



A note on optimal allocation with costly verification[☆]

Albin Erlanson^{a,*}, Andreas Kleiner^b

^a Department of Economics, Stockholm School of Economics, Sweden

^b Department of Economics, W.P. Carey School of Business, Arizona State University, United States of America

ARTICLE INFO

Article history:

Received 5 July 2018

Received in revised form 11 May 2019

Accepted 24 June 2019

Available online 15 July 2019

Keywords:

Optimal mechanisms

Costly verification

BIC and EPIC equivalence

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 ex-post incentive compatible mechanisms. We also provide a simple proof showing that the optimal mechanism is a threshold mechanism.

© 2019 Elsevier B.V. All rights reserved.

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 he was truthful, and the object is randomly allocated according to a given probability distribution otherwise. If an agent is caught lying he does not receive the object.

BDL point out that, somewhat surprisingly, the optimal mechanism is ex-post 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 ex-post equilibrium and induces the same expected verification costs. Proving the equivalence might help to analyze this or related models because it allows to optimize over a smaller class of ex-post incentive compatible mechanisms. Since these mechanisms are conceptually simpler the corresponding analysis might sometimes be more tractable. Similar equivalence results exist in the standard one-dimensional mechanism design setting with single-crossing and quasi-linear utilities (Mookherjee and Reichelstein, 1992; 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 interim allocation rules.³ To prove the optimality of threshold mechanisms we observe that the relevant incentive constraints are formulated in terms of interim allocation and verification rules. Thus, we can restate the optimization problem using only interim rules and optimize over them directly, which is significantly easier. A characterization of feasible interim allocation rules is readily available due to Border (1991), and we can show that threshold mechanisms are optimal. Our approach of using interim allocation rules to solve for optimal mechanisms is one example among several recent papers, for other examples see Mierendorff (2016), Mylovanov and Zapechelnuyk (2017) and Pai and Vohra (2014).

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

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

[☆] We are grateful to the Editor, an Associate Editor and a referee for helpful suggestions and comments. We would also like to thank 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. Andreas Kleiner gratefully acknowledges financial support from SFB-TR 15.

* Corresponding author.

E-mail addresses: albin.erlanson@hhs.se (A. Erlanson), andreas.kleiner@asu.edu (A. Kleiner).

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

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 ex-post incentive compatible mechanisms. In Section 4 we provide our alternative proof for the optimality of threshold mechanisms.

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 $t_i \in \mathcal{T}_i \equiv [t_i, \bar{t}_i]$, where $-\infty < t_i < \bar{t}_i < \infty$. 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 will consider both ex-post and Bayesian incentive compatibility.

Definition 1. A mechanism (p, q) is Bayesian incentive compatible (BIC) if, for all $i \in \mathcal{I}$ and all $t_i, t'_i \in \mathcal{T}_i$

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

The right hand side of Eq. (1) is the probability of being allocated the object $\mathbb{E}_{t_{-i}}[p_i(t'_i, t_{-i})]$ conditional upon not being verified minus the probability of being verified $\mathbb{E}_{t_{-i}}[q_i(t'_i, t_{-i})]$. This is the probability of getting an object when an agent of type t_i reports to be of type $t'_i \neq t_i$, since once verified agent i does not get the object. For (p, q) to be incentive compatible this has to be lower than the probability of getting an object when reporting the type t_i truthfully.

Since the right-hand side of the inequality is independent of t_i , this inequality holds if and only if it holds for the type that minimizes the left-hand side, for which the left-hand side assumes the value $\inf_{t'_i \in \mathcal{T}_i} \mathbb{E}_{t_{-i}}[p_i(t'_i, t_{-i})]$. This characterizes BIC mechanisms by giving us a lower bound on how often we need to check an agent with type t_i to make an allocation rule p incentive compatible.

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

$$\mathbb{E}_{t_{-i}}[q_i(t_i, t_{-i})] \geq \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})] - \inf_{t'_i \in \mathcal{T}_i} \mathbb{E}_{t_{-i}}[p_i(t'_i, t_{-i})]. \tag{2}$$

The stronger notion of ex-post incentive compatibility requires that Eq. (1) holds pointwise and we get the analogous characterization of ex-post incentive compatibility.

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

$$q_i(t_i, t_{-i}) \geq p_i(t_i, t_{-i}) - \inf_{t'_i \in \mathcal{T}_i} p_i(t'_i, t_{-i}). \tag{3}$$

In this note, instead of using the ex-post allocation rule p , we will use the corresponding *interim allocation rule* (also called reduced form) \hat{p} , where $\hat{p}_i(t_i) = \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$ and the interim verification rule $\hat{q}_i(t_i) = \mathbb{E}_{t_{-i}}[q_i(t_i, t_{-i})]$. These are lower dimensional objects and are simpler to maximize over than ex-post rules. Moving from ex-post allocation rules to interim allocation rules we need to know whether a given interim allocation rule 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 interim allocation rules: a monotone interim allocation rule is feasible if and only if, for all $(\alpha_1, \dots, \alpha_n) \in \mathcal{T}$,

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

This condition is necessary for an interim allocation rule 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.⁵ The content of Border's theorem is to show that the above condition is also sufficient for a nondecreasing interim allocation rule to be feasible.

3. BIC-EPIC equivalence

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

BDL showed that the optimal BIC mechanism satisfies the stronger notion of ex-post incentive compatibility.⁶ 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” EPIC mechanism (\bar{p}, \bar{q}) that induces the same expected verification costs and the same interim allocation rules. While this equivalence holds in this model for any mechanism, it fails in model variations with imperfect verification, interdependent preferences, or more general decision problems (see Erlanson and Kleiner (2019)).

To understand the logic behind the equivalence and to compare Bayesian with ex-post incentive constraints, we begin with averaging the pointwise constraints in (3) by taking expectations over reports of the other agents. Consequently, any EPIC mechanism must satisfy

$$\mathbb{E}_{t_{-i}}[q_i(t_i, t_{-i})] \geq \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})] - \mathbb{E}_{t_{-i}}[\inf_{t'_i \in \mathcal{T}_i} p_i(t'_i, t_{-i})]; \tag{4}$$

moreover, for any allocation rule p we can define q by $q_i(t_i, t_{-i}) = p_i(t_i, t_{-i}) - \inf_{t'_i \in \mathcal{T}_i} p_i(t'_i, t_{-i})$ so that (p, q) is EPIC and satisfies (4) as an equality.

Note that $\inf_{t'_i \in \mathcal{T}_i} \mathbb{E}_{t_{-i}}[p_i(t'_i, t_{-i})] \geq \mathbb{E}_{t_{-i}}[\inf_{t'_i \in \mathcal{T}_i} p_i(t'_i, t_{-i})]$, and hence the right-hand side of (4) is in general larger than the right-hand side of the Bayesian constraints in (2). Thus, for a

⁵ We use this direction to construct an upper bound on the objective function in Section 4.

⁶ BDL show that the optimal mechanism is dominant-strategy incentive compatible (DIC). Our equivalence result below extends to DIC instead of EPIC if we specify more precisely the allocation if an agent is found lying, see BDL.

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

given allocation rule p that is BIC with the verification rule q , we may need to verify agents more frequently to satisfy the more stringent ex-post incentive constraints. Therefore, in general it is more costly to implement an allocation rule in ex-post equilibrium than in a Bayes Nash equilibrium; only if the expectation operator commutes with the infimum operator, that is if $\inf_{t'_i \in \mathcal{T}_i} \mathbb{E}_{t_{-i}}[p_i(t'_i, t_{-i})] = \mathbb{E}_{t_{-i}}[\inf_{t'_i \in \mathcal{T}_i} p_i(t'_i, t_{-i})]$, can a rule be implemented at the same verification costs in ex-post equilibrium. **Example 1** illustrates an allocation rule p which requires strictly more verifications to be implemented in ex-post equilibrium than in Bayes Nash equilibrium.

Example 1. Suppose that $\mathcal{I} = \{1, 2\}$ and that each type profile is equally likely and consider the incentive constraints for agent 2. The allocation rule p_2 for agent 2 is shown in Fig. 1(a). We can think of agent 2 as having three intervals to report in, T_2^L, T_2^I and T_2^H , since the allocation rule is the same for any report in these three intervals. We first compute the interim allocation rule as $\mathbb{E}_{t_1}[p_2(t_1, t_2)] = 0.8$ for $t_2 \in T_2^L$, $\mathbb{E}_{t_1}[p_2(t_1, t_2)] = 0.2$ for $t_2 \in T_2^I$, and $\mathbb{E}_{t_1}[p_2(t_1, t_2)] = 0.4$ for $t_2 \in T_2^H$. To find a verification rule q_2 such that p_2 and q_2 satisfy BIC, we observe that $\inf_{t_2 \in T_2} \mathbb{E}_{t_1}[p_2(t_1, t_2)] = 0.2$. By Lemma 1, BIC is satisfied if q_2 satisfies $\mathbb{E}_{t_1}[q_2(t_1, t_2)] = 0.6$ for $t_2 \in T_2^L$, $\mathbb{E}_{t_1}[q_2(t_1, t_2)] = 0$ for $t_2 \in T_2^I$, and $\mathbb{E}_{t_1}[q_2(t_1, t_2)] = 0.2$ for $t_2 \in T_2^H$.

If we instead want to satisfy the EPIC constraints, the expected number of verifications and hence the cost of verification increases. To see this, note that $\inf_{t_2} p_2(t_1, t_2) = 0.2$ for $t_1 \in T_1^L$, $\inf_{t_2} p_2(t_1, t_2) = 0$ for $t_1 \in T_1^I$, and $\inf_{t_2} p_2(t_1, t_2) = 0.1$ for $t_1 \in T_1^H$. Therefore, $\mathbb{E}_{t_1}[\inf_{t_2} p_2(t_1, t_2)] = 0.1 < \inf_{t_2} \mathbb{E}_{t_1}[p_2(t_1, t_2)]$. Inequality (4) then implies that $\mathbb{E}_{t_1}[q_2(t_1, t_2)] \geq 0.8 - 0.1$ for $t_2 \in T_2^L$, $\mathbb{E}_{t_1}[q_2(t_1, t_2)] = 0.2 - 0.1$ for $t_2 \in T_2^I$, and $\mathbb{E}_{t_1}[q_2(t_1, t_2)] = 0.4 - 0.1$ for $t_2 \in T_2^H$ (the verification probabilities for each profile of reports are given in Fig. 1(b)).

Example 1 shows that for some allocation rules it is strictly more expensive to implement them in ex-post equilibrium than in Bayes Nash equilibrium. It is therefore not obvious why the optimal mechanism in the class of BIC mechanisms is also EPIC. However, we show below that for any allocation rule p there exists another allocation rule \tilde{p} , which is equivalent in the sense that it has the same interim allocation rules and therefore induces the same expected utilities, such that the expectation operator commutes with the infimum operator. It follows from our arguments leading to (4) that this rule can be implemented in ex-post equilibrium using the same expected verification probabilities, and hence the same expected verification costs, that are necessary to implement it in Bayes Nash equilibrium.

Example 1 (Ctd.). Fig. 2(a) illustrates an allocation rule \tilde{p}_2 that is equivalent to p_2 and for which the expectation operator and the infimum operator commute. It is readily verified that $\mathbb{E}_{t_1}[\tilde{p}_2(t_1, t_2)] = \mathbb{E}_{t_1}[p_2(t_1, t_2)]$ for all t_2 . Moreover, $\inf_{t_2} \tilde{p}_2(t_1, t_2) = 0.3$ for $t_1 \in T_1^L$, $\inf_{t_2} \tilde{p}_2(t_1, t_2) = 0.3$ for $t_1 \in T_1^I$, and $\inf_{t_2} \tilde{p}_2(t_1, t_2) = 0$ for $t_1 \in T_1^H$. Therefore, $\mathbb{E}_{t_1}[\inf_{t_2} \tilde{p}_2(t_1, t_2)] = 0.2 = \inf_{t_2} \mathbb{E}_{t_1}[\tilde{p}_2(t_1, t_2)]$ and we can implement \tilde{p} with EPIC using the same expected number of verifications that are necessary for BIC. Fig. 2(b) shows the verification rule that achieves this.

To state our result, we define our equivalence notion formally.

Definition 2. 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 .

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

To explain our proof approach, suppose that the type space of each agent is finite. We can then reorder the types so that $\mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$ is nondecreasing in t_i for each i . It then follows from Theorem 1 in Gershkov et al. (2013) that there exists an equivalent allocation rule \tilde{p} such that $\tilde{p}_i(t_i, t_{-i})$ is nondecreasing in t_i for each t_{-i} and i . It follows that $\inf_{t_i} \tilde{p}_i(t_i, t_{-i})$ is attained, for any choice of t_{-i} , for the smallest reordered type, denoted by t_i^* . It follows that

$$\mathbb{E}_{t_{-i}}[\inf_{t_i} \tilde{p}_i(t_i, t_{-i})] = \mathbb{E}_{t_{-i}}[\tilde{p}_i(t_i^*, t_{-i})] = \inf_{t_i} \mathbb{E}_{t_{-i}}[\tilde{p}_i(t_i, t_{-i})];$$

that is, the expectation operator and the infimum operator commute for \tilde{p} . Step 1 below provides the argument for the case of infinite type spaces.

We can then define in the second step verification rules for \tilde{p} that make (\tilde{p}, \tilde{q}) EPIC while satisfying $\mathbb{E}_{t_{-i}}[\tilde{q}_i(t_i, t_{-i})] = \mathbb{E}_{t_{-i}}[q_i(t_i, t_{-i})]$.

Proof of Theorem 1. Let (p, q) be a BIC mechanism. We will construct an equivalent EPIC mechanism (\tilde{p}, \tilde{q}) .

Step 1: Constructing an equivalent allocation rule \tilde{p} such that expectation and infimum commute

We will define a new type space and an allocation rule p' that is nondecreasing in its marginals on this new type space. Given this, we can construct another allocation rule p'' that is pointwise nondecreasing. Finally we will construct the new allocation rule \tilde{p} using p'' and the relation between the original and the new type space.

Let, for all $i \in \mathcal{I}$, $\sigma_i : \mathcal{T}_i \rightarrow \mathbb{R}$ be defined as $\sigma_i(t_i) = \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$ and let the new type space be $\tilde{\mathcal{T}}_i = \{x \in \mathbb{R} \mid \sigma_i(t_i) = x \text{ for some } t_i \in \mathcal{T}_i\}$, 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 $\mathbb{E}_{t_{-i}}[p'_i(\tilde{t}_i, \tilde{t}_{-i})] = \tilde{t}_i$ for all $\tilde{t}_i \in \tilde{\mathcal{T}}_i$, recall that by construction $\tilde{t}_i = \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$ for some $t_i \in \mathcal{T}_i$. The existence of such an allocation rule p' is guaranteed by a characterization of interim allocation rules (Mierendorff, 2011).⁷ Thus, the interim allocation rules of p'_i are nondecreasing and Theorem 1 in Gershkov et al. (2013) implies that there 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

$$\tilde{p}_i(t_i, t_{-i}) = p''_i(\sigma_i(t_i), \sigma_{-i}(t_{-i})) \text{ for all } i \in \mathcal{I}. \tag{5}$$

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, we get that the interim allocation rules are the same for \tilde{p} and p , that is $\mathbb{E}_{t_{-i}}[\tilde{p}_i(t_i, t_{-i})] = \mathbb{E}_{t_{-i}}[p_i(t_i, t_{-i})]$, as desired.

⁷ Let μ_i denote the measure induced by the cdf. F_i , we get for all $\alpha_i \in \tilde{\mathcal{T}}_i$ (setting $A_i = \{\tilde{t}_i \in \tilde{\mathcal{T}}_i \mid \tilde{t}_i \geq \alpha_i\}$):

$$\begin{aligned} \sum_i \int_{A_i} \tilde{t}_i dG_i(\tilde{t}_i) &= \sum_i \int_{\sigma^{-1}(A_i)} \mathbb{E}_{t_{-i}} p_i(t_i, t_{-i}) dF_i(t_i) \\ &\leq 1 - \prod_i [1 - \mu_i(\sigma^{-1}(A_i))] = 1 - \prod_i G_i(\alpha_i) \end{aligned}$$

The equalities hold by construction and the inequality holds because p is a feasible interim allocation rule.

⁸ Gershkov 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).

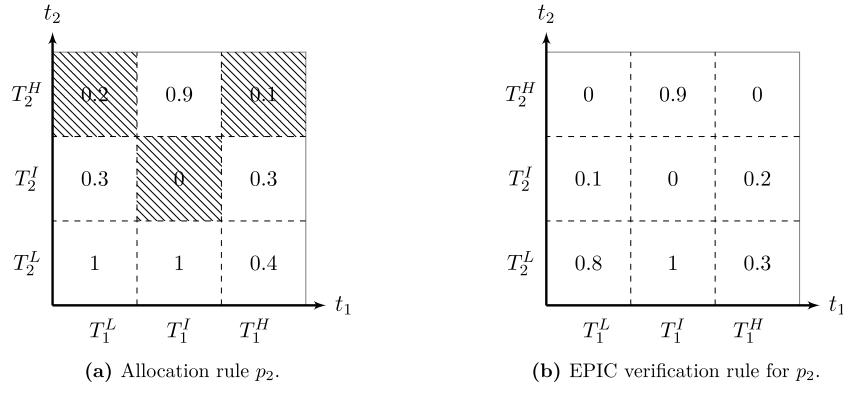


Fig. 1. BIC-EPIC non-equivalence.

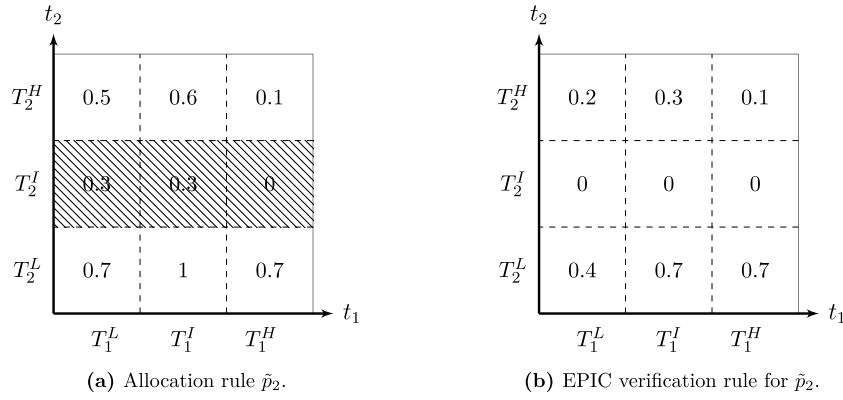


Fig. 2. BIC-EPIC equivalence.

Because p_i'' is pointwise nondecreasing, we obtain that the expectation operator and the infimum operator commute. Formally, for arbitrary $\varepsilon > 0$,

$$\forall \tilde{t}_{-i} \in \tilde{\mathcal{T}}_{-i} : \exists \tilde{t}_i^* \in \tilde{\mathcal{T}}_i : p_i''(\tilde{t}_i(\tilde{t}_{-i}), \tilde{t}_{-i}) \leq \inf_{\tilde{t}_i} p_i''(\tilde{t}_i, \tilde{t}_{-i}) + \varepsilon/2.$$

Therefore, if we choose $\tilde{t}_i^* \in \tilde{\mathcal{T}}_i$ such that $\Pr_{\tilde{t}_{-i}}(\tilde{t}_i(\tilde{t}_{-i}) \geq \tilde{t}_i^*) \geq 1 - \varepsilon/2$ then the fact that p_i'' is nondecreasing implies

$$\inf_{\tilde{t}_i} \mathbb{E}_{\tilde{t}_{-i}}[p_i''(\tilde{t}_i, \tilde{t}_{-i})] \leq \mathbb{E}_{\tilde{t}_{-i}}[p_i''(\tilde{t}_i^*, \tilde{t}_{-i})] \leq \mathbb{E}_{\tilde{t}_{-i}}[\inf_{\tilde{t}_i} p_i''(\tilde{t}_i, \tilde{t}_{-i})] + \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, we conclude that $\inf_{\tilde{t}_i} \mathbb{E}_{\tilde{t}_{-i}}[p_i''(\tilde{t}_i, \tilde{t}_{-i})] = \mathbb{E}_{\tilde{t}_{-i}}[\inf_{\tilde{t}_i} p_i''(\tilde{t}_i, \tilde{t}_{-i})]$.

Further by definition of \tilde{p}_i we have that $\tilde{p}_i(t) = p_i''(\sigma(t))$. Thus,

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

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 EPIC, i.e., satisfies (3), and induces the same expected verification costs. The verification rules are defined as

$$\tilde{q}_i(t_i, t_{-i}) = \tilde{p}_i(t_i, t_{-i}) - \inf_{t'_i \in \mathcal{T}_i} \tilde{p}_i(t'_i, t_{-i}) \text{ for all } i \in \mathcal{I}. \quad (6)$$

By construction, the incentive constraints in (3) hold as equalities, and it remains to show that the expected verification probabilities are the same.

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

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

The first equality follows by definition of \tilde{q}_i . The second equality follows from Step 1 and the third by construction of \tilde{p}_i . 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 analyze which mechanism maximizes the expected payoff for the principal. We provide a simple proof showing that it is optimal to use a threshold mechanism.

The principal's problem can be stated using interim allocation rules as

$$\begin{aligned} \max_{0 \leq \hat{p}, \hat{q} \leq 1} & \sum_i \mathbb{E}_{t_i} [t_i \hat{p}_i(t_i) - c_i \hat{q}_i(t_i)] \\ \text{s.t.} & \inf_{t'_i \in \mathcal{T}_i} \hat{p}_i(t'_i) \geq \hat{p}_i(t_i) - \hat{q}_i(t_i) \\ & \hat{p} \text{ is feasible} \end{aligned}$$

The objective function is the expected payoff for the principal, taking into account the expected costs of verification. The first set of constraints are the incentive compatibility constraints, and the second constraint ensures that the interim allocation rules are actually implementable.

Note that the objective function is decreasing in \hat{q} and any optimal solution therefore satisfies $\hat{q}_i(t_i) = \hat{p}_i(t_i) - \inf_{t'_i \in \mathcal{T}_i} \hat{p}_i(t'_i)$. We can then restate the problem as a simpler problem where we only optimize over interim allocation rules \hat{p} and use the

characterization of feasible interim allocation rules:

$$\begin{aligned} \max_{0 \leq \hat{p} \leq 1} \sum_i \mathbb{E}_{t_i} [(t_i - c_i) \hat{p}_i(t_i) + c_i \inf_{t'_i \in \mathcal{T}_i} \hat{p}_i(t'_i)] \\ \text{s.t. (Border)} \end{aligned}$$

While this problem is not a linear program, we can formulate a related program that is linear by replacing $\inf_{t'_i \in \mathcal{T}_i} \hat{p}_i(t'_i)$ by φ_i :

$$\begin{aligned} \max_{\hat{p}, \varphi} \sum_i \mathbb{E}_{t_i} [(t_i - c_i) \hat{p}_i(t_i) + c_i \varphi_i] \\ \text{s.t. } \hat{p}_i(t_i) \geq \varphi_i \\ \text{(Border)} \end{aligned}$$

Indeed, if (\hat{p}, φ) solves the latter problem, $\varphi_i = \inf_{t'_i \in \mathcal{T}_i} \hat{p}_i(t'_i)$ must hold as we could increase φ_i otherwise to obtain a higher objective value. Therefore, \hat{p} obtains the same objective value in the former problem. Conversely, if \hat{p} solves the former problem we can define $\varphi_i := \inf_{t'_i \in \mathcal{T}_i} \hat{p}_i(t'_i)$ and then (\hat{p}, φ) will obtain the same value in the latter problem, establishing that both problems achieve the same value.

The last step before solving the problem is to consider net-types $t_i - c_i$ instead of the original types t_i . This gives us an equivalent formulation of the optimization problem that is notationally easier to work with. Henceforth, with slight abuse of notation, the type t_i is the net-type $t_i - c_i$.⁹

Below, we show that a so-called threshold mechanism is optimal. To do so, we show that for any choice of $\{\varphi_i\}_i$ a threshold mechanism is optimal in the above problem. We therefore analyze the following parametrized problem¹⁰: Given $\{\varphi_i\}_i$ such that $\sum_i \varphi_i \leq 1$,

$$\begin{aligned} \max_{\hat{p}} \sum_i \mathbb{E}_{t_i} [t_i \hat{p}_i(t_i)] \\ \text{s.t. } \hat{p}_i(t_i) \geq \varphi_i \\ \text{(Border)} \end{aligned} \tag{R}$$

Following BDL, we define a *threshold mechanism* with threshold α to be a mechanism p with the following interim allocation rule: $\hat{p}_i(t_i) = \prod_{j \neq i} F_j(t_i)$ for $t_i > \alpha$ and $\hat{p}_i(t_i) = \varphi_i$ otherwise. This interim allocation rule 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

$$\alpha^* = \inf\{\alpha \in \mathbb{R}_+ \mid \sum_i \varphi_i F_i(\alpha) \leq \prod_i F_i(\alpha) \text{ and } F_i(\alpha) > 0 \text{ for all } i\} \tag{7}$$

and denote by p^* the threshold mechanism with threshold α^* .¹¹ Given the parameters φ_i , the designer would like to choose the threshold as low as possible: if all reports are below 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

⁹ Formally, we could define $\tilde{t}_i = t_i - c_i$ and write the whole problem in terms of \tilde{t}_i . In particular, this would change the distribution of types from F_i to another distribution \tilde{F}_i of net types $t_i - c_i$. Since a threshold mechanism is optimal independent of the particular distribution, we do not formalize this step. To simplify notation, we use t_i in place of \tilde{t}_i .

¹⁰ Since φ_i 's are constant we can drop any terms that only involve φ_i 's from the objective function.

¹¹ The constraint set is nonempty, and hence α^* is well-defined, because given $\sum_i \varphi_i \leq 1$, $\alpha = \max_i \tilde{t}_i$ satisfies all constraints. Note also that, by continuity of F_i and the definition of α^* , $\sum_i \varphi_i F_i(\alpha^*) = \prod_i F_i(\alpha^*)$.

low. We prove below that 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.*

Our proof is significantly shorter than the original proof provided in BDL, provides some insight for the economic properties of the optimal mechanism, and suggests an approach that might be applicable in related allocation problems. First, we state and prove the result using only interim allocation rules. While BDL use interim allocation rules in some places, for example to state the incentive constraints, they switch back and forth between interim allocation rules and ex-post allocation rules, which adds complexity. Using Border's characterization of interim allocation rules allows us to use exclusively interim allocation rules, which are simpler to analyze. Our second key insight is that the optimization problem, written using interim allocation rules, has a tractable structure. To illustrate this, suppose that each type space is finite. The optimization problem is then a linear program and the constraint set is a generalized polymatroid, and it is well-known that such a problem can be solved by a greedy algorithm.¹² A greedy algorithm makes locally optimal choices, not taking into account the effect this has on the feasible choices for other types. For our linear program this implies that the allocation probability for the highest net type is set as high as possible, then the allocation probability for the next highest net type is set as high as possible given the previous choice and so on. This explains the particular economic structure of the optimal allocation: the threshold mechanism allocates the object to the agent with the highest type as long as someone reports above the threshold. The threshold is chosen so that this mechanism still satisfies the constraint $\hat{p}_i(t_i) \geq \varphi_i$.

Since the type spaces in our model are infinite, the constraint set is the infinite-dimensional analog of a polymatroid. Our proof of Theorem 2 therefore follows the outline of arguments showing that a greedy algorithm solves the finite-dimensional linear program, but adapts these arguments to infinite types. More specifically, we show in Step 1 that no feasible mechanism can achieve a higher objective value than the threshold mechanism \hat{p}^* . 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 2, we show that every optimal interim allocation rule must in fact be equal to \hat{p}^* almost everywhere. Finally, In Step 3 we verify 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.

Proof of Theorem 2.

Step 1: Optimality

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

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

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

¹² See, for example, Che et al. (2013).

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

$$\sum_i \int_{\underline{t}_i}^{\alpha'} \varphi_i f_i(t_i) dt_i + \sum_i \int_{\alpha'}^{\bar{t}_i} f_i(t_i) \tilde{p}_i(t_i) dt_i \leq 1 - \prod_i F_i(\bar{t}_i) = 1$$

or, equivalently,

$$\sum_i \int_{\alpha'}^{\bar{t}_i} f_i(t_i) \tilde{p}_i(t_i) dt_i \leq 1 - \sum_i \varphi_i F_i(\alpha') \text{ for all } \alpha' \in \mathbb{R}. \quad (9)$$

Denoting $\bar{t} = \max_i \{\bar{t}_i\}$, we therefore get:

$$\begin{aligned} & \sum_i \int_0^{\bar{t}_i} f_i(t_i) \tilde{p}_i(t_i) t_i dt_i \\ &= \sum_i t_i \int_0^{t_i} f_i(s) \tilde{p}_i(s) ds \Big|_{t_i=0}^{\bar{t}_i} - \sum_i \int_0^{\bar{t}_i} \int_0^{t_i} f_i(s) \tilde{p}_i(s) ds dt_i \\ &= \int_0^{\bar{t}} \sum_i \int_{\alpha}^{\bar{t}_i} f_i(s) \tilde{p}_i(s) ds d\alpha \\ &\leq \int_0^{\alpha^*} [1 - \sum_i \varphi_i F_i(\alpha)] d\alpha + \int_{\alpha^*}^{\bar{t}} [1 - \prod_i F_i(\alpha)] d\alpha, \end{aligned} \quad (10)$$

where the first equality follows from integration by parts, the second by rearranging terms and the inequality follows from (8) and (9).

Note that $\tilde{p}_i(t_i) \geq \varphi_i$ implies $\int_{\underline{t}_i}^0 f_i(t_i) \tilde{p}_i(t_i) t_i dt_i \leq \int_{\underline{t}_i}^0 f_i(t_i) \varphi_i t_i dt_i$. 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 (10) 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

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

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

Step 2: Uniqueness

Note that any feasible interim allocation rule \tilde{p} satisfies the following inequality:

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

Since G is monotone, it is differentiable almost everywhere, and H is differentiable by assumption. Because \hat{p}^* satisfies inequality (10) as an equality, we conclude that any optimal interim allocation rule must satisfy (10) as an equality as well. Hence, any optimal interim allocation rule 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 (8) 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.

Step 3: Feasibility

We will construct a feasible ex-post rule inducing the interim allocation 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 $p_i^*(t) = \frac{\varphi_i}{\prod_{j \neq i} F_j(\alpha^*)}$.¹³ 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, $\sum_i \frac{\varphi_i}{\prod_{j \neq i} F_j(\alpha^*)} \leq 1$.¹⁴ 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 (7). Differentiating h , $h'(x) = \sum_i f_i(x) [\varphi_i - \prod_{j \neq i} F_j(x)] \leq \sum_i f_i(x) [\varphi_i - \prod_{j \neq i} F_j(\alpha^*)] \leq \sum_i \frac{f_i(x)}{F_i(\alpha^*)} [\sum_i \varphi_i F_i(\alpha^*) - \prod_i F_i(\alpha^*)] = 0$, since $h(\alpha^*) = 0$. Thus, $h'(x) \leq 0$ for all $x > \alpha^*$ and as desired h is non-increasing for all $x > \alpha^*$. Thus, $\hat{p}_i(t_i) \geq \varphi_i$ for all $t_i \in \mathcal{T}_i$ and $i \in \mathcal{I}$, and we conclude that \hat{p} is feasible in (R). \square

References

Ben-Porath, E., Dekel, E., Lipman, B.L., 2014. Optimal allocation with costly verification. *Amer. Econ. Rev.* 104, 3779–3813.
 Border, K.C., 1991. Implementation of reduced form auctions: a geometric approach. *Econometrica* 59, 1175–1187.

¹³ If $F_j(\alpha^*) = 0$ for some j , we define the ex-post rule to always allocate to the agent with the highest type.

¹⁴ For 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.

- Che, Y.-K., Kim, J., Mierendorff, K., 2013. Generalized reduced-form auctions: a network-flow approach. *Econometrica* 81, 2487–2520.
- Erlanson, A., Kleiner, A., 2019. Costly verification in collective decisions. Working Paper.
- Gershkov, A., Goeree, J.K., Kushnir, A., Moldovanu, B., Shi, X., 2013. On the equivalence of Bayesian and dominant strategy implementation. *Econometrica* 81 (1), 197–220.
- Lipman, B., 2015. An elementary proof of the optimality of threshold mechanisms.
- Manelli, A.M., Vincent, D.R., 2010. Bayesian and dominant-strategy implementation in the independent private-values model. *Econometrica* 78 (6), 1905–1938.
- Mierendorff, K., 2011. Asymmetric reduced form auctions. *Econ. Lett.* 110 (1), 41–44.
- Mierendorff, K., 2016. Optimal dynamic mechanism design with deadlines. *J. Econ. Theory* 161, 190–222.
- Mookherjee, D., Reichelstein, S., 1992. Dominant strategy implementation of bayesian incentive compatible allocation rules. *J. Econ. Theory* 56 (2), 378–399.
- Mylovanov, T., Zapechelnyuk, A., 2017. Optimal allocation with ex-post verification and limited penalties. *Amer. Econ. Rev.* 107, 2666–2694.
- Pai, M., Vohra, R., 2014. Optimal auctions with financially constrained bidders. *J. Econ. Theory* 150, 383–425.