \leq 0$ for all $x > \alpha^*$ and as desired h is non-increasing for all $x > \alpha^*$.

Step 2: Optimality

We first establish an upper bound for the objective function and then show that the reduced form \hat{p}^* achieves this upper bound.

Let \tilde{p} be any feasible reduced form, which therefore satisfies the Border conditions for all $\alpha \in \mathbb{R}$:

$$\sum_i \int_{\alpha}^{\bar{t}_i} f_i(t_i) \tilde{p}_i(t_i) dt_i \leq 1 - \prod_i F_i(\alpha). \quad (7)$$

Since $\tilde{p}_i(t_i) \geq \varphi_i$, the Border conditions also imply that for all α ,

$$\sum_i \int_{\underline{t}_i}^{\alpha} \varphi_i f_i(t_i) dt_i + \sum_i \int_{\alpha}^{\bar{t}_i} f_i(t_i) \tilde{p}_i(t_i) dt_i \leq 1,$$

or, equivalently,

$$\sum_i \int_{\alpha}^{\bar{t}_i} f_i(t_i) \tilde{p}_i(t_i) dt_i \leq 1 - \sum_i \varphi_i F_i(\alpha). \quad (8)$$

Note that if $\underline{t}_i < 0$, $\tilde{p}_i(t_i) \geq \varphi_i$ implies $\int_{\underline{t}_i}^0 f_i(t_i) \tilde{p}_i(t_i) t_i dt_i \leq \int_{\underline{t}_i}^0 f_i(t_i) \varphi_i t_i dt_i$. Moreover, denoting $\bar{t} = \max_i \{\bar{t}_i\}$ we get:

$$\begin{aligned} & \sum_i \int_0^{\bar{t}_i} f_i(t_i) \tilde{p}_i(t_i) t_i dt_i \\ &= \sum_i t_i \int_0^{t_i} f_i(s) \tilde{p}_i(s) ds \Big|_{t_i=0}^{\bar{t}_i} - \sum_i \int_0^{\bar{t}_i} \int_0^{t_i} f_i(s) \tilde{p}_i(s) ds dt_i \\ &= \int_0^{\bar{t}} \sum_i \int_{\alpha}^{\bar{t}_i} f_i(s) \tilde{p}_i(s) ds d\alpha \\ &\leq \int_0^{\alpha^*} [1 - \sum_i \varphi_i F_i(\alpha)] d\alpha + \int_{\alpha^*}^{\bar{t}} [1 - \prod_i F_i(\alpha)] d\alpha, \end{aligned} \quad (9)$$

where the first equality follows from integration by parts, the second by rearranging terms and the inequality follows from (7) and (8). This gives the following upper bound on the objective function for any feasible solution \tilde{p} ,

$$\begin{aligned} & \sum_i \int_{\underline{t}_i}^{\bar{t}_i} f_i(t_i) \tilde{p}_i(t_i) t_i dt_i \\ &\leq \sum_i \int_{\underline{t}_i}^0 [f_i(\alpha) \varphi_i \alpha] d\alpha + \int_0^{\alpha^*} [1 - \sum_i \varphi_i F_i(\alpha)] d\alpha + \int_{\alpha^*}^{\bar{t}} [1 - \prod_i F_i(\alpha)] d\alpha. \end{aligned}$$

We claim that \hat{p}^* achieves this upper bound because inequality (9) is binding for \hat{p}^* . First, for $\alpha \geq \alpha^*$, $\sum_i \int_{\alpha}^{\bar{t}_i} f_i(s) \hat{p}_i^*(s) ds = \sum_i \int_{\alpha}^{\bar{t}_i} f_i(s) \prod_{j \neq i} F_j(s) ds = 1 - \prod_i F_i(\alpha)$. Moreover, for $\alpha < \alpha^*$,

$$\begin{aligned} \sum_i \int_{\alpha}^{\bar{t}_i} f_i(s) \hat{p}_i^*(s) ds &= \sum_i \int_{\alpha}^{\alpha^*} f_i(s) \varphi_i ds + 1 - \prod_i F_i(\alpha^*) \\ &= \sum_i \varphi_i [F_i(\alpha^*) - F_i(\alpha)] + 1 - \prod_i F_i(\alpha^*) = 1 - \sum_i \varphi_i F_i(\alpha) \end{aligned}$$

since, by definition of α^* , $\sum_i \varphi_i F_i(\alpha^*) = \prod_i F_i(\alpha^*)$. Summarizing, we get

$$\begin{aligned} &\sum_i \int_{t_i}^{\bar{t}_i} f_i(t_i) \hat{p}_i(t_i) t_i dt_i \\ &= \sum_i \int_{t_i}^0 [f_i(\alpha) \varphi_i \alpha] d\alpha + \int_0^{\alpha^*} [1 - \sum_i \varphi_i F_i(\alpha)] d\alpha + \int_{\alpha^*}^{\bar{t}} [1 - \prod_i F_i(\alpha)] d\alpha, \end{aligned}$$

which shows that \hat{p}^* is indeed an optimal solution.

Step 3: Uniqueness

Note that any feasible reduced form \tilde{p} satisfies the following inequality:

$$G(\alpha_1, \dots, \alpha_n) := \sum_i \int_{\alpha_i}^{\bar{t}_i} f_i(s) \tilde{p}_i(s) ds \leq 1 - \prod_i F_i(\alpha_i) =: H(\alpha_1, \dots, \alpha_n).$$

Since G is monotone, it is differentiable almost everywhere, and H is differentiable by assumption. Because \hat{p}^* satisfies inequality (9) as an equality, we conclude that any optimal reduced form must satisfy (9) as an equality as well. Hence, any optimal reduced form must satisfy, for almost every $\alpha \geq \alpha^*$, $G(\alpha, \dots, \alpha) = H(\alpha, \dots, \alpha)$ and that G and H are differentiable in α_i for all i at (α, \dots, α) . Since H is an upper bound for G , this implies that their derivatives must coincide at (α, \dots, α) :

$$-\tilde{p}_i(\alpha) f_i(\alpha) = - \prod_{j \neq i} F_j(\alpha) f_i(\alpha).$$

Moreover, by (7) and $\tilde{p}_i \geq \varphi_i$, $\tilde{p}_i(t_i) = \varphi_i$ for $t_i < \alpha^*$. We conclude that \tilde{p} equals \hat{p}^* almost everywhere. \square

References

- Ben-Porath, E., E. Dekel, and B. L. Lipman (2014). Optimal allocation with costly verification. *American Economic Review* 104, 3779–3813.
- Border, K. C. (1991). Implementation of reduced form auctions: A geometric approach. *Econometrica* 59, 1175–1187.
- Che, Y.-K., J. Kim, and K. Mierendorff (2013). Generalized reduced-form auctions: A network-flow approach. *Econometrica* 81, 2487–2520.

- Gershkov, A., J. K. Goeree, A. Kushnir, B. Moldovanu, and X. Shi (2013). On the equivalence of bayesian and dominant strategy implementation. *Econometrica*.
- Lipman, B. L. (2015). An elementary proof of the optimality of threshold mechanisms. *Working paper*.
- Manelli, A. M. and D. R. Vincent (2010). Bayesian and dominant-strategy implementation in the independent private-values model. *Econometrica*, 1905–1938.
- Mierendorff, K. (2011). Asymmetric reduced form auctions. *Economics Letters* 110.
- Mierendorff, K. (2016). Optimal dynamic mechanism design with deadlines. *The Journal of Economic Theory* 161.
- Mylovanov, T. and A. Zapechelnuk (2015). Mechanism design with ex-post verification and limited punishments. *Working paper*.
- Pai, M. and R. Vohra (2014). Optimal auctions with financially constrained bidders. *Journal of Economic Theory* 150, 383–425.