



PIER

PENN INSTITUTE *for* ECONOMIC RESEARCH
UNIVERSITY *of* PENNSYLVANIA

The Ronald O. Perelman Center for
Political Science and Economics (PCPSE)
133 South 36th Street
Philadelphia, PA 19104-6297

pier@econ.upenn.edu

<http://economics.sas.upenn.edu/pier>

PIER Working Paper 19-002

Dynamic Mechanisms with Verification

MARKOS EPITROPOU
University of Pennsylvania

RAKESH VOHRA
University of Pennsylvania

February 18, 2019

<https://ssrn.com/abstract=3337074>

Dynamic Mechanisms with Verification

Markos Epitropou¹ and Rakesh Vohra^{1,2*}

¹ Department of Electrical and Systems Engineering, University of Pennsylvania

² Department of Economics, University of Pennsylvania

Abstract. We consider a principal who allocates an indivisible object among a finite number of agents who arrive on-line, each of whom prefers to have the object than not. Each agent has access to private information about the principal’s payoff if he receives the object. The decision to allocate the object to an agent must be made upon arrival of an agent and is irreversible. There are no monetary transfers but the principal can inspect agents’ reports at a cost and punish them. A novelty of this paper is a reformulation of this dynamic problem as a compact linear program. Using the formulation we characterize the form of the optimal mechanism and reduce the dynamic version of the inspection problem with identical distributions to an instance of the secretary problem with one fewer secretary and a modified value distribution. This reduction also allows us to derive a prophet inequality for the dynamic version of the inspection problem.

1 Introduction

In many large organizations scarce resources must be allocated internally without the benefit of prices. Examples include, the headquarters of a firm that must choose between multiple investment proposals from each of its division managers and funding agencies allocating a grant among researchers. In these settings the private information needed to determine the right allocation resides with the agents and the principal must rely on verification of agents’ claims, which can be costly. We interpret verification as acquiring information (e.g., requesting documentation, interviewing an agent, or monitoring an agent at work), which can be costly. The headquarters of the diversified firm can hire an external firm to conduct an assessment of any division manager’s claims, for example. The funding agency must allocate time to evaluate the claims of the researcher applying for a grant. Furthermore, in these settings, the principal can punish an agent if his claim is found to be false. For example, the head of personnel can reject an applicant, fire an employee or deny a promotion. Funding agencies can cut off funding.

Prior work considered the static version of this problem only. Ben-Porath et al. [2014], which introduced the question, assumes punishment is unlimited in the sense that an agent can be rejected and not receive the resource. Punishment can be limited because verification is imperfect or information arrives only after an agent has been hired for a while. In Mylovanov and Zapechelnuyk [2017], verification is free, but punishment is limited. Li [2017] generalizes both papers by incorporating costly verification *and* limited punishment.

This paper introduces and analyzes a dynamic version of this problem. Specifically, there is a principal who has to allocate one indivisible object among a finite number of agents. The value to the principal of assigning the object to a particular agent is the private information of the agent. Each agent prefers to possess the object than not. The principal would like to give the object to the agent who has the highest value to her.

The agents arrive and depart one at a time, and the decision to allocate the object to an agent must be made upon arrival of an agent. If the principal declines to allocate the object

* Research supported in part by DARPA grant HR001118S0045.

to an agent, the agent departs and cannot be recalled. If the principal allocates the object to an agent, the decision is irreversible.

If each agent were to truthfully report the value to the principal, the principal faces a cardinal version of the secretary problem Krenzel and Sucheston [1977, 1978]: one must select online an element (a ‘secretary’) with maximum value from a randomly ordered sequence. An element has to be selected or discarded upon its arrival, and this decision is irrevocable. If the principal knows the value distribution associated with each agent, she solves a straightforward stopping problem to determine which agent should receive the object. The solution would involve a sequence of thresholds, indexed by the agent, and the principal allocates the object to the first agent whose reported value exceeds their corresponding threshold.

If the principal were to adopt such a policy in our setting it would encourage all agents to exaggerate their values. To discourage agents from exaggerating, the principal can ration at the top of the distribution of values or the principal can verify an agent’s claim and punish him if his claim is found to be false. The first reduces allocative efficiency while the second is costly. The goal of this paper is to find the optimal way to provide incentives via these two devices in a dynamic setting. A novelty of this paper is a reformulation of this dynamic problem as a compact linear program that may be useful in other applications. Using the formulation we can elaborate on the form of the optimal mechanism and reduce the dynamic version of the inspection problem with identical distributions to an instance of the secretary problem with one fewer secretary and a modified value distribution. This reduction also allows us to derive a prophet inequality (Samuel-Cahn [1984]) for the dynamic version of the inspection problem.

Our work is related to the extensive literature on versions of the secretary problem where the principal can rely on prices that was initiated in Hajiaghayi et al. [2007], Chawla et al. [2010]. This was subsequently extended to include additional constraints such as cardinality constraints Hajiaghayi et al. [2007], Alaei [2011], matroids Kleinberg and Weinberg [2012], matchings Alaei et al. [2012], and knapsack constraints Feldman et al. [2015], Duetting et al. [2017]. The absence of money in our setting means that the results from this paper do not apply. However, linear programming approach may be useful in analyzing problems when the principal has access to prices.

In section 2 we introduce our setting and the linear programming formulation. In section 3 we characterize the form of the optimal mechanism and provide a corresponding prophet inequality. In section 4 we study the variation of the problem with limited punishment.

2 Model

There is a single indivisible good to allocate among a set of agents denoted by $I = \{1, \dots, n\}$. The type of agent $i \in I$ is t_i which is the value to the principal of allocating the object to agent i . We assume that the agents’ types are independently distributed. The distribution of agent’s i type has strictly positive density f_i over the interval $T_i = [\underline{t}_i, \bar{t}_i]$. The preferences of the agents are simple: each prefers to possess the object to not. The actual private benefit enjoyed by an agent from receiving the object does not need to be specified.

The agents arrive one by one and report their type, not necessarily truthfully. The principal can inspect the reported type of agent i at cost $c > 0$ and determine perfectly if the agent has lied. In the event an agent is discovered to have lied, we withhold the object from them. This is the case of unlimited punishment. The case of limited punishment is considered later.

By the revelation principle we can restrict attention to direct mechanisms. Denote by t_i the profile of reported types made by all agents upto and including agent i . We write $t_{<i}$ to

denote the profile of reported types made by all agents upto but not including i . A directed mechanism specifies for each profile of type reports, two probabilities for each agent i : the actual probability $q_i(t_i)$ he is assigned the good conditional on the event that the good is not already allocated. Specifically, $q_i(t_i) = Pr[\text{choose } t_i | 1, \dots, i-1 \text{ not allocated}]$. This fully captures the set of online strategies, since independence means there is no need to condition the decision to allocate the good to agent i upon $t_{<i}$. Similarly, let $a_i(t_i)$ be the probability agent i is assigned the good *and* inspected conditional on the event that the good is not already allocated.

These variables must satisfy the following conditions:

$$0 \leq a_i(t_i) \leq q_i(t_i) \leq 1 \quad \forall i \in I \quad \forall t_i \in T_i \quad (1)$$

The incentive compatibility constraints are as follows:

$$q_i(t_i) \geq q_i(t'_i) - a_i(t'_i) \quad \forall i \in I \quad \forall t_i, t'_i \in T_i \quad (2)$$

The principal would like to choose the allocation probabilities q and a to maximize:

$$\sum_{i \in I} \mathbb{E}_{t_{<i}} \left[\prod_{j < i} (1 - q_j(t_j)) \right] \mathbb{E}_{t_i} [r_i q_i(t_i)].$$

2.1 Reduced Form Representation

We work with a reduced form representation of the allocation and inspection rules (see for example C Border [1991], Vohra [2012], Li [2017]). Given a mechanism (q, a) , let $Q_i(t_i) = q_i(t_i) \mathbb{E}_{t_{<i}} \left[\prod_{j < i} (1 - q_j(t_j)) \right]$ and $A_i(t_i) = a_i(t_i) \mathbb{E}_{t_{<i}} \left[\prod_{j < i} (1 - q_j(t_j)) \right]$ be the interim allocation and inspection probabilities respectively. The interim allocation and inspection probabilities are related to the ex-post allocation and inspection probabilities as follows:

Lemma 1.

$$q_i(t_i) = \frac{Q_i(t_i)}{1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)]} \quad (3)$$

$$a_i(t_i) = \frac{A_i(t_i)}{1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)]} \quad (4)$$

Proof. We prove (3). The proof of (4) is similar. Now, $Q_i(t_i) = q_i(t_i) \mathbb{E}_{t_{<i}} \left[\prod_{j < i} (1 - q_j(t_j)) \right]$.

Thus,

$$q_i(t_i) = \frac{Q_i(t_i)}{\mathbb{E}_{t_{<i}} \left[\prod_{j < i} (1 - q_j(t_j)) \right]}.$$

It suffices to prove the following:

$$\mathbb{E}_{t_{\leq i}} \left[\prod_{j \leq i} (1 - q_j(t_j)) \right] = 1 - \sum_{j \leq i} \mathbb{E}_{t_j} [Q_j(t_j)]$$

We do so by induction. For $i = 1$, the equality reduces to

$$\mathbb{E}_{t_1} [1 - q_1(t_1)] = 1 - \mathbb{E}_{t_1} [Q_1(t_1)]$$

which holds since $Q_1(t_1) = q_1(t_1)$. Let's now prove the equality for i . This holds since

$$\begin{aligned}
\mathbb{E}_{t_{\leq i}}\left[\prod_{j \leq i} (1 - q_j(t_j))\right] &= \mathbb{E}_{t_i}[(1 - q_i(t_i))]\mathbb{E}_{t_{< i}}\left[\prod_{j < i} (1 - q_j(t_j))\right] \\
&= (1 - \mathbb{E}_{t_i}[q_i(t_i)])\mathbb{E}_{t_{< i}}\left[\prod_{j < i} (1 - q_j(t_j))\right] \\
&= \mathbb{E}_{t_{< i}}\left[\prod_{j < i} (1 - q_j(t_j))\right] \\
&\quad - \mathbb{E}_{t_i}[q_i(t_i)]\mathbb{E}_{t_{< i}}\left[\prod_{j < i} (1 - q_j(t_j))\right] \\
&= 1 - \sum_{j=1}^{i-1} \mathbb{E}_{t_j}[Q_j(t_j)] - \mathbb{E}_{t_i}[Q_i(t_i)] \\
&= 1 - \sum_{j=1}^i \mathbb{E}_{t_j}[Q_j(t_j)]
\end{aligned}$$

where the first equality follows from independence, the second equality follows from linearity of expectations, and the fourth equality follows from the inductive step and the definition of the interim allocation.

It follows from lemma 1 that the set of constraints (1) can be reduced to

$$\begin{aligned}
Q_i(t_i) + \sum_{j < i} \mathbb{E}_{t_j}[Q_j(t_j)] &\leq 1 \quad \forall i \in I \quad \forall t_i \in T_i \\
0 \leq A_i(t_i) &\leq Q_i(t_i) \quad \forall i \in I \quad \forall t_i \in T_i
\end{aligned}$$

Using the reduced form representation we can formulate the principal's problem as the following linear program (denoted LP):

$$\begin{aligned}
\max_{Q, A} \quad &\sum_{i \in I} \mathbb{E}_{t_i}[t_i Q_i(t_i) - c A_i(t_i)] \\
s.t. \quad &Q_i(t_i) + \sum_{j < i} \mathbb{E}_{t_j}[Q_j(t_j)] \leq 1 \quad \forall i \quad \forall t_i \in T_i \\
&Q_i(t_i) \geq Q_i(t'_i) - A_i(t'_i) \quad \forall i \in I \quad \forall t_i, t'_i \in T_i \\
&0 \leq A_i(t_i) \leq Q_i(t_i) \quad \forall i \in I \quad \forall t_i \in T_i
\end{aligned}$$

3 The Optimal Mechanism

In this section we examine the optimal inspection and allocation rules. Let ϕ_i be the minimum possible share that agent i can receive, i.e.,

$$\phi_i = \inf_{t_i \in T_i} Q_i(t_i). \tag{5}$$

Then, in any optimal solution to the above LP, each agent getting more than the minimum possible must be inspected for this extra winning probability, i.e.

$$A_i(t_i) = Q_i(t_i) - \phi_i. \tag{6}$$

We can use this to eliminate the inspection variables from LP:

$$\begin{aligned} & \max_{\phi, Q} \sum_{i \in I} \mathbb{E}_{t_i} [Q_i(t_i)(t_i - c)] + \sum_{i \in I} \phi_i c \\ & \text{s.t. } Q_i(t_i) + \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)] \leq 1 \quad \forall i \in I \quad \forall t_i \in T_i \\ & \quad Q_i(t_i) \geq \phi_i \quad \forall i \in I \quad \forall t_i \in T_i \end{aligned}$$

Given the vector ϕ , denote this program by $\text{LP}(\phi)$ and let $V(\phi)$ be the optimal objective value. This program has a feasible solution for all ϕ such that $\sum_i \phi_i \leq 1$. Thus, the initial online mechanism with verification can be reduced to

$$\max_{\phi: \sum_{i \in I} \phi_i \leq 1} V(\phi)$$

We will characterize the optimal allocation and inspection rules given ϕ .

Lemma 2. *The optimal solution of $\text{LP}(\phi)$ is monotonic, i.e.*

$$Q_i(t_i) \leq Q_i(t'_i) \quad \forall i \in I \quad \forall t_i \leq t'_i$$

Proof. Suppose not. Then, there is an i and pair (t_i, t'_i) such that $Q_i(t_i) > Q_i(t'_i)$. We pick an $\epsilon > 0$ such that

$$\begin{aligned} - & Q_i(t_i) - \frac{\epsilon}{f_i(t_i)} \geq Q_i(t'_i), \\ - & Q_i(t'_i) + \frac{\epsilon}{f_i(t'_i)} \leq Q_i(t_i). \end{aligned}$$

If we reduce $Q_i(t_i)$ by $\frac{\epsilon}{f_i(t_i)}$ and increase $Q_i(t'_i)$ by $\frac{\epsilon}{f_i(t'_i)}$ feasibility is preserved. The objective function value increases by $\epsilon(t'_i - t_i) > 0$, which is a contradiction.

Hence, there exists a threshold \hat{t}_i for all i such that $Q_i(t_i) = \phi_i$ for $t_i < \hat{t}_i$ and $Q_i(t_i) > \phi_i$ otherwise. In the optimal solution to $\text{LP}(\phi^*)$ it will hold $Q_i(\hat{t}_i) = \phi_i^*$.

We further identify the optimal strategy as a threshold strategy in each round. A transformation of variables will prove convenient:

$$Q_i(t_i) = \phi_i + x_i(t_i) \tag{7}$$

Given ϕ , we can find the optimal strategy by indentifying the solution to the following linear program:

$$\begin{aligned} & \max_x \sum_{i \in I} \mathbb{E}_{t_i} [x_i(t_i)(t_i - c)] \\ & \text{s.t. } x_i(t_i) + \sum_{j < i} \mathbb{E}_{t_j} [x_j(t_j)] \leq 1 - \sum_{j \leq i} \phi_j \quad \forall i \in I \quad \forall t_i \in T_i \\ & \quad x_i(t_i) \geq 0 \quad \forall i \in I \quad \forall t_i \in T_i \end{aligned}$$

Lemma 3. *Suppose that Q is the optimal solution to $\text{LP}(\phi)$. Then for each agent i , there exists a threshold \hat{t}_i , such that*

$$Q_i(t_i) = \begin{cases} 1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)] & \text{if } t_i \geq \hat{t}_i \\ \phi_i & \text{otherwise} \end{cases} \tag{8}$$

Proof. Suppose we are interested in the allocation and inspection rules when we reach agent i . Fix all other variables to be optimal. We are interested in solving the following linear program

$$\begin{aligned}
& \max_{x_i} \mathbb{E}_{t_i}[x_i(t_i)(t_i - c)] \\
& \text{s.t. } x_i(t_i) \leq 1 - \sum_{j \leq i} \phi_j - \sum_{j < i} \mathbb{E}_{t_j}[x_j(t_j)] \quad \forall i \in I \quad \forall t_i \in T_i \\
& \mathbb{E}_{t_i}[x_i(t_i)] \leq 1 - \sum_{j \leq k} \phi_j - x_k(t_k) - \sum_{j < k, j \neq i} \mathbb{E}_{t_j}[x_j(t_j)] \quad \forall k > i \quad \forall t_k \in T_k \\
& x_i(t_i) \geq 0 \quad \forall i \in I \quad \forall t_i \in T_i
\end{aligned}$$

Now, it is clear that the optimal solution can actually be characterized by a threshold. All high types will be assigned their upper limit till the constraint of the aggregate allocation binds. Thus, the optimal solution x is given by

$$x_i(t_i) = \begin{cases} 1 - \sum_{j \leq i} \phi_j - \sum_{j < i} \mathbb{E}_{t_j}[x_j(t_j)] & \text{if } t_i \geq \hat{t}_i \\ 0 & \text{otherwise} \end{cases}$$

Returning back to Q variables completes the proof.

Now that we can characterize the optimal interim allocation in terms of parameters ϕ, \hat{t} , we provide a form for the ex post allocations.

Corollary 1. *For each agent i there exists a threshold \hat{t}_i and constant α_i , such that the optimal ex post allocation can be written as follows:*

$$q_i(t_i) = \begin{cases} 1 & \text{if } t_i \geq \hat{t}_i \\ \alpha_i & \text{otherwise} \end{cases} \quad a_i(t_i) = \begin{cases} 1 - \alpha_i & \text{if } t_i \geq \hat{t}_i \\ 0 & \text{otherwise} \end{cases}$$

Proof. We use lemma 1 to get the form of the ex post allocation:

$$q_i(t_i) = \frac{Q_i(t_i)}{1 - \sum_{j < i} \mathbb{E}_{t_j}[Q_j(t_j)]} = \begin{cases} \frac{1 - \sum_{j < i} \mathbb{E}_{t_j}[Q_j(t_j)]}{1 - \sum_{j < i} \mathbb{E}_{t_j}[Q_j(t_j)]} & \text{if } t_i \geq \hat{t}_i \\ \frac{\phi_i}{1 - \sum_{j < i} \mathbb{E}_{t_j}[Q_j(t_j)]} & \text{otherwise} \end{cases} = \begin{cases} 1 & \text{if } t_i \geq \hat{t}_i \\ \alpha_i & \text{otherwise} \end{cases}$$

where $\alpha_i = \frac{\phi_i}{1 - \sum_{j < i} \mathbb{E}_{t_j}[Q_j(t_j)]}$.

The form for the ex post verification rule follows by (6).

3.1 Identical Distributions

We examine the case of identical distributions, i.e., $T_i = T$ and $f_i(t) = f(t)$ for all $i \in I, t_i \in T$. In this case, we can give a neat representation of the optimal strategy.

Let $\mu = \mathbb{E}_t[t]$. Now, $\text{LP}(\phi)$ can be written as

$$\begin{aligned}
& \max_x \sum_{i \in I} \mathbb{E}_t[x_i(t)(t - c)] + \mu \sum_i \phi_i \\
& \text{s.t. } x_i(t) + \sum_{j < i} \mathbb{E}_t[x_j(t)] \leq 1 - \sum_{j \leq i} \phi_j \quad \forall i \in I \quad \forall t \in T \\
& x_i(t) \geq 0 \quad \forall i \in I \quad \forall t \in T
\end{aligned}$$

Let ϕ^* be the vector ϕ that maximizes $V(\phi)$. A simple perturbation argument shows that there exists ϕ^* such that $\phi_i^* = 0$ for all $i < n$. Furthermore, by equation (5), $x_n(t) = 0$ for all $t \in T_n$ in an optimal solution. We can now reduce $\text{LP}(\phi^*)$ to the following linear program for identifying the strategies for the first $n - 1$ agents:

$$\begin{aligned} \max_x \quad & \sum_{i \in I \setminus \{n\}} \mathbb{E}_t[x_i(t)(t - c)] \\ \text{s.t.} \quad & x_i(t) + \sum_{j < i} \mathbb{E}_t[x_j(t)] \leq 1 - \phi_n^* \quad \forall i \in I \setminus \{n\} \quad \forall t \in T \\ & x_i(t) \geq 0 \quad \forall i \in I \setminus \{n\} \quad \forall t \in T \end{aligned}$$

By normalizing the resource variables this linear program can be interpreted as arising from a secretary problem with $n - 1$ secretaries, where the value of each ‘secretary’ is $t - c$ drawn according to a density function f . In case the object is still available in the last round it is given to the last agent.

3.2 Prophet Inequality

We derive a prophet inequality for the setting with inspection using the reduced form. It scales the optimal offline solution so as to make it a feasible solution for the online setting. This technique can also be used in the standard setting.

Theorem 1. *The optimal online algorithm achieves at least 1/2 of the performance of the optimal offline algorithm on expectation.*

Proof. Let $Q_i^*(t_i)$ the interim expected probability with which agent i with type t_i receives the item in the optimal off-line solution. Let $\phi_i^* = \inf_{t_i} Q_i^*(t_i)$ as proposed in Ben-Porath et al. [2014]. The expected total value to the principal is given by

$$\sum_{i \in I} [\mathbb{E}_{t_i}[Q_i^*(t_i)(t_i - c)] + \phi_i^* c].$$

Pick online values $Q_i(t_i) = \frac{1}{2}Q_i^*(t_i)$ and $\phi_i = \frac{1}{2}\phi_i^*$. It is clear that the objective function with respect to the reduced form for both problems is linear and coincides. Thus, a simple scaling approximates the optimal objective:

$$\sum_{i \in I} [\mathbb{E}_{t_i}[Q_i(t_i)(t_i - c)] + \phi_i c] = \frac{1}{2} \sum_{i \in I} [\mathbb{E}_{t_i}[Q_i^*(t_i)(t_i - c)] + \phi_i^* c]$$

It suffices to prove that the proposed solution is feasible for the online problem.

- $Q_i(t_i) + \sum_{j < i} \mathbb{E}_{t_j}[Q_j(t_j)] = \frac{1}{2}Q_i^*(t_i) + \frac{1}{2} \sum_{j < i} \mathbb{E}_{t_j}[Q_j^*(t_j)] \leq 1$: This holds since $Q_i^*(t_i) \leq 1$ and the expected offline allocation for the first $i - 1$ agents is also less than 1.
- $Q_i(t) \geq \phi_i$: The constraint coincides with the offline constraint. Nothing changes by scaling both sides of the inequality.

4 Limited Punishment

We further illustrate the applicability of our technique under a slight change of the verification process. We say that punishment is limited if the principal cannot reduce an agent's payoff to his outside option by punishing him. If we interpret verification as acquiring information, then punishment can be limited because information is imperfect.³ We assume that punishment is proportional to the private benefit enjoyed by the agent from receiving the object. If v_i is the private benefit enjoyed by agent i , punishment is $k_i v_i$, where each $k_i \in [0, 1]$. These are the same assumptions as in Li [2017]. As we show below, limited punishment will cause the principal to 'ration at the top' as well. All types above some threshold face the same probability of receiving the good.

By the Revelation Principle we can focus on direct mechanisms. In this case, if an agent is inspected, it is optimal to penalize him if and only if he is found to have lied. After the allocation is made, the planner will observe the agent's type and destroy a fraction k_i of the agent's payoff. A direct mechanism specifies for each profile of type reports the probability $q_i(t_i)$ that the good is assigned to agent i conditional on the event that it is not already assigned. These variables must satisfy the following feasibility conditions:

$$0 \leq q_i(t_i) \leq 1 \quad \forall i \in I \quad \forall t_i \in T_i \quad (9)$$

The incentive compatibility constraints are as follows:

$$q_i(t_i) \geq (1 - k_i)q_i(t'_i) \quad \forall i \in I \quad \forall t_i, t'_i \in T_i \quad (10)$$

The principal would like to choose the allocation probabilities q to maximize:

$$\sum \mathbb{E}_{t_{<i}} \left[\prod_{j<i} (1 - q_j(t_j)) \mathbb{E}_{t_i} [t_i q_i(t_i)] \right].$$

As before we work with a reduced form representation. This allows us to formulate the optimal mechanism as the following linear program :

$$\begin{aligned} \max_Q \quad & \sum_{i \in I} \mathbb{E}_{t_i} [t_i Q_i(t_i)] \\ \text{s.t.} \quad & Q_i(t_i) + \sum_{j<i} \mathbb{E}_{t_i} [Q_j(t_j)] \leq 1 \quad \forall i \quad \forall t_i \in T_i \\ & Q_i(t_i) \geq (1 - k_i)Q_i(t'_i) \quad \forall i \quad \forall t_i \in T_i \quad \forall t'_i \in T_i \\ & Q_i(t_i) \geq 0 \quad \forall i \quad \forall t_i \in T_i \end{aligned}$$

4.1 The Optimal Mechanism

We simplify the incentive constraint, as in Mylovanov and Zapechelnyuk [2017]. We include the proof for completeness.

³ We take verification cost and punishment level as exogenous but it is possible that the principal can get more precise information by incurring a higher information acquisition cost, which, in turn, leads to a severer expected punishment. The results in this paper readily extend to the case where the principal can jointly optimize over verification cost and punishment level.

Lemma 4. *An allocation rule satisfies incentive compatibility if and only if for all i there exists χ_i such that*

$$(1 - k_i)\chi_i \leq Q_i(t_i) \leq \chi_i \quad \forall t_i \in T_i \quad (11)$$

Proof. If IC holds then (11) holds with $\chi_i = \sup_{t_i} Q_i(t_i)$. Conversely, if (11) holds for some χ_i , then it also holds with $\chi'_i = \sup_{t_i} Q_i(t_i)$, which implies incentive compatibility.

We now write down a linear program which finds the optimal strategy. We know that for optimal χ this linear program is going to return the optimal strategy.

$$\begin{aligned} \max_{Q, \chi} \quad & \sum_{i \in I} \mathbb{E}_{t_i} [t_i Q_i(t_i)] \\ \text{s.t.} \quad & Q_i(t_i) + \sum_{j < i} \mathbb{E}_{t_i} [Q_j(t_j)] \leq 1 \quad \forall i \in I \quad \forall t_i \in T_i \\ & (1 - k_i)\chi_i \leq Q_i(t_i) \leq \chi_i \quad \forall i \in I \quad \forall t_i \in T_i \\ & Q_i(t_i) \geq 0 \quad \forall i \in I \quad \forall t_i \in T_i \end{aligned}$$

We now describe the optimal strategy.

Lemma 5. *Suppose that Q is the optimal online solution. Let $\chi_i = \sup_{t_i \in T_i} Q_i(t_i)$. Then for each agent i , there exists a threshold \hat{t}_i such that*

$$Q_i(t_i) = \begin{cases} \chi_i & \text{if } t_i \geq \hat{t}_i \\ (1 - k_i)\chi_i & \text{otherwise} \end{cases} \quad (12)$$

Proof. Suppose we are interested in the allocation rule when we reach agent i . Fix all other variables to be optimal. We are interested in solving the following linear program:

$$\begin{aligned} \max_{Q_i} \quad & \mathbb{E}_{t_i} [t_i Q_i(t_i)] \\ \text{s.t.} \quad & Q_i(t_i) \leq 1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)] \quad \forall i \in I \quad \forall t_i \in T_i \\ & \mathbb{E}_{t_i} [Q_i(t_i)] \leq 1 - Q_k(t_k) - \sum_{j < k, j \neq i} \mathbb{E}_{t_j} [Q_j(t_j)] \quad \forall k > i \quad \forall t_k \in T_k \\ & (1 - k_i)\chi_i \leq Q_i(t_i) \leq \chi_i \quad \forall t_i \in T_i \\ & Q_i(t_i) \geq 0 \quad \forall t_i \in T_i \end{aligned}$$

Now, it is clear that the optimal solution can actually be characterized by a threshold. All high types will be assigned their upper limit till a constraint for the aggregate allocation binds. We can now say that the optimal online solution has the following form:

$$Q_i(t_i) = \begin{cases} \min\{\chi_i, 1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)]\} & \text{if } t_i \geq \hat{t}_i \\ (1 - k_i)\chi_i & \text{otherwise} \end{cases} \quad (13)$$

The upper limit can be simplified. We prove that

$$\chi_i \leq 1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)] \quad \forall i \in I.$$

Suppose the contrary. We pick $\chi' = 1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)]$. This makes the constraints less strict since the upper bound remains the same but the lower bound reduces. Thus we can reduce the allocation for lower types and increase the allocation of higher types while holding the aggregate allocation steady. This is a contradiction since such a change will increase the total welfare.

In the limited penalties case the ex post allocation will have a slightly different form.

Corollary 2. *For each agent i there exists a threshold \hat{t}_i , and constant β_i , such that the optimal ex post allocation can be written as follows:*

$$q_i(t_i) = \begin{cases} \beta_i & \text{if } t_i \geq \hat{t}_i \\ (1 - k_i)\beta_i & \text{otherwise} \end{cases}$$

Proof. We use lemma 1 to get the form of the ex post allocation:

$$q_i(t_i) = \frac{Q_i(t_i)}{1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)]} = \begin{cases} \frac{\chi_i}{1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)]} & \text{if } t_i \geq \hat{t}_i \\ \frac{(1 - k_i)\chi_i}{1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)]} & \text{otherwise} \end{cases} = \begin{cases} \beta_i & \text{if } t_i \geq \hat{t}_i \\ (1 - k_i)\beta_i & \text{otherwise} \end{cases}$$

where $\beta_i = \frac{\chi_i}{1 - \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)]}$.

4.2 Prophet Inequality

We use the same machinery as before to further illustrate that extra constraints that restrict the optimal solution in both offline and online cases, do not have an effect on the prophet inequality.

Theorem 2. *The optimal online algorithm achieves at least 1/2 of the performance of the optimal offline algorithm on expectation.*

Proof. Let $Q_i^*(t_i)$ be the interim probability with which agent i with type t_i receives the item in the optimal off-line solution. Let $\chi_i^* = \sup_{t_i \in T_i} Q_i^*(t_i)$ as proposed in Mylovanov and Zapechelnuk [2017]. The expected total value to the principal is given by

$$\sum_{i \in I} \mathbb{E}_{t_i} [t_i Q_i^*(t_i)]$$

Pick online values $Q_i(t_i) = \frac{1}{2} Q_i^*(t_i)$ and $\chi_i = \frac{1}{2} \chi_i^*$ for all $i \in I$. It is clear that the objective function with respect to the reduced form for both problems is linear and coincides. Thus, a simple scaling approximates the optimal objective:

$$\sum_{i \in I} \mathbb{E}_{t_i} [t_i Q_i(t_i)] = \frac{1}{2} \sum_{i \in I} \mathbb{E}_{t_i} [t_i Q_i^*(t_i)]$$

it suffices to prove that the proposed solution is feasible for the online problem.

- $Q_i(t_i) + \sum_{j < i} \mathbb{E}_{t_j} [Q_j(t_j)] = \frac{1}{2} Q_i^*(t_i) + \frac{1}{2} \sum_{j < i} \mathbb{E}_{t_j} [Q_j^*(t_j)] \leq 1$: This holds since $Q_i^*(t_i) \leq 1$ and the expected offline allocation for the first $i - 1$ agents is also less than 1.
- $(1 - k_i)\chi_i \leq Q_i(t) \leq \chi_i$: The constraint coincides with the offline constraint. Nothing changes by scaling both sides of the inequalities.

Bibliography

- S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *52nd IEEE Annual Symposium on Foundations of Computer Science*, pages 512–521, 2011.
- Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 18–35, 2012.
- Elchanan Ben-Porath, Eddie Dekel, and Barton L. Lipman. Optimal allocation with costly verification. *American Economic Review*, 104(12):3779–3813, 2014.
- Kim C Border. Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59:1175–87, 1991.
- Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, pages 311–320, 2010.
- Paul Duetting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *58th IEEE Annual Symposium on Foundations of Computer Science*, pages 540–551, 2017.
- Moran Feldman, Ola Svensson, and Rico Zenklusen. A simple $o(\log\log(\text{rank}))$ -competitive algorithm for the matroid secretary problem. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1189–1201, 2015.
- Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd National Conference on Artificial Intelligence*, pages 58–65, 2007.
- Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, pages 123–136, 2012.
- Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bull. Amer. Math. Soc.*, 83(4):745–747, 1977.
- Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Adv. in Prob. Related Topics*, 4:197–266, 1978.
- Yunan Li. Mechanism Design with Costly Verification and Limited Punishments. *Working Paper*, 2017.
- Tymofiy Mylovanov and Andriy Zapechelnyuk. Optimal allocation with ex post verification and limited penalties. *American Economic Review*, 107(9):2666–94, 2017.
- Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent non-negative random variables. *Ann. Probab.*, 12(4):1213–1216, 11 1984.
- Rakesh V. Vohra. Dynamic mechanism design. *Surveys in Operations Research and Management Science*, 17(1):60–68, 2012.