Sequential Information Acquisition and Optimal Search

Nawaaz Khalfan* Rakesh Vohra†

October 3, 2024

- click here for the most recent version -

Abstract

A principal receives an unknown reward from allocating to an agent who has private information about the reward. Prior to allocating, the principal may elicit a report from the agent and inspect them at a cost, but must do so without transfers. When the private information is noisy, the mechanism that maximizes the principal's expected return at most segments signals into two groups; inspecting and conditionally allocating to high types, and partially allocating to low types. We prove this by reformulating Weitzman (1979) and extending it to a mechanism design problem through the reduced-form.

1 Overview

Appraising the value of an asset is an essential precursor to its exchange. Employers interview potential employees, public funds assess grant applications, venture capitalists evaluate investment opportunities. This process is often costly, and information that could be used to lower, or even circumvent, these costs, is often privately held.

This paper considers a principal whose return from allocating to an agent is inherent, though uncertain, to the agent they allocate to. The principal has the ability to inspect agents sequentially at a cost and learn about the true return, as well as the opportunity to receive a report from the agent. The agent, independent of their private information, strictly prefers to be allocated to.

This early version describes a set of technical results that allow us to find the principal optimal mechanism, and demonstrates this by deriving the optimal mechanism for a single-agent setting - essentially reproving Khalfan (2023) and generalising Ben-Porath, Dekel and Lipman (2014) to noisy information. The paper concludes by outlining how this approach can be used to find the

*Monash University; nawaaz.khalfan@monash.edu
†University of Pennsylvania; rvohra@seas.upenn.edu

optimal mechanism in the multi-agent setting, as well as describing a conjectured solution.

The environment encompasses many important settings. Consider the following.

1. **Hiring**: a firm, the principal, seeks to fill an open position in their operation with a potential employee, the agent. The agent would like to be hired and, aware of their own characteristics, has an estimate of their fit in the position. The principal can ask for this estimate, and interview the agent themselves, discovering a better forecast of their productivity. The interview, however, is costly for the principal.

2. **Funding**: a governing board, the principal, sets the rules by which it allocates a scarce, publicly owned resource, such as funding for an applicant's project, the agent. The agent is interested in being approved, valuing their own use above rival users, and knows the most about the project's characteristics, likelihood of success and the project's social value. The principal wants to fund positive net value projects.

3. **Investing**: an investor, the principal, determines the way it evaluates and finances early investment opportunities, the agents. The investor may be governed by the motivation to strengthen an existing portfolio and even personal philanthropic concerns, but is restricted in outlining these preferences publicly. The agent wishes to be financed and expand their enterprise, and has the most information about the enterprise. This information doesn't fully determine the investor's value for the opportunity without an appraisal.

In this analysis, the principal's reward is binary, and they must inspect the agent in order to allocate and receive the reward. These assumptions simplify the exposition but do not restrict the application, as discussed in section 5.

For the single-agent setting, the mechanism that maximizes the expected return for the principal has a simple structure. The principal segments agents into two groups: those with sufficiently favourable information - *high signals* - and those who do not pass this threshold - *low signals*. Agents with high signals are always inspected and only allocated to post-inspection if the discovered reward is positive. Agents with low signals are compensated for their report with a smaller probability of unconditional allocation.

This segmentation is the only optimal mechanism that does not entirely *pool* signals. That is, if the principal finds it too costly to partially-separate, they treat all agents non-preferentially. They do this by either never inspecting and rejecting all signals, or always inspecting and allocating only when the realized return is positive. The choice between one of these pooling mechanisms and the partially-separating mechanism depends on the prior over the rewards, the agent's signal accuracy, and the cost of inspection.

The conjectured multi-agent solution is analogous. From an interim perspective, agents with high signals are *maximally* inspected, and allocated to only if the discovered reward is positive. Agents with low signals are not inspected but compensated with a small probability of unconditional allocation. This can be implemented by an ex post mechanism that segments reports, randomly inspects high types, and conditional on not allocating among the high types, randomly selects a low type to allocate to.

To prove this, we map the canonical search problem of Weitzman (1979) into a linear program following Bertsimas and Niño-Mora (1996), and re-derive *Pandora's rule*. We call this Pandora's linear program and, with the inclusion of additional side constraints, allows for the adaption of Pandora's rule to the allocation problem at hand. This demonstrates the applicability of the polyhedral approach, and contributes to the adaption of Pandora's rule to a wider class of problems.

Aside from optimal search, this is related to a branch of the mechanism design literature devoted to costly inspection without transfers and demonstrates how we can recover interesting observations about search behaviour with noise, a feature mostly missing from the branch. With the additional assumption that an agent must be inspected prior to selection, this paper seeks to generalise Khalfan (2023) to multiple agents and Ben-Porath, Dekel and Lipman (2014) to imperfect signals.

Rather than the customary environment setup, we begin with the technical contributions, which conveniently outlines the full information benchmark for the model that follows. We then move onto the single-agent setting and conclude with the multi-agent discussion and summary.

2 Pandora's linear program

The principal's problem when the agent's do not have any private information[1] is identical to *Pandora's problem*, introduced by Weitzman (1979). Adopting the language of Weitzman, a searcher, *Pandora*, has a finite set of *boxes* to select from. Each box contains a prize, drawn independently from a known distribution, whose realization is initially unknown to Pandora. Opening a box reveals the prize and allows her to ultimately select the prize to keep, but is costly. She may open whichever boxes she likes, in whichever order, though she may only select at most one prize.

Weitzman (1979) shows that optimal search is characterized by a simple *index rule*: assign each box an index, dependent only on the box's known characteristics, and sequentially open the box with the highest index among unopened boxes until the best found prize exceeds the highest index remaining, in which case Pandora should stop and collect her best found prize. This result is known as *Pandora's rule*, and the associated index is known as the *Weitzman index*.

[1] or, equivalently, when the private information is sufficiently uninformative

In this section, Pandora's problem is recast as a *dynamic scheduling problem*. Subsequently, this problem can be written as a linear program following Bertsimas and Niño-Mora (1996). This particular program satisfies the sufficient conditions for a solution to be *indexable*, reproving Weitzman's initial result. This section concludes by explicitly deriving the *Weitzman index*.

The purpose here is to demonstrate how to use Bertsimas and Niño-Mora (1996) to transform dynamic choice problems into convenient linear programs, and present a version of Pandora's problem that is more readily adaptable to other settings - in particular, optimal allocation with noisy inspection. This can be thought of as analogous to the Border approach of representing optimal auctions in terms of interim allocations.

2.1 Pandora's problem and rule

Establishing notation, suppose box $i \in \{a, b, \ldots, \omega\} = \Omega$ conceals a prize of x_n, where $n \in \{1, 2, \ldots, N\} = \mathcal{N}$, with probability π_{in}, and costs $c_i > 0$ to open. Order $\mathcal{N}$ such that $n < m$ whenever $x_n > x_m$. Pandora's outside option is given by an already open box with prize $x_0 \geq 0$. These details for a 2 box, N prize environment are displayed in Figure 1.

Pandora seeks to maximize the value of her highest found prize, net of any search costs. Given a set of opened boxes S and a realized vector of prizes x_S, Pandora's stopping value is:

$$v(S, x_S) = \max_{i \in S \cup \{0\}} x_i - \sum_{j \in S} c_j$$

Pandora's objective is to select a sequential, *nonanticipative* opening policy that maximizes her expected stopping value. Here, nonanticipative refers to the requirement that her policy cannot depend on the prize realizations of unopened boxes.

For each box i, its index, denoted z_i, is the solution to the following:

$$c_i = \sum_{x_n > z_i} \pi_{in}(x_n - z_i)$$

Weitzman (1979) shows that Pandora's optimal policy is described by the following algorithm:

1. Set the *reserve value* as the value of the outside option, x_0, and continue to step 2.

2. If the highest index among the set of unopened boxes exceeds the reserve value, open that box and proceed to step 3. Otherwise, stop, take the reserve value and terminate the algorithm.

3. If the realized prize exceeds the reserve value, replace the reserve value with this new prize and return to step 2.

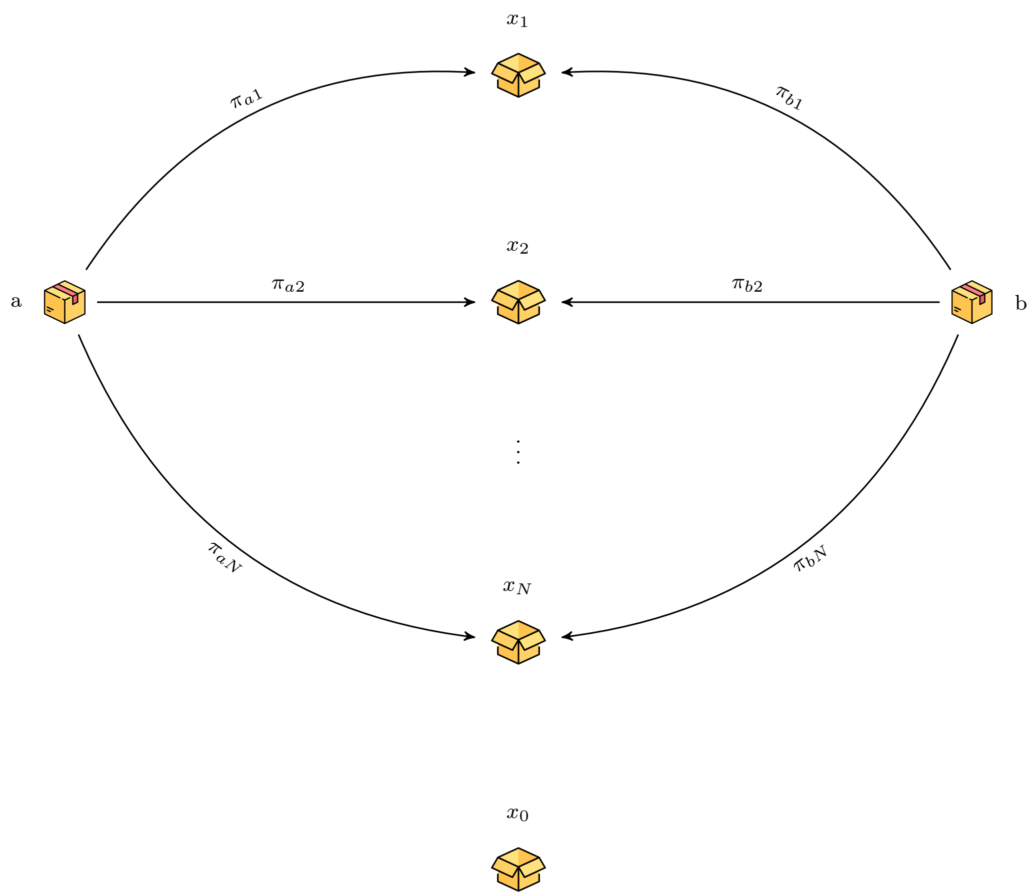

Figure 1: A 2 box, N prize *search* problem

Refer to this policy as **Pandora's rule** and the index, z_i, as **Weitzman's index**.

A dynamic problem is called **indexable** if optimal behaviour is fully described by assigning all potential options an index, ranking these indices into a permutation, σ, and exercising the highest ranked option available at any decision point. It is also said to be **decomposable** if the index of each option depends on characteristics of that option alone. For Pandora's problem, the potential options are *open box* i, for each $i \in \Omega$, and *stop and select prize* n, for each $n \in \mathcal{N} \cup \{0\}$. Upon assigning indices for any revealed prize as simply the value of that prize, $z_n = x_n$ for each $n \in \mathcal{N} \cup \{0\}$, it's clear that Pandora's problem is both indexable and decomposable.

2.2 A dynamic scheduling problem

We can recast Pandora's problem as a dynamic scheduling problem, following Bertsimas and Niño-Mora (1996). In the terminology of queuing theory, a *class* is a particular state and if a *job* is in class $j \in \mathcal{J}$ and is *serviced*, then the scheduler receives a *service reward* r_j and the job transitions to a new class $k \in \mathcal{J}$ with *transition probability* p_{jk}. At each time step, the scheduler must choose which job to be serviced, determining the reward collected at that step, with the goal to maximize the expected sum of discounted rewards.

For Pandora, introduce $\beta \in (0,1)$, a discount factor that exchanges her one-shot returns to equivalent expected returns over an infinite horizon. We now need to define classes, service rewards, and transition probabilities that correspond to her decision nodes. Let $j \in \Omega \cup \mathcal{N} \cup \{0\} = \mathcal{J}$, and,

- if $j \in \Omega$: $r_j = -c_j$ and $p_{jn} = \pi_{jn}$ for all $n \in \mathcal{N}$,
- if $j \in \mathcal{N}$: $r_j = x_j(1-\beta)$ and $p_{jj} = 1$, and
- if $j = 0$: $r_j = x_0(1-\beta)$ and $p_{00} = 1$.

That is, if $j \in \Omega$, then the job is an unopened box, and servicing the job is costly but transitions it into an opened box allowing her to reap a future reward. If $j \in \mathcal{N}$, then the job is an opened box, and servicing the job yields the associated (discounted) reward and remains in that class. Finally, if $j = 0$, then the job is the outside option and has the same properties as an opened box.

Initializing the system with a single job in each class $j \in \Omega \cup \{0\}$, the infinite horizon optimal dynamic servicing problem of these jobs is a discounted analogue of Pandora's problem that coincides as $\beta \to 1$. The key details of Pandora's scheduling problem for a 2 box, N prize environment are depicted in figure 2, where the coloured nodes are classes with available jobs to service.

To illustrate, suppose at the first time step, the Pandora selects a job in class a (the node coloured blue) for servicing. She incurs a reward of $-c_a$, that is, she incurs the cost of opening that box. After servicing, the job transitions stochastically to another class, say, class 1, as depicted in Figure 2 by the arc $(a, 1)$ terminating in node 1. This transition captures the realized reward from having opened box a. Suppose at the next time step, Pandora selects the job now in class 1 for servicing. She collects a reward of $x_1(1-\beta)$ and this job never transitions to a new class. If in all subsequent time steps she always select a job in class 1 for servicing, the discounted payoff of the policy described is $-c_a + \beta x_1$. As $\beta \to 1$ this corresponds to the payoff from opening box a, seeing prize 1, stopping and keeping that prize.

Alternatively, after servicing a job in class a, Pandora may find it preferable to service a job in class b (the node coloured red). She then incurs a reward of $-c_b$ and the job transitions stochastically to another class. Suppose she ultimately decides to service the job in class 0 (the node coloured green) from then on. The discounted payoff from this policy is $-c_a - \beta c_b + \beta^2 x_0$, and as $\beta \to 1$ this corresponds to the payoff from opening both boxes and taking the outside option.

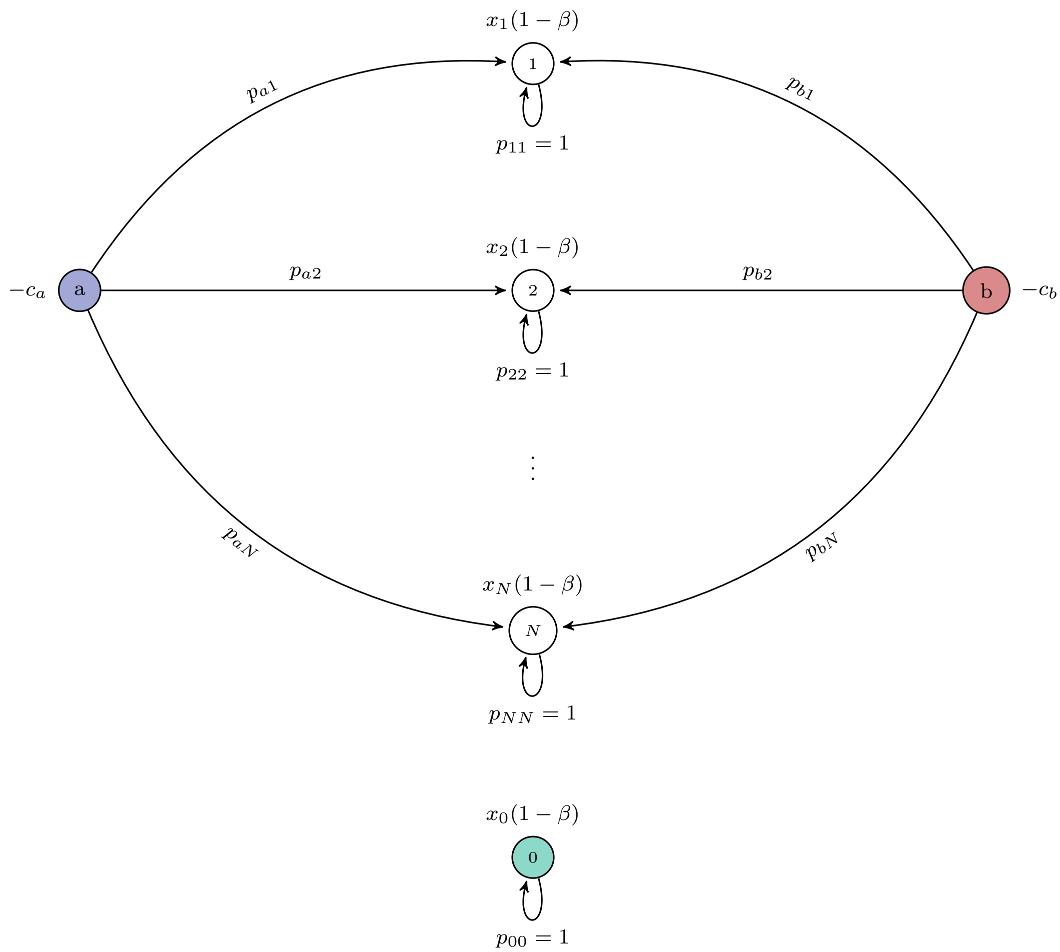

Figure 2: A 2 box, N prize *scheduling* problem

Let $\mathcal{U}$ be the set of all admissible *policies* - which jobs to service after every possible history - and define $\mathbb{1}_j(t)$ as the indicator variable for whether a job in class j is serviced at time t. Then $\mathcal{Z}$, the discounted sum of rewards from the optimal dynamic schedule, is given by,

$$\mathcal{Z} = \max\left\{\mathbb{E}_u\left[\sum_{t=0}^{\infty}\sum_{j\in\mathcal{J}} r_j \mathbb{1}_j(t)\beta^t\right] \mid u \in \mathcal{U}\right\} \tag{DSP}$$

Call this the *dynamic scheduling problem* (DSP). Define λ_j^u as the expected, discounted *service time* of a job in class j under policy $u \in \mathcal{U}$. That is,

$$\lambda_j^u := \mathbb{E}_u\left[\sum_{t=0}^{\infty} \mathbb{1}_j(t)\beta^t\right]$$

Then the vector of these service times lives in,

$$\Lambda = \left\{ (\lambda_j^u)_{\mathcal{J}} \mid u \in \mathcal{U} \right\}$$

We can then outline an equivalent mathematical program (MP):

$$\mathcal{Z} = \max \sum_{j \in \mathcal{J}} r_j \lambda_j \quad \text{s.t.} \quad (\lambda_j)_{\mathcal{J}} \in \Lambda \tag{MP}$$

We now have a linear objective. For this reformulation to be useful, we need a characterization of the feasible space of service times, Λ. As it turns out, for this dynamic schedule, Λ is a convex polytope - the intersection of a finite number of half-spaces - and thus, MP is a linear program. Furthermore, Λ is an extended polymatroid - a convex polytope defined by a submodular function - and thus the linear program can be solved with a greedy algorithm which corresponds to an index policy. Note that extended polymatroids generalize the notion of polymatroids, which are used in Border's reduced form characterization.

2.2.1 Conservation laws

To characterize the feasible space of service times, Λ, we first derive a set of conservation laws regarding the service times to identify inequalities that any job must satisfy. Subsequently, we are going to show that these inequalities are not only necessary but sufficient for describing Λ. This derivation is detailed in appendix A.1, but a brief outline is included here.

First, with the inclusion of the outside option, the scheduler never stands to strictly benefit from being idle. As such, one job must be serviced in each time period:

$$\sum_{j \in \mathcal{J}} \lambda_j^u = \frac{1}{1-\beta} \quad \forall u \in \mathcal{U} \tag{1}$$

Secondly, we include an expression about the maximum possible service time under all admissible policies, $u \in \mathcal{U}$, as they pertain to classes in subset S. From the sets individual classes, we can build conservation laws for all subsets, S, providing upper bounds on the total service times for jobs in those subsets (see appendix A.1 and the inequalities 2.1, 2.2, and 2.3). Conveniently, these subsets fall into two interesting cases.

In the first, if $0 \in S$, then the subset of jobs can be serviced immediately and indefinitely by servicing the outside option. This gives us the following, mostly redundant, inequality:

$$\sum_{j \in S} \lambda_j^u \leq \frac{1}{1-\beta} \quad \forall u \in \mathcal{U} \tag{3.1}$$

For the second, suppose $0 \notin S$. Then S can be partitioned into two sets: $S^\Omega := S \cap \Omega$, and

$S^{\mathcal{N}} := S \cap \mathcal{N}$. The following priority policy intuitively achieves this maximum:

1. Service all jobs in S^{Ω}.
2. If any job from S^{Ω} has transitioned into a class from $S^{\mathcal{N}}$, service this job indefinitely.
3. If no job from S^{Ω} has transitioned into a class from $S^{\mathcal{N}}$, then service remaining jobs in $\Omega \backslash S^{\Omega}$, until a job transitions into a class from $S^{\mathcal{N}}$ and then service this job indefinitely.

Evaluating this policy and rearranging to follow convention that choice variables are collected on the left-hand side of the constraints, we have our final conservation law:

$$\sum_{j \in S} \lambda_j^u - \beta^{|S^{\Omega}|} \Big[\prod_{i \in S^{\Omega}} (1 - \sum_{j \in S^{\mathcal{N}}} p_{ij}) \Big] \Big[\sum_{i \in \Omega \backslash S^{\Omega}} \lambda_i^u \sum_{j \in S^{\mathcal{N}}} p_{ij} \frac{\beta}{1-\beta} \Big] \leq \frac{1}{1-\beta} \Big[1 - \beta^{|S^{\Omega}|} \Big[\prod_{i \in S^{\Omega}} (1 - \sum_{j \in S^{\mathcal{N}}} p_{ij}) \Big] \Big] \quad \forall u \in \mathcal{U} \tag{3.2}$$

Inequalities 3.1 and 3.2 generalise the singleton set expressions, and as such, along with 1, they define the *maximal service time region*, $\mathcal{L}$:

$$\mathcal{L} := \{ \lambda \in \mathbb{R}_+^{|\mathcal{J}|} \mid \lambda \text{ satisfy } 1, 3.1, \text{ and } 3.2 \}$$

2.2.2 Sufficiency

Bertsimas and Niño-Mora (1996) show that if the feasible space of service times satisfy *generalized conservation laws*, then they can be represented as an *extended polymatroid*. This means that optimizing a linear objective over the feasible space will not only yield an extreme solution, but the extreme solution is characterized by a priority-index over classes. Conditional on one further assumption on the structure of the feasible space, the index is also decomposable. These are the subject of Bertsimas and Niño-Mora (1996)'s Theorem 1, 2, and 3.

The scheduling problem presented here is a special case of the classic *multiarmed bandit problem*. In particular, it is a multiarm bandit with one intransient arm, *the outside option*, and N transient arms, *the unopened boxes*. The transient arms are available from the start, and stochastically map into payoff-varied intransient arms, *the opened boxes*. The multiarmed bandit problem is shown to be a satisfy Bertsimas and Niño-Mora's generalized conservation laws in Proposition 8 and Theorem 11. Given our problem is simpler than the general multiarm bandit problem, these results are refined and restated in the following theorem.

Theorem 1 *The performance region, Λ, is the extended polymatroid defined by the maximal service region, $\mathcal{L}$. Further, for each class $j \in \mathcal{J}$ there exists indices, z, depending only on characteristics of that class, such that an optimal policy is to schedule a job with the largest current index.*

Note that, in its original generality, this reproves the celebrated result of Gittins and Jones (1974): *there exists a decomposable index that dictates the optimal scheduling of projects in the multiarmed bandit problem.* For our problem, this proves that the schedule's solution is also characterized by a decomposable priority-index, which we can directly take advantage of.

2.3 Pandora's schedule

Making the most of these properties, we can explicitly solve for the optimal schedule.

Theorem 2 *The solution to Pandora's schedule is given by an indexing rule defined by an index, z, where if $j \in \mathcal{N} \cup \{0\}$ then $z_j = r_j$, and, if $j \in \Omega$, then z_j is the solution to:*

$$-r_j = \sum_{\substack{i \in \mathcal{N} \\ r_i > z_j}} p_{ji}(r_i - z_j)$$

As outlined in section A.2, Pandora's scheduling problem satisfies Bertsimas and Niño-Mora's generalized conservation laws and so the feasible space of achievable performance is an extended polymatroid. Optimization of a linear objective over an extended polymatroid is solved by an adaptive greedy algorithm, which leads to an optimal solution having the indexability property. Further, this extends to the decomposability property, which means the indices must apply to all sub-problems. Then, by normalizing the index for the outside option as the value of the outside option itself: $z_0 = r_0$, we can build the optimal index by considering sub-problems that only include pairs of possible jobs. The full proof is included in appendix A.3.

Pandora's rule is then given by the limiting behaviour of the optimal schedule, referred to as **Pandora's schedule**.

Corollary 1 *As $\beta \to 1$, Pandora's schedule is precisely Pandora's rule.*

Pandora's problem is now fully described by this linear program which takes as inputs a vector of returns for opening or selecting a box in each class, r, and a transition matrix of unopened boxes into prizes, p, and outputs a vector of service times for each class, λ. While these service times describe discounted expected behaviour over an infinite horizon, they can easily be translated into the undiscounted, finite horizon of Weitzman by taking the limiting behaviour as the discount approaches one, $\beta \to 1$.

Adapting this program to other environments then involves remapping the returns from opening and selecting, r, or introducing additional constraints on the service times, λ. For example:

- we may want to adapt r to reflect the concern that Pandora may have discovered a prize but subsequently lost it to a rival searcher, or

- we may want to restrict the choice of λ so that the boxes and their known characteristics are indeed available to Pandora; that is individual rationality and incentive compatibility.

Then, the simplest adaptions involve an alteration to the index, reflecting changes to the cost-benefit analysis, and more advanced adaptations involve the introduction of *side constraints*, that must live within the generalized conservation laws for the problem to continue to have the properties highlighted here. To demonstrate this, consider the single-agent environment of the informed-agent allocation problem.

3 Single-agent environment

Now take a single box, a, that contains a high, h, prize with probability $p \in [0,1]$, and a low, ℓ, prize with complementary probability $1-p$. For a cost $c > 0$, the searcher, Pandora, can discover which prize the box contains and have the option of keeping the prize. Alternatively, at any point, she can take an outside option. Let Pandora's payoff from accepting the high prize be 1, the low prize, -1, and the outside option, 0. Following *Pandora's Linear Program*, this problem can be recast as the scheduling problem depicted in the Figure 3.

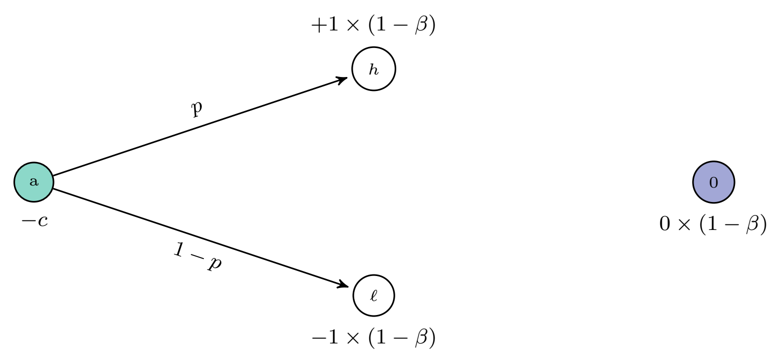

Figure 3: A simple 1 box, 2 prize scheduling problem

3.1 Full information

Suppose Pandora knows p. Then, following the formulation in *Pandora's linear program*, her **interim first best value** is given as:

$$
\begin{array}{rccccccccc}
\overline{v}(p,\beta) = \max\limits_{\lambda \geq 0} & -c\lambda_a & + & (1-\beta)\lambda_h & - & (1-\beta)\lambda_\ell & & & & \\
\text{s.t.} & \lambda_a & + & \lambda_h & + & \lambda_\ell & + & \lambda_0 & = & \frac{1}{1-\beta} \\
& \lambda_a & & & & & & & \leq & 1 \\
& -\frac{p\beta}{1-\beta}\lambda_a & + & \lambda_h & & & & & \leq & 0 \\
& -\frac{(1-p)\beta}{1-\beta}\lambda_a & & & + & \lambda_\ell & & & \leq & 0
\end{array}
$$

A solution to this problem is $\lambda^\star$, where:

- if $p\beta - c \geq 0$, then $\lambda^\star_a = 1$, $\lambda^\star_h = \frac{p\beta}{1-\beta}$, and $\lambda^\star_\ell = 0$, and
- if $p\beta - c < 0$, then $\lambda^\star_0 = \frac{1}{1-\beta}$.

Evaluating the value as $\beta \to 1$,

$$\overline{v}(p) = \max\{p - c, 0\}$$

Reading this as a solution to the original problem: Pandora opens the box if her expected reward from keeping only the high prize exceeds the cost of opening the box. Otherwise, she takes the outside option.

Suppose that this box is initially drawn from a finite distribution Π with support on $\mathcal{P}$, such that $\pi(p)$ is the probability the box has type p. For reference, let $|\mathcal{P}| =: n$. Then the **ex ante first best value** is given as:

$$\overline{V} = \sum_{p \in \mathcal{P}} \pi(p)\overline{v}(p)$$

Enumerating by summing over the interim first best:

$$\overline{V} = \sum_{p \geq c} \pi(p)(p - c)$$

The interim first best value and policy are summarised in figures 4 and 5 respectfully.

3.1.1 A general solution

It is important in what follows to know what the solution looks like when the rewards for selecting are generalised. Suppose the reward for scheduling node h is $r_h(1-\beta)$ and node ℓ, $r_\ell(1-\beta)$. Then the problem is given by:

$$\begin{array}{rrcrcrcrcl}
v(p,\beta) = \max\limits_{\lambda \geq 0} & -c\lambda_a & + & r_h(1-\beta)\lambda_h & + & r_\ell(1-\beta)\lambda_\ell & & & & \\
\text{s.t.} & \lambda_a & + & \lambda_h & + & \lambda_\ell & + & \lambda_0 & = & \frac{1}{1-\beta} \\
& \lambda_a & & & & & & & \leq & 1 \\
& -\frac{p\beta}{1-\beta}\lambda_a & + & \lambda_h & & & & & \leq & 0 \\
& -\frac{(1-p)\beta}{1-\beta}\lambda_a & & & + & \lambda_\ell & & & \leq & 0
\end{array}$$

A solution is given by $\lambda^\star$, where:

1. if $r_h, r_\ell \geq 0$, then $\lambda^\star_a = \mathbb{1}\{-c + r_h p\beta + r_\ell(1-p)\beta \geq 0\}$, $\lambda^\star_h = \frac{p\beta}{1-\beta}\lambda^\star_a$ and $\lambda^\star_\ell = \frac{(1-p)\beta}{1-\beta}\lambda^\star_a$,
2. if $r_h \geq 0 > r_\ell$, then $\lambda^\star_a = \mathbb{1}\{-c + r_h p\beta \geq 0\}$, $\lambda^\star_h = \frac{p\beta}{1-\beta}\lambda^\star_a$ and $\lambda^\star_\ell = 0$,
3. if $r_\ell \geq 0 > r_h$, then $\lambda^\star_a = \mathbb{1}\{-c + r_\ell(1-p)\beta \geq 0\}$, $\lambda^\star_h = 0$ and $\lambda^\star_\ell = \frac{(1-p)\beta}{1-\beta}\lambda^\star_a$, and

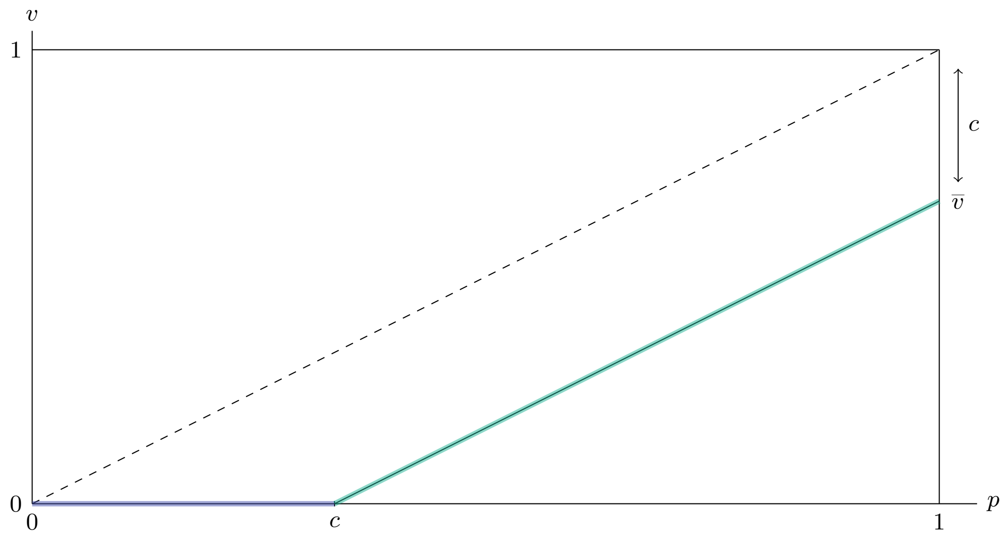

Figure 4: interim first best value, $\overline{v}$, as a function of high prize probability, p, given a cost, $c > 0$

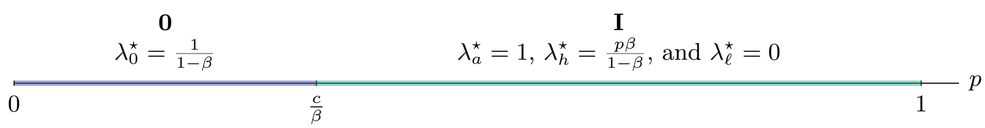

Figure 5: first-best policy, $(\lambda_a^\star, \lambda_h^\star, \lambda_\ell^\star, \lambda_0^\star)$

4. if $0 > r_h, r_\ell$, then $\lambda_h^\star = 0$, $\lambda_\ell^\star = 0$ and $\lambda_a^\star = 0$.

With $\mathbb{1}\{\mathcal{Q}\}$ the indicator function that is equal to 1 if the statement $\mathcal{Q}$ is true given the arguments, and 0 otherwise.[2]

3.2 Private information

Now suppose Pandora does not know p, but the box does. Refer to this as the boxes **type**. The distribution of types, as before, is Π and common knowledge. The box strictly prefers to be selected than not and has no direct payoff from their type. Let their **ex post** payoff be 1 if selected, and 0 otherwise.

If Pandora knows their type, p, and institutes a solution λ, then the boxes *interim payoff* is simply the gross probability that they are selected and so:

$$u(p, \beta) = (1-\beta)\lambda_h(p) + (1-\beta)\lambda_\ell(p)$$

Suppose Pandora tries to implement the first best policy. Then, conditional on the boxes report-

[2]The standard definition of an indicator function is $\mathbb{1}_A(x) \coloneqq 1$ if $x \in A$ and 0 if $x \notin A$. We're more interested in the set A and less in the argument x, so suppress the argument and promote the set.

ing truthfully, their (normalised) payoffs are given in figure 6.

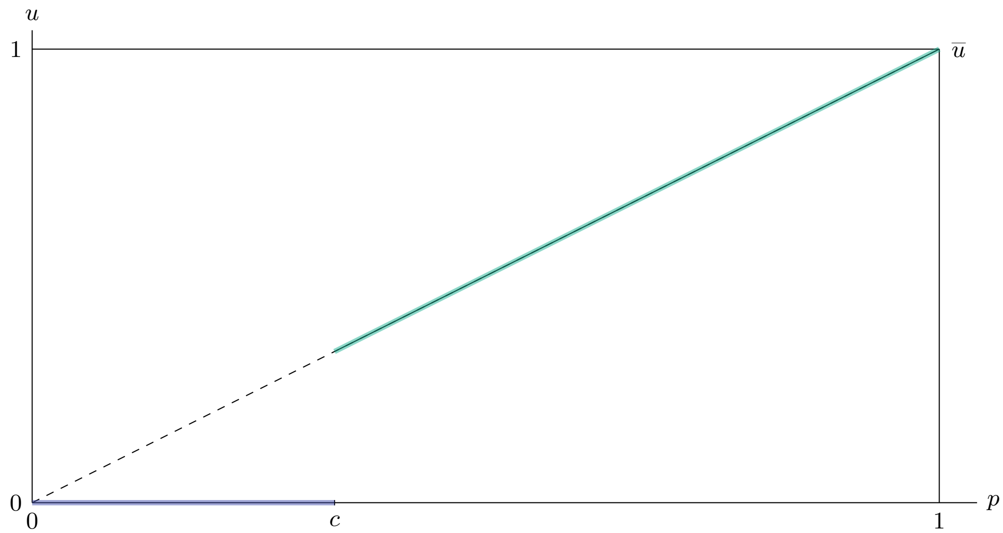

Figure 6: interim payoffs under the first best policy, $\overline{u}$, as a function of type, p, given a cost, c

There are two reasons, however, why Pandora may not be able to implement this rule. Suppose there are only two types, $p_0 = c - \varepsilon$ and $p_1 = c + \varepsilon$. Then clearly this policy isn't incentive compatible as any report that gives p_0 a positive probability of being searched and ultimately selected, has a payoff that exceeds that of reporting they are p_0, the unsearchable type. Secondly, it isn't immediately clear whether Pandora can award such a deviation payoff, as the policy is in terms of ex-ante service times, achievable by a convex combination of extreme priority policies, and thus a function of the box's true type.

As such, let us carefully derive outline the deviation payoffs and then study the incentive compatibility constraints. By normalising for $(1 - \beta)$ we get:

$$u(q|p) = \frac{p}{q}\lambda_h(q) + \frac{1-p}{1-q}\lambda_\ell(q)$$

Collecting terms that depend on the true type p gives us the following convenient form:

$$u(q|p) = \frac{\lambda_\ell(q)}{1-q} + p \cdot \left[\frac{\lambda_h(q)}{q} - \frac{\lambda_\ell(q)}{1-q}\right]$$

The first term is a guaranteed allocation from reporting q, and the second term is an allocation **differential**, or *bonus*, from receiving a high draw, scaled by the agents true type p.[3] Let the

[3]Of course the bonus could, in principle, be negative and thus the first term isn't necessarily *guaranteed*. We'll see later that bonuses are optimally positive.

differential for an arbitrary q be denoted by $\Delta(q)$. That is:

$$\Delta(q) = \frac{\lambda_h(q)}{q} - \frac{\lambda_\ell(q)}{1-q}$$

Further, to condense notation, let $y_\ell(q)$ denote the scaled low draw allocation, and $y_h(q)$ the high. That is:

$$y_\ell(q) = \frac{\lambda_\ell(q)}{1-q} \quad , \quad y_h(q) = \frac{\lambda_h(q)}{q}$$

Then the **incentive compatibility constraints** can be written as:

$$IC_{pq}: \; u(p|p) = y_\ell(p) + p\Delta(p) \geq y_\ell(q) + p\Delta(q) = u(q|p) \quad \forall \; p, q \in \mathcal{P}$$

3.2.1 Virtual-ultity hypothesis

Order and index $\mathcal{P}$ so that $p_i > p_j$ if $i < j$. We will promote the index i to the argument where it does not cause confusion. The following lemma allows us to reduce the number of incentive compatibility constraints from $n(n-1)$ global constraints to $n-1$ equality constraints, $IC_{i,i+1}$ binds for all $i < n$, and $n-1$ inequality constraints, $\Delta(i) \leq \Delta(i+1)$ for all $i < n$.

Lemma 1 *Local upward incentive compatibility binds under the second-best policy. That is, if $\lambda = \lambda^*$, then $IC_{i,i+1}$ binds for every $i < n$.*

This is detailed in appendix B.1 but is the inspection and allocation equivalent of the analogous result from optimal auctions: *local upward incentive compatibility and monotonicity of the allocation are both sufficient and necessary.* This is less straightforward without the inclusion of transfers, but is none-the-less true.

Before outlining and proving the optimal mechanism, let's consider the standard *virtual-utility hypothesis*. Proceed by assuming monotonocity, and working only with the upward local IC constraints. From lemma 4, we have that for each $i \in \{1, \dots, n-1\}$:

$$\lambda_h(i) + \lambda_\ell(i) = \frac{p_i}{p_{i+1}}\lambda_h(i+1) + \frac{1-p_i}{1-p_{i+1}}\lambda_\ell(i+1)$$

Subtracting the left hand-side from the right, and attaching a non-zero multiplier μ_i:

$$0 = \mu_i \cdot \left[\frac{p_i}{p_{i+1}}\lambda_h(i+1) + \frac{1-p_i}{1-p_{i+1}}\lambda_\ell(i+1) - \lambda_h(i) - \lambda_\ell(i)\right]$$

Including the expression on the right-hand side into our objective means that we can separate the ex ante second-best problem type-wise where the multipliers link the interim problems. For symmetry, also include a dummy constraint for $i = 0$ and $i = n$, where $\mu_0 = 0$ and $\mu_n = 0$. These sub-problems now resembles the general problem outlined in section 3.1.1:

$$\begin{array}{rccccccccc} v(i,\beta) = \max\limits_{\lambda(i)\geq 0} & -c\lambda_a(i) & + & r_h(i)(1-\beta)\lambda_h(i) & + & r_\ell(i)(1-\beta)\lambda_\ell(i) & & & & \\ \text{s.t.} & \lambda_a(i) & + & \lambda_h(i) & + & \lambda_\ell(i) & + & \lambda_0(i) & = & \frac{1}{1-\beta} \\ & \lambda_a(i) & & & & & & & \leq & 1 \\ & -\frac{p_i\beta}{1-\beta}\lambda_a(i) & + & \lambda_h(i) & & & & & \leq & 0 \\ & -\frac{(1-p_i)\beta}{1-\beta}\lambda_a(i) & & & + & \lambda_\ell(i) & & & \leq & 0 \end{array}$$

where:

$$r_h(i) = 1 + \mu_{i-1}\frac{p_{i-1}}{p_i} - \mu_i$$

and

$$r_\ell(i) = -1 + \mu_{i-1}\frac{1-p_{i-1}}{1-p_i} - \mu_i$$

The solution is given by $\lambda^\star$, where:

1. if $r_h, r_\ell \geq 0$, then $\lambda_a^\star = \mathbb{1}\{-c + r_h p\beta + r_\ell(1-p)\beta \geq 0\}$, $\lambda_h^\star = \frac{p\beta}{1-\beta}\lambda_a^\star$ and $\lambda_\ell^\star = \frac{(1-p)\beta}{1-\beta}\lambda_a^\star$,
2. if $r_h \geq 0 \geq r_\ell$, then $\lambda_a^\star = \mathbb{1}\{-c + r_h p\beta \geq 0\}$, $\lambda_h^\star = \frac{p\beta}{1-\beta}\lambda_a^\star$ and $\lambda_\ell^\star = 0$,
3. if $r_\ell \geq 0 \geq r_h$, then $\lambda_a^\star = \mathbb{1}\{-c + r_\ell(1-p)\beta \geq 0\}$, $\lambda_h^\star = 0$ and $\lambda_\ell^\star = \frac{(1-p)\beta}{1-\beta}\lambda_a^\star$, and
4. if $0 \geq r_h, r_\ell$, then $\lambda_h^\star = 0$, $\lambda_\ell^\star = 0$ and $\lambda_a^\star = 0$.

The task is then to find multipliers that minimize this expression. We will do this directly with the dual in the next section, however note that we know from 2 that the solution must be represented by a decomposable index. As a result, any analogous problems, such as the interim version of the multi-agent setting, must also be described by a decomposable index.

3.2.2 Optimal mechanism

We can now show that there are only three types of solution to our problem. For reference, let's label these as follows:

- *Full ideal inspection*: Open any box i and keep only if i has a high draw,

$$\lambda_a(i) = 1, \lambda_{ah}(i) = p_i\frac{\beta}{1-\beta} \text{ and } \lambda_{a\ell}(i) = 0 \text{ for all } i$$

- *Partial separation*: Given a threshold t, open any box i that is above t and keep only if i has a high draw, and open box any i below t with probability p_t and unconditionally allocate,

$$\lambda_a(i) = 1, \lambda_{ah}(i) = p_i\frac{\beta}{1-\beta} \text{ and } \lambda_{a\ell}(i) = 0 \text{ for all } i > t$$

$$\lambda_a(i) = p_t, \lambda_{ah}(i) = p_t p_i\frac{\beta}{1-\beta} \text{ and } \lambda_{a\ell}(i) = p_t(1-p_i)\frac{\beta}{1-\beta} \text{ for all } i \leq t$$

- *No allocation*: Never open any box i,

$$\lambda_a(i) = \lambda_{ah}(i) = \lambda_{a\ell}(i) = 0 \text{ for all } i$$

Then the optimal mechanism is given by the following theorem.

Theorem 3 *Full ideal inspection is optimal if all types are sufficiently high, and no allocation is optimal if all types are sufficiently low. That is,*

- *if* $p_i\beta - c \geq 0$ *for all* i*, full ideal inspection is optimal, and*
- *if* $p_i\beta - c \leq 0$ *for all* i*, no allocation is optimal.*

Partial separation is optimal if the first best treats different types differently and if the value of partial separation is positive. That is, if $p_1\beta - c < 0$ *and* $p_n\beta - c > 0$*, and* $V_t \geq 0$*, partial separation is optimal where* t *is set to maximise*

$$V_t = \sum_{i \leq t} \pi_i p_t (m_i\beta - c) + \sum_{i > t} \pi_i (p_i\beta - c)$$

Finally, no allocation is optimal if the first best treats different types differently and if the value of partial separation is negative.

We prove this by first changing variables into a convenient representation, and then showing there is an equivalent solution for the dual in each case. This is done in appendix B.2 and B.3 and the first part is listed as proposition 1 and the second and third as proposition 2.

This first part of the theorem should be unsurprising: if the first best is incentive compatible, then the first best is achievable. The interesting case is when types fall on either side of this threshold and the first best isn't achievable, which is described by the second and third part of the theorem. Note that partial separation generalises full ideal inspection, and so in an operational sense, there are only two policy's to evaluate and compare: partial separation and no allocation.

The second-best payoffs and policy are summarised in figures 7 and 8 respectfully for a distribution where the first best treats different types differently and if the value of partial separation is positive. The underlying distribution for these representations is inherently continuous (though it need not be), so let $\tau^* \coloneqq p_{t^*}$.

Note that, with the additional assumption that the principal must inspect before they allocate and the restriction to binary prizes, this reproves the main result of Khalfan (2023). The mechanism also aligns with Ben-Porath, Dekel and Lipman (2014) as the noise in the signal tends to zero with the exception that the return to the principal here is reduced by the cost of inspecting types that do not meet the threshold.

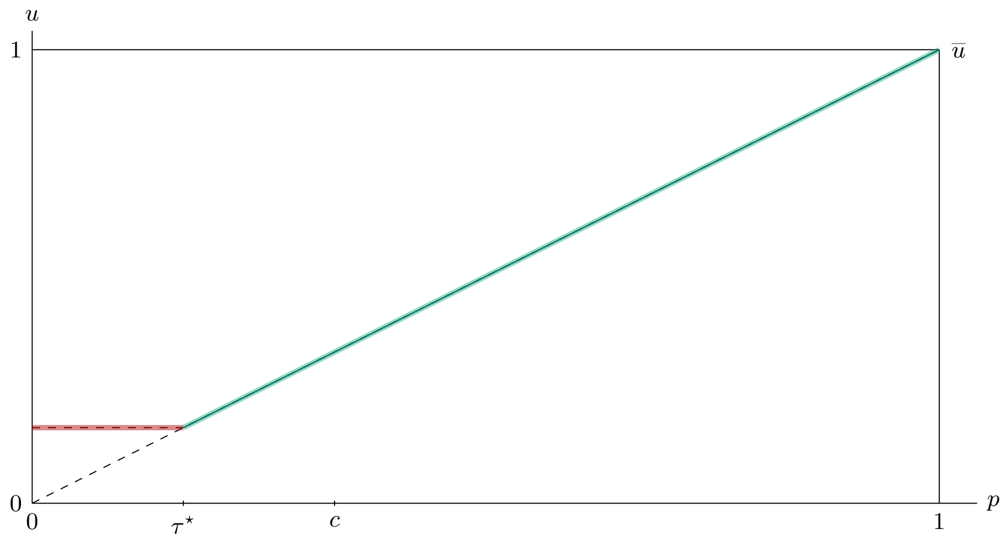

Figure 7: second-best payoffs, u^*, as a function of high prize probability, p, given a cost, $c > 0$

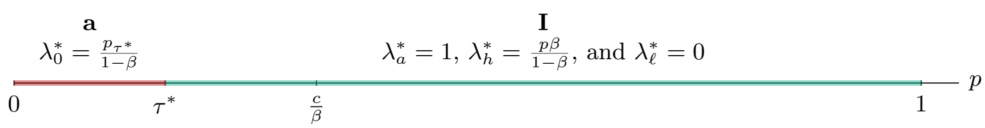

Figure 8: second-best policy, $(\lambda_a^*, \lambda_h^*, \lambda_\ell^*, \lambda_0^*)$

4 Multi-agent environment

Now suppose there are n agents, indexed $i \in \{1, \ldots, n\} := N$. Let T be the set of possible types of an agent. Denote by $x_{ij}(t)$ the allocation to agent $i \in N$ at state $j \in \Omega$ at the profile $t = (t_1, \ldots, t_n)$ of types. In our case, there is a special state, denoted '0' and the probability of transitioning from state 0 to state $j \in \Omega \setminus \{0\}$ for agent i of type t_i is denoted $p(j|t_i)$. Denote the set of all profiles of types by T^n. In our case, state will be a symbol used to denote whether the box has been opened and, if opened, the reward associated with the box.

4.1 Reduced form

For every profile of types t we have a non-negative, non-decreasing, submodular function $g(\cdot|t)$ defined on subsets of $N \times \Omega$. An allocation rule is ex-post feasible if

$$\sum_{(i,j) \in S} x_{ij}(t) \leq g(S|t) \quad \forall S \subseteq N \times \Omega,\ t \in T^n$$

and

$$x_{ij}(t) - a_j(t_i) x_{i0}(t) \leq 0 \quad \forall i \in N\ t_i \in T,\ j \in \Omega \setminus \{0\}$$

Note that $a_j(t_i)$ does not depend on the entire profile of types. Call the first set *total conservation constraints* as they restrict the total service time of states within a particular subset, and the second *transition conservation constraints* as they only restrict the service time of states in transition.

It will be more convenient to write these constraints by weighting them by the likelihood of the type profile, $\pi(\cdot)$:

$$\sum_{(i,j)\in S} \pi(t)x_{ij}(t) \leq \pi(t)g(S|t) \quad \forall S \subseteq N \times \Omega,\ t \in T^n$$

$$\pi(t)x_{ij}(t) - \pi(t)a_j(t_i)x_{i0}(t) \leq 0 \quad \forall i \in N\ t_i \in T,\ j \in \Omega \setminus \{0\}$$

Let Q denote the interim allocation variables so that:

$$\pi(t_i)Q_{ij}(t_i) = \sum_{t_{-i}} \pi(t_i)\pi(t_{-i})x_{ij}(t_i, t_{-i})$$

Given any $S \subseteq N \times \Omega$ let S_N be the 'projection' of S into N. In words the set of agents associated with S.

The feasible Qs must satisfy the interim version of these two sets of constraints.

Consider first the total conservation constraints. For a particular subset S and profile t:

$$\sum_{(i,j)\in S} x_{ij}(t) \leq g(S|t)$$

Weigh this by the likelihood of profile t and sum over all profiles:

$$\sum_{(i,j)\in S} \sum_{t} \pi(t)x_{ij}(t) \leq \sum_{t} \pi(t)g(S|t).$$

An important feature of $g(S|t_{S_N}, t_{-S_N})$ in our application is that it is independent of t_{-S_N}, i.e. $g(S|t_{S_N}, t_{-S_N}) = g(S|t_{S_N})$. Replacing for our interim variables and simplifying the right-hand side:

$$\sum_{(i,j)\in S} \sum_{t_i} \pi(t_i)Q_{ij}(t_i) \leq \sum_{t_{S_N}} \pi(t_{S_N})g(S|t_{S_N})$$

The class of submodular functions is closed under non-negative linear combinations, so the right-hand side is still submodular. Let $G(S) = \sum_{t_{S_N}} \pi(t_{S_N})g(S|t_{S_N})$. We then have our *interim total conservation constraint*:

$$\sum_{(i,j)\in S} \sum_{t_i} \pi(t_i)Q_{ij}(t_i) \leq G(S) \quad \forall S \subseteq N \times \Omega \tag{1}$$

The *interim transition conservation constraint* are more straightforward to derive, simply weighting by the likelihood of profile t:

$$\pi(t_i)Q_{ij}(t_i) - \pi(t_i)a_j(t_i)Q_{i0}(t_i) \leq 0 \quad \forall i \in N \ t_i \in T, \ j \in \Omega \setminus \{0\} \tag{2}$$

$$\Rightarrow Q_{ij}(t_i) - a_j(t_i)Q_{i0}(t_i) \leq 0$$

Bayesian incentive compatibility for each agent i with type t_i is also straightforward:

$$\sum_{j \in \Omega \setminus \{0\}} Q_{ij}(t_i) \geq \sum_{j \in \Omega \setminus \{0\}} \frac{p(j|t_i)}{p(j|t_i')} Q_{ij}(t_i') \quad \forall t_i'. \tag{3}$$

Remember, that while all feasible Qs must satisfy these three constraints, it is not necessarily true that all Qs that satisfy them are feasible. That is, there may not exist an ex-post feasible x that generate any particular Q, and so must be verified later. Further, it is not guaranteed that there exists an ex-post feasible x that generates Q that is additionally dominant strategy incentive compatible, and must be verified (or disproven) later.

The ex ante objective function written with ex post variables is the following:

$$V = \sum_{t \in T^N} \left[\sum_{j \in \Omega} \sum_{i \in N} r_j x_{ij}(t) \right] \pi(t) = \sum_i \sum_j r_j \left[\sum_t x_{ij}(t)\pi(t) \right]$$

We can then split the type space to get a familiar interim form:

$$V = \sum_i \sum_j r_j \left[\sum_{t_i} \sum_{t_{-i}} x_{ij}(t)\pi(t_i)\pi(t_{-i}) \right] = \sum_i \sum_j r_j \left[\sum_{t_i} Q_{ij}(t_i)\pi(t_i) \right]$$

As such we have an interim version of the objective function which is just the sum of interim values over agents and their types:

$$V = \sum_i \sum_{t_i} \left[\sum_j r_j Q_{ij}(t_i) \right] \pi(t_i)$$

In our binary setup, we set $p(2|t_i) = t_i$ and $p(1|t_i) = 1 - t_i$. Thus, type is the probability distribution over states. Then let $r_0, r_1 < 0$ and $r_2 > 0$.

Inequality (2) becomes:

$$\pi(t_i)Q_{ij}(t_i) - \pi(t_i)\frac{\beta p(j|t_i)}{1-\beta}Q_{i0}(t_i) \leq 0 \quad \forall i \in N \ t_i \in T, \ j \in \{1, 2\}$$

Inequality (3) reduces to

$$Q_{i1}(t_i) + Q_{i2}(t_i) \geq \frac{(1-t_i)}{(1-t'_i)} Q_{i1}(t'_i) + \frac{t_i}{t'_i} Q_{i2}(t'_i) \quad \forall t'_i.$$

The usual argument tells us that IC implies:

$$\frac{Q_{i2}(t_i)}{t_i} - \frac{Q_{i1}(t_i)}{1-t_i} \geq \frac{Q_{i2}(t'_i)}{t'_i} - \frac{Q_{i1}(t'_i)}{1-t'_i} \quad \forall t_i \geq t'_i \tag{4}$$

The converse need not be true. We can argue that the adjacent upward IC and constraint (4) suffice just as in the single-agent setting.

Suppose the rewards at each of the states are r_0, r_1 and r_2 and are agent independent. Suppose the probability that an agent is in state 2 can be $\{0, \frac{1}{m}, \frac{2}{m}, \ldots, \frac{m-1}{m}, 1\}$. Then, a generic upward adjacent IC constraint for agent i with type $k < m$ is

$$Q_{i1}(k) + Q_{i2}(k) \geq \frac{(1-\frac{k}{m})}{(1-\frac{k+1}{m})} Q_{i1}(k+1) + \frac{\frac{k}{m}}{\frac{k+1}{m}} Q_{i2}(k+1).$$

$$\Rightarrow Q_{i1}(k) + Q_{i2}(k) \geq \frac{m-k}{m-k-1} Q_{i1}(k+1) + \frac{k}{k+1} Q_{i2}(k+1).$$

When $k+1 = m$, then, $\frac{m-k}{m-k-1}$ is undefined. However, this corresponds to the case when the probability of state 1 is zero, so assume $Q_{i1}(m) = 0$. And similarly, monotonicity

$$\frac{Q_{i2}(k+1)}{\frac{k+1}{m}} - \frac{Q_{i1}(k+1)}{1-\frac{k+1}{m}} \geq \frac{Q_{i2}(k)}{\frac{k}{m}} - \frac{Q_{i1}(k)}{1-\frac{k}{m}}$$

$$\Rightarrow \frac{Q_{i2}(k+1)}{k+1} - \frac{Q_{i1}(k+1)}{m-k-1} \geq \frac{Q_{i2}(k)}{k} - \frac{Q_{i1}(k)}{m-k}$$

4.2 Lagrangian

Before continuing, let's recall the Lagrangian approach. Suppose we wish to solve $Z = \max\{cx : \text{ s.t. } Ax \leq b, Cx \leq d, x \geq 0\}$. Let $\mu \geq 0$ be the shadow price associated with $Cx \leq d$. Then

$$L(\mu) = \max\{cx + \mu(d - Cx) : \text{ s.t. } Ax \leq b, x \geq 0\}.$$

Further, $Z = \min_{\mu \geq 0} L(\mu)$. Let $\mu_i(k, k+1) \geq 0$ be the shadow price associated with the upward adjacent IC constraint. Note that we will have variables of this kind for $0 \leq k \leq m-1$ only. The objective function we are maximizing will be

$$\sum_i \sum_{t_i} \pi(t_i)[r_0 Q_{i0}(t_i) + r_1 Q_{i1}(t_i) + r_2 Q_{i2}(t_i)].$$

If we take the IC constraint into the objective function, we will be introducing terms of the following kind into the objective function:

$$\mu_i(k,k+1)[-\frac{m-k}{m-k-1}Q_{i1}(k+1) - \frac{k}{k+1}Q_{i2}(k+1) + Q_{i1}(k) + Q_{i2}(k)]$$

This gives us an objective,

$$\begin{aligned} L(\mu,Q) = &\sum_i \sum_{k=0}^{m} \pi(k)[r_0 + r_1 Q_{i1}(k) + r_2 Q_{i2}(k)] \\ &+ \sum_i \sum_{k=0}^{m-1} \mu_i(k,k+1)[Q_{i1}(k) + Q_{i2}(k) - \frac{m-k}{m-k-1}Q_{i1}(k+1) - \frac{k}{k+1}Q_{i2}(k+1)] \end{aligned}$$

Setting $Q_{i1}(m) = 0$ and collecting terms, the objective becomes:

$$\begin{aligned} L(\mu,Q) = \sum_i \langle &\pi(0)r_0 Q_{i0}(0) + [\pi(0)r_1 + \mu_i(0,1)]Q_{i1}(0) + [\pi(0)r_2 + \mu_i(0,1)]Q_{i2}(0) \\ &+ \sum_{k=1}^{m-1} \{\pi(k)r_0 Q_{i0}(k) + [\pi(k)r_1 + \mu_i(k,k+1) - \frac{m-k+1}{m-k}\mu_i(k-1,k)]Q_{i1}(k) \\ &\qquad + [\pi(k)r_2 + \mu_i(k,k+1) - \frac{k-1}{k}\mu_i(k-1,k)]Q_{i2}(k)\} \\ &+ \pi(m)r_0 Q_{i0}(m) + [\pi(0)r_2 - \frac{m-1}{m}\mu_i(m-1,m)]Q_{i2}(m)\rangle \end{aligned}$$

4.3 Threshold conjecture

As we can see, the structure of the reduced form mirrors the single-agent setting and as such, it's reasonable to guess that it too must follow a threshold structure. This gives us the following conjecture.

Conjecture 1 *The optimal mechanism sets a threshold t, such that for each agent i,*

- *if $k \geq t$, $Q_{i0}(k) = Q_{i0}(N)$, $Q_{i1}(k) = 0$ and $Q_{i2}(k) = \frac{\beta \frac{k}{m}}{1-\beta}Q_{i0}(k)$, and*
- *if $k < t$, $Q_{i0}(k) = Q_{i0}(0)$, $Q_{i1}(k) = \frac{\beta(1-\frac{k}{m})}{1-\beta}Q_{i0}(k)$ and $Q_{i2}(k) = \frac{\beta \frac{k}{m}}{1-\beta}Q_{i0}(k)$.*

That is, *high* types are inspected maximally and allocated to efficiently, and low types are compensated with a small probability of inspection and unconditional allocation that makes the marginal type indifferent. The task is then to solve for $Q_{i0}(N)$ and $Q_{i0}(0)$, show that this solution also minimizes the analogous dual problem, and verify that there is an ex post mechanism that implements the conjecture - also presumably a threshold mechanism. We leave this task for future versions of the paper.

5 Summary

When a principal wishes to allocate a scarce resource among privately informed, rival users without transfers, they must balance the discovery of new information via inspection with the verification of private information. Optimally, they do this by at most partially exploiting private information to guide search. In particular, the principal *over-inspects* high and low types, *under-allocates* to agents who are worthy of inspection, and *over-allocates* to agents who are not.

This is related to a branch of the mechanism design literature devoted to costly inspection without transfers and demonstrates how we can recover interesting observations about search behaviour with noise, a feature mostly missing from the branch. With the additional assumption that an agent must be inspected prior to selection, this paper generalises Khalfan (2023) to multiple agents and Ben-Porath, Dekel and Lipman (2014) to imperfect signals.

A feature of the setup that may seem restrictive is that rewards are binary. Recall that in optimal auctions - considering Myerson (1981) in particular - the type space is typically single dimensional. Here, if the reward space is expanded, our types are now a a distribution over many rewards so may be many dimensional. As such, some structure must be imposed to give us a uni-dimensional order, essentially reducing the problem back to the analogous setup here. For example, Khalfan (2023) demonstrates that when rewards are general and signal distributions are ordered by the monotone likelihood ratio property, the same result applies with an additional threshold with respect to the rewards to determine which rewards are *high* and which are *low*.

A second feature is that the principal must inspect the agent in order to allocate and receive the value. There are several reasons we are comfortable studying this restriction. Firstly, in many institutional settings, the principal is required to complete this step prior to allocation, i.e. *due diligence*, and so is a realistic feature in and of itself. Secondly, this assumption is consistent with, and a necessary feature of, the treatment in Bertsimas and Niño-Mora (1996) - and by extension Weitzman (1979). This assumption makes Pandora's problem tractable and allows for the technical treatment we have presented. For the complications that arise with this problem when there is not necessary inspection, see Doval (2018). Finally, in comparing the result here to Khalfan (2023) and Ben-Porath, Dekel and Lipman (2014), we see that the structure of the solution is the same with the caveat that it is more costly to allocate to low type agents. As such, one could then take the results here, increase the value to the principal by taking away this requirement for inspecting low agents, and be confident that the resulting mechanism is likely to also be optimal in this new environment.

Future versions of this paper will evaluate the multi-agent conjecture and confirm that it's ex post implementable.

References

Ben-Porath, Elchanan, Eddie Dekel and Barton L. Lipman. 2014. "Optimal Allocation with Costly Verification." *American Economic Review* 104(12):3779–3813.

Bertsimas, Dimitris and José Niño-Mora. 1996. "Conservation Laws, Extended Polymatroids and Multiarmed Bandit Problems; A Polyhedral Approach to Indexable Systems." *Mathematics of Operations Research* 21(2):257–306.

Border, Kim C. 1991. "Implementation of Reduced Form Auctions: A Geometric Approach." *Econometrica* 59(4):1175–1187.

Doval, Laura. 2018. "Whether or not to open Pandora's box." *Journal of Economic Theory* 175:127–158.

Gittins, J. C. and D. M. Jones. 1974. *A Dynamic Allocation Index for the Sequential Design of Experiments*. Amsterdam: North-Holland pp. 241–266.

Khalfan, Nawaaz. 2023. "Optimal Allocation with Noisy Inspection.".

Myerson, Roger B. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6(1):58–73.

Weitzman, M. 1979. "Optimal Search for the Best Alternative." *Econometrica* 47(3):641–654.

A Pandora's linear program

In this section of the appendix, the details for the linear program translation are expanded upon. In particular, section A.1 steps through the conservation laws, section A.2 proves they're sufficient, and A.3 completes the proof.

A.1 Conservation laws

To characterize the feasible space of service times, Λ, we first derive a set of conservation laws regarding the service times to identify inequalities that any job must satisfy. Subsequently, we are going to show that these inequalities are not only necessary but sufficient for describing Λ.

First, with the inclusion of the outside option, the scheduler never stands to strictly benefit from being idle. As such, one job must be serviced in each time period:

$$\sum_{j \in \mathcal{J}} \lambda_j^u = \frac{1}{1-\beta} \quad \forall u \in \mathcal{U} \tag{1}$$

Second, let u_S for $S \subseteq \mathcal{J}$ denote the policy that always services a job in class $j \in S$ whenever one is available. This policy generates an expression about the maximum possible service time and as such gives us an upper bound for all admissible policies, $u \in \mathcal{U}$, as they pertain to classes in S.

If $S = j \in \Omega$, then the job is an unopened box and can thus be serviced at most once, so:

$$\lambda_j^{u_j} = 1 \quad \Rightarrow \quad \lambda_j^u \leq 1 \quad \forall u \in \mathcal{U} \tag{2.1}$$

If $S = j \in \mathcal{N}$, then the job is an opened box with associated prize x_j. This will only be available when an unopened box reveals such a prize, and can then be selected. In the scheduling problem, this is equivalent to being infinitely serviced from then on. As such:

$$\lambda_j^{u_j} = \sum_{i \in \Omega} \lambda_i^{u_j} p_{ij} \frac{\beta}{1-\beta} \quad \Rightarrow \quad \lambda_j^u - \sum_{i \in \Omega} \lambda_i^u p_{ij} \frac{\beta}{1-\beta} \leq 0 \quad \forall u \in \mathcal{U} \tag{2.2}$$

To see why, note we'd like to calculate the service time for j under u_j which involves the likelihood of reaching j from any unopened box i. Suppose under u_j, the service time for i, $\lambda_i^{u_j}$, can be decomposed into a stream of service probabilities, α_t^i. That is, $\lambda_i^{u_j} = \alpha_0^i + \alpha_1^i \beta + \ldots + \alpha_t^i \beta^t + \ldots$.

Then we can write $\lambda_j^{u_j}$ as:

$$\begin{aligned}
\lambda_j^{u_j} &= \sum_{i\in\Omega}\big[\alpha_0^i p_{ij}(\beta+\beta^2+\ldots)+\alpha_1^i p_{ij}(\beta^2+\beta^3+\ldots)+\alpha_2^i p_{ij}(\beta^3+\beta^4+\ldots)+\ldots\ \big] \\
&= \sum_{i\in\Omega}\big[\alpha_0^i p_{ij}\beta\frac{1}{1-\beta}+\alpha_1^i p_{ij}\beta^2\frac{1}{1-\beta}+\alpha_2^i p_{ij}\beta^3\frac{1}{1-\beta}+\ldots\ \big] \\
&= \sum_{i\in\Omega}\lambda_i^{u_j} p_{ij}\frac{\beta}{1-\beta}
\end{aligned}$$

Finally, for the singleton sets, if $S = j = 0$ then the job is the outside option which is available from the start and can be serviced immediately and indefinitely. Then:

$$\lambda_j^{u_j} = \frac{1}{1-\beta} \quad \Rightarrow \quad \lambda_j^{u} \le \frac{1}{1-\beta} \quad \forall u \in \mathcal{U} \tag{2.3}$$

With inequalities 2.1, 2.2, and 2.3, we can build conservation laws for all subsets, S, providing upper bounds on the total service times for jobs in those subsets. Conveniently, these subsets fall into only two interesting cases.

In the first, if $0 \in S$, then the subset of jobs can be serviced immediately and indefinitely by servicing the outside option. This gives us the following, mostly redundant, inequality:

$$\sum_{j\in S}\lambda_j^{u_S} = \frac{1}{1-\beta} \quad \Rightarrow \quad \sum_{j\in S}\lambda_j^{u} \le \frac{1}{1-\beta} \quad \forall u \in \mathcal{U} \tag{3.1}$$

For the second, suppose $0 \notin S$. Then S can be partitioned into two sets: $S^\Omega := S \cap \Omega$, and $S^\mathcal{N} := S \cap \mathcal{N}$. As before, we'd like to find the maximum possible total service time. The following priority policy intuitively achieves this maximum:

1. Service all jobs in S^Ω.
2. If any job from S^Ω has transitioned into a class from $S^\mathcal{N}$, service this job indefinitely.
3. If no job from S^Ω has transitioned into a class from $S^\mathcal{N}$, then service remaining jobs in $\Omega \backslash S^\Omega$, until a job transitions into a class from $S^\mathcal{N}$ and then service this job indefinitely.

We can evaluate this policy as:

$$\begin{aligned}
\sum_{j\in S}\lambda_j^{u_S} = \frac{1-\beta^{|S^\Omega|}}{1-\beta} &+ \beta^{|S^\Omega|}\Big[1-\prod_{i\in S^\Omega}(1-\sum_{j\in S^\mathcal{N}} p_{ij})\Big]\Big[\frac{1}{1-\beta}\Big] \\
&+ \Big[\prod_{i\in S^\Omega}(1-\sum_{j\in S^\mathcal{N}} p_{ij})\Big]\Big[\sum_{i\in\Omega\backslash S^\Omega}\lambda_i^{u_S}\sum_{j\in S^\mathcal{N}} p_{ij}\frac{\beta}{1-\beta}\Big]
\end{aligned}$$

The first expression is the discounted value of the first $|S^\Omega|$ time steps, achieved by servicing the unopened boxes in S. Simplifying this expression:

$$\sum_{j\in S}\lambda_j^{u_S} = \frac{1}{1-\beta}\left[1-\beta^{|S^\Omega|}\Big[\prod_{i\in S^\Omega}(1-\sum_{j\in S^\mathcal{N}} p_{ij})\Big]\Big[1-\beta\sum_{i\in\Omega\setminus S^\Omega}\lambda_i^{u_S}\sum_{j\in S^\mathcal{N}} p_{ij}\Big]\right]$$

Rearranging to follow convention that choice variables are collected on the left-hand side of the constraints, we have our final conservation law:

$$\sum_{j\in S}\lambda_j^u-\beta^{|S^\Omega|}\Big[\prod_{i\in S^\Omega}(1-\sum_{j\in S^\mathcal{N}} p_{ij})\Big]\Big[\sum_{i\in\Omega\setminus S^\Omega}\lambda_i^u\sum_{j\in S^\mathcal{N}} p_{ij}\frac{\beta}{1-\beta}\Big]\leq\frac{1}{1-\beta}\Big[1-\beta^{|S^\Omega|}\Big[\prod_{i\in S^\Omega}(1-\sum_{j\in S^\mathcal{N}} p_{ij})\Big]\Big]\quad\forall u\in\mathcal{U} \tag{3.2}$$

Inequalities 3.1 and 3.2 generalise the singleton set expressions, and as such, along with 1, they define the *maximal service time region*, $\mathcal{L}$:

$$\mathcal{L} := \{\lambda\in\mathbb{R}_+^{|\mathcal{J}|}\mid\lambda \text{ satisfy } 1,\ 3.1,\text{ and } 3.2\}$$

A.2 Sufficiency

Bertsimas and Niño-Mora (1996) show that if the feasible space of service times satisfy *generalized conservation laws*, then they can be represented as an *extended polymatroid.* This means that optimizing a linear objective over the feasible space will not only yield an extreme solution, but the extreme solution is characterized by a priority-index over classes. Conditional on one further assumption on the structure of the feasible space, the index is also decomposable. These are the subject of Bertsimas and Niño-Mora (1996)'s Theorem 1, 2, and 3.

The scheduling problem presented here is a special case of the classic *multiarmed bandit problem.* In particular, it is a multiarm bandit with one intransient arm, *the outside option*, and N transient arms, *the unopened boxes.* The transient arms are available from the start, and stochastically map into payoff-varied intransient arms, *the opened boxes.* The multiarmed bandit problem is shown to be a satisfy Bertsimas and Niño-Mora's generalized conservation laws in Proposition 8 and Theorem 11. Given our problem is simpler than the general multiarm bandit problem, these results are refined and restated in the following theorem.

Theorem 1 *The performance region, Λ, is the extended polymatroid defined by the maximal service region, $\mathcal{L}$. Further, for each class $j\in\mathcal{J}$ there exists indices, z, depending only on characteristics of that class, such that an optimal policy is to schedule a job with the largest current index.*

Proof: See Bertsimas and Niño-Mora (1996), Proposition 8 and Theorem 11. □

Note that, in its original generality, this reproves the celebrated result of Gittins and Jones (1974): *there exists a decomposable index that dictates the optimal scheduling of projects in the multiarmed bandit problem.* For our problem, this proves that the schedule's solution is also characterized by a decomposable priority-index, which we can directly take advantage of.

A.3 Pandora's schedule

Making the most of these properties, we can explicitly solve for the optimal schedule.

Theorem 2 *The solution to Pandora's schedule is given by an indexing rule defined by an index, z, where if $j \in \mathcal{N} \cup \{0\}$ then $z_j = r_j$, and, if $j \in \Omega$, then z_j is the solution to:*

$$-r_j = \sum_{\substack{i \in \mathcal{N} \\ r_i > z_j}} p_{ji}(r_i - z_j)$$

Proof: As outlined in section A.2, Pandora's scheduling problem satisfies Bertsimas and Niño-Mora's generalized conservation laws and so the feasible space of achievable performance is an extended polymatroid. Optimization of a linear objective over an extended polymatroid is solved by an adaptive greedy algorithm, which leads to an optimal solution having the indexability property. Further, this extends to the decomposability property, which means the indices must apply to all sub-problems.

Normalize the index for the outside option as the value of the outside option itself: $z_0 = r_0$. Now, if $j \in \mathcal{N} \cup \{0\}$, then $z_j > z_0$ if and only if $r_j > r_0$ as, if $\mathcal{J} = \{0, j\}$, then Pandora must service the highest prize thereafter for her schedule to be optimal. As r_0 could take any value, then it must be that, for all $j \in \mathcal{N} \cup \{0\}$, $z_j = r_j$.

Now suppose $j \in \Omega$ and again consider $\mathcal{J} = \{0, j\}$. If $z_j = z_0$, then Pandora must be indifferent between servicing j and servicing the outside option indefinitely:

$$r_j + \sum_{r_i > r_0,\ i \in \mathcal{N}} p_{ji} r_i + \sum_{r_i \leq r_0,\ i \in \mathcal{N}} p_{ji} r_0 = r_0$$

where the continuation value after servicing j is fixed by the indexibility of the solution and already pinned down values of z_j for all $j \in \mathcal{N} \cup \{0\}$. Rearranging and substituting $z_j = z_0$ and $z_0 = r_0$:

$$-r_j = \sum_{r_i > z_j,\ i \in \mathcal{N}} p_{ji}(r_i - z_j)$$

Then, we've pinned down the value of the index, modulo the normalized index for the outside option. ∎

Pandora's rule is then given by the limiting behaviour of the optimal schedule, referred to as **Pandora's schedule**.

Corollary 2 *As $\beta \to 1$, Pandora's schedule is precisely Pandora's rule.*

B Single-agent environment

B.1 Local incentive compatibility

Firstly, we have $n(n-1)$ constraints to obey. The following lemmas reduce the number of constraints to track.

Lemma 1 *Incentive compatibility implies the differential is monotone increasing. That is, IC_{pq} and IC_{qp} imply $\Delta(p) \geq \Delta(q)$ for all $p > q$.*

Proof: Rearranging IC_{pq} and IC_{qp}:

$$IC_{pq}: \; y_\ell(p) + p\Delta(p) \geq y_\ell(q) + p\Delta(q)$$
$$IC_{qp}: \; y_\ell(q) + q\Delta(q) \geq y_\ell(p) + q\Delta(p)$$

Subtracting the right hand side from IC_{qp} from the left hand side from IC_{pq} and vice versa gives us the following inequality:

$$p\Delta(p) - q\Delta(p) \geq p\Delta(q) - q\Delta(q)$$

And finally rearranging:

$$[p-q] \cdot [\Delta(p) - \Delta(q)] \geq 0$$

Thus, if $p > q$ it must be that $\Delta(p) \geq \Delta(q)$ and vice versa. □

Lemma 2 *Local upward (downward) incentive compatibility and differential monotonicity implies global upward (downward) incentive compatibility. That is, IC_{pq}, IC_{qr} and $\Delta(p)$ increasing implies IC_{pr} for all $p < q < r$ (or $p > q > r$).*

Proof: Collecting the guaranteed allocations and the differential in IC_{pq}, IC_{qr} gives us:

$$IC_{pq}: \; y_\ell(p) - y_\ell(q) \geq p[\Delta(q) - \Delta(p)]$$
$$IC_{qr}: \; y_\ell(q) - y_\ell(r) \geq q[\Delta(r) - \Delta(q)]$$

Adding the left hand and right hand sides gives us the following inequality:

$$y_\ell(p) - y_\ell(r) \geq p[\Delta(q) - \Delta(p)] + q[\Delta(r) - \Delta(q)]$$

Adding and subtracting $p\Delta(r)$ to the right hand side and rearranging:

$$y_\ell(p) - y_\ell(r) \geq p[\Delta(r) - \Delta(p)] + [q - p][\Delta(r) - \Delta(q)]$$

The final term is positive when $p < q < r$ (or when $p > q > r$) as Δ is monotone. Then it must be true that:

$$y_\ell(p) - y_\ell(r) \geq p[\Delta(r) - \Delta(p)]$$

Which is precisely the collected form of IC_{pr}. □

Lemma 3 *Binding upward (downward) incentive compatibility and differential monotonicty implies downward (upward) incentive compatibility. That is, IC_{pq} binds and $\Delta(p)$ increasing implies IC_{qp} for all $p < q$ ($p > q$).*

Proof: The collected representation of IC_{pq} is:

$$IC_{pq}: \; y_\ell(p) - y_\ell(q) = p[\Delta(q) - \Delta(p)]$$

If $q > p$, then $\Delta(q) - \Delta(p) \geq 0$ as Δ is monotone, (and if $q < p$, then $\Delta(q) - \Delta(p) \leq 0$) so replacing p with q implies:

$$y_\ell(p) - y_\ell(q) \leq q[\Delta(q) - \Delta(p)]$$
$$y_\ell(q) - y_\ell(p) \geq q[\Delta(p) - \Delta(q)]$$

Which is precisely the collected representation of IC_{qp}. □

So far, these lemmas show that local incentive compatibility and the allocation differential matter. Order and index $\mathcal{P}$ so that $p_i > p_j$ if $i < j$. We will also promote the index i to the argument where it does not cause confusion.

Lemma 4 *Local upward incentive compatibility binds under the second-best policy. That is, if $\lambda = \lambda^*$, then $IC_{i,i+1}$ binds for every $i < n$.*

Proof: Let $\underline{p} = \frac{c}{\beta}$, and recall that if $p \geq \underline{p}$ then Pandora's value increases if λ_a and λ_h are raised at a ratio of $1 : \frac{p\beta}{1-\beta}$, and if $p < \underline{p}$, then Pandora's value increases if λ_a and λ_h are proportionately lowered. Also note, that under $\lambda = \lambda^*$ for any p, it cannot be that $\frac{\beta}{1-\beta}\lambda_a > \max\{\frac{\lambda_h}{p}, \frac{\lambda_\ell}{1-p}\}$, as then lowering λ_a directly increases Pandora's value without violating any constraints.

There are two cases to consider: $p_i < \underline{p}$ and $p_i \geq \underline{p}$.

Suppose $p_i < \underline{p}$, and $IC_{i,i+1}$ does not bind. Then:

$$IC_{i-1,i} : \lambda_h(i-1) + \lambda_\ell(i-1) \geq \frac{p_{i-1}}{p_i}\lambda_h(i) + \frac{1-p_{i-1}}{1-p_i}\lambda_\ell(i)$$
$$IC_{i,i+1} : \lambda_h(i) + \lambda_\ell(i) > \frac{p_i}{p_{i+1}}\lambda_h(i+1) + \frac{1-p_i}{1-p_{i+1}}\lambda_\ell(i+1)$$

Now, lowering $\lambda_\ell(i)$ increases Pandora's value, tightens $IC_{i,i+1}$ and relaxes $IC_{i-1,i}$ (if it exists), a contradiction. Then, either $IC_{i,i+1}$ binds or $\lambda_\ell(i) = 0$.

If $\lambda_\ell(i) = 0$, then lowering $\lambda_a(i)$ and $\lambda_h(i)$ by a ratio of $1 : \frac{p\beta}{1-\beta}$ increases Pandora's value, tightens $IC_{i,i+1}$ and relaxes $IC_{i-1,i}$ (if it exists), a contradiction. Then, either $IC_{i,i+1}$ binds or $\lambda_a(i) = \lambda_h(i) = 0$.

If $\lambda_a(i) = \lambda_h(i) = \lambda_\ell(i) = 0$, then $u(i) = 0$ which is the lower bound for payoffs and thus cannot strictly exceed $u(i+1|i)$, a contradiction. So $IC_{i,i+1}$ must bind if $p_i < \underline{p}$.

Suppose $p_i \geq \underline{p}$, and $IC_{i,i+1}$ does not bind. Then:

$$IC_{i-1,i} : \lambda_h(i-1) + \lambda_\ell(i-1) \geq \frac{p_{i-1}}{p_i}\lambda_h(i) + \frac{1-p_{i-1}}{1-p_i}\lambda_\ell(i)$$
$$IC_{i,i+1} : \lambda_h(i) + \lambda_\ell(i) > \frac{p_i}{p_{i+1}}\lambda_h(i+1) + \frac{1-p_i}{1-p_{i+1}}\lambda_\ell(i+1)$$

As before, lowering $\lambda_\ell(i)$ increases Pandora's value, tightens $IC_{i,i+1}$ and relaxes $IC_{i-1,i}$ (if it exists), a contradiction. Then, either $IC_{i,i+1}$ binds or $\lambda_\ell(i) = 0$.

If $\lambda_\ell(i) = 0$, then raising $\lambda_a(i+1)$ and $\lambda_h(i+1)$ by a ratio of $1 : \frac{p_{i+1}\beta}{1-\beta}$ increases Pandora's value, tightens $IC_{i,i+1}$ and relaxes $IC_{i+1,i+2}$ (if it exists), a contradiction. Then, either $IC_{i,i+1}$ binds or $\lambda_a(i+1) = 1$ and $\lambda_h(i+1) = \frac{p_{i+1}\beta}{1-\beta}$.

If $\lambda_\ell(i) = 0$, $\lambda_a(i+1) = 1$ and $\lambda_h(i+1) = \frac{p_{i+1}\beta}{1-\beta}$ then the right-hand side of $IC_{i,i+1}$ is given by:

$$p_i\frac{\beta}{1-\beta} + \frac{1-p_{i+1}}{1-p_{i+2}}\lambda_\ell(i+1)$$

and the left-hand side is $\lambda_h(i)$ which is necessarily less than or equal to $p_i\frac{\beta}{1-\beta}$, a contradiction. So $IC_{i,i+1}$ must bind if $p_i \geq \underline{p}$. □

These lemmas reduce the number of incentive compatibility constraints from $n(n-1)$ inequality constraints, to $n-1$ equality constraints, $IC_{i,i+1}$ binds for all $i < n$, and $n-1$ inequality constraints, $\Delta(i) \leq \Delta(i+1)$ for all $i < n$.

B.2 A change of variables

The pair of inequalities,

$$-\frac{p_i\beta}{1-\beta}\lambda_a(i)+\lambda_h(i)\leq 0$$

$$-\frac{(1-p_i)\beta}{1-\beta}\lambda_a(i)+\lambda_\ell(i)\leq 0$$

describe a cone in $\mathbb{R}^3$.

To find the generators we can focus on the extreme points of:

$$-\frac{p_i\beta}{1-\beta}\lambda_a(i)+\lambda_h(i)+s_1=0$$

$$-\frac{(1-p_i)\beta}{1-\beta}\lambda_a(i)+\lambda_\ell(i)+s_2=0$$

$$\lambda_a(i)+\lambda_h(i)+\lambda_\ell(i)=1$$

$$\lambda_a(i),\lambda_h(i),\lambda_\ell(i),s_1,s_2\geq 0$$

We have three equations and five variables. So, in an extreme point solution, at least two variables must be zero.

1. $s_1=s_2=0$

$$\lambda_a(i)=\frac{1}{\frac{p_i\beta}{1-\beta}+\frac{(1-p_i)\beta}{1-\beta}+1}$$

$$\lambda_h(i)=\frac{\frac{p_i\beta}{1-\beta}}{\frac{p_i\beta}{1-\beta}+\frac{(1-p_i)\beta}{1-\beta}+1}$$

$$\lambda_\ell(i)=\frac{\frac{(1-p_i)\beta}{1-\beta}}{\frac{p_i\beta}{1-\beta}+\frac{(1-p_i)\beta}{1-\beta}+1}$$

2. $s_1=0$, $\lambda_\ell(i)=0$

$$\lambda_a(i)=\frac{1}{1+\frac{p_i\beta}{1-\beta}}$$

$$\lambda_h(i)=\frac{\frac{p_i\beta}{1-\beta}}{1+\frac{p_i\beta}{1-\beta}}$$

3. $s_2=0$, $\lambda_h(i)=0$

$$\lambda_a(i)=\frac{1}{1+\frac{(1-p_i)\beta}{1-\beta}}$$

$$\lambda_\ell(i) = \frac{\frac{(1-p_i)\beta}{1-\beta}}{1 + \frac{(1-p_i)\beta}{1-\beta}}$$

4. $\lambda_h(i) = \lambda_\ell(i) = 0$

$$\lambda_a(i) = 1$$

5. For all other combinations there is no feasible solution.

Therefore, we can take as our generators:

$$(1,0,0), (1, \frac{p_i\beta}{1-\beta}, \frac{(1-p_i)\beta}{1-\beta}), (1, \frac{p_i\beta}{1-\beta}, 0), (1, 0, \frac{(1-p_i)\beta}{1-\beta})$$

Any $(\lambda_a(i), \lambda_h(i), \lambda_\ell(i))$ in the cone can be expressed as a non-negative linear combination of the generators. Denote the weights of the linear combination by w_i, x_i, y_i, z_i. Hence,

$$\lambda_a(i) = w_i + x_i + y_i + z_i$$

$$\lambda_h(i) = [x_i + y_i]\frac{p_i\beta}{1-\beta}$$

$$\lambda_\ell(i) = [x_i + z_i]\frac{(1-p_i)\beta}{1-\beta}$$

These have a nice interpretation. For a box with type i,

- w_i is the mass assigned to opening them without keeping them,
- x_i is the mass assigned to opening them and unconditionally keeping,
- y_i is the mass assigned to opening them and keeping only if i has a *high* draw, and
- z_i is the mass assigned to opening them and keeping only if i has a *low* draw.

B.3 Duality

With this change of variables we can write our optimization problem for type i as

$$\max -c[w_i + x_i + y_i + z_i] + (1-\beta)[x_i + y_i]\frac{p_i\beta}{1-\beta} - (1-\beta)[x_i + z_i]\frac{(1-p_i)\beta}{1-\beta}$$

$$\text{s.t. } [w_i + x_i + y_i + z_i] + [x_i + y_i]\frac{p_i\beta}{1-\beta} + [x_i + z_i]\frac{(1-p_i)\beta}{1-\beta} + \lambda_0(i) = \frac{1}{1-\beta}$$

$$w_i + x_i + y_i + z_i \leq 1$$

$$w_i, x_i, z_i, \lambda_0(i) \geq 0$$

Let the expected return from unconditionally allocating to i be $m_i := p_i - (1 - p_i)$. Then this simplifies to

$$\max -cw_i + (m_i\beta - c)x_i + (p_i\beta - c)y_i + (-(1-p_i)\beta - c)z_i$$

$$\text{s.t. } w_i + (1 + \frac{\beta}{1-\beta})x_i + (1 + \frac{p_i\beta}{1-\beta})y_i + (1 + \frac{(1-p_i)\beta}{1-\beta}))z_i + \lambda_0(i) = \frac{1}{1-\beta}$$

$$w_i + x_i + y_i + z_i \leq 1$$

Simplifying further

$$\max -cw_i + (m_i\beta - c)x_i + (p_i\beta - c)y_i + (-(1-p_i)\beta - c)z_i$$

$$\text{s.t. } (1-\beta)w_i + x_i + (1 - (1-p_i)\beta)y_i + (1 - p_i\beta)z_i + (1-\beta)\lambda_0(i) = 1$$

$$w_i + x_i + y_i + z_i \leq 1$$

It is easy to see that we always set $w_i = z_i = 0$. The problem reduces to

$$\max(m_i\beta - c)x_i + (p_i\beta - c)y_i$$

$$\text{s.t. } x_i + (1 - (1-p_i)\beta)y_i + (1-\beta)\lambda_0(i) = 1$$

$$x_i + y_i \leq 1$$

And, as $p_i > m_i$ for all $p_i < 1$, the solution is given by $x_i = 0$ and $y_i = 1$ if $p_i\beta - c > 0$, and 0 otherwise, with $\lambda_0(i)$ set to the residual.

Now let's incorporate montonicity and incentive compatibility. Under the change of variables,

$$\Delta(p_i) = \frac{\lambda_h(p_i)}{p_i} - \frac{\lambda_\ell(p_i)}{1-p_i} = [x_i + y_i]\frac{\beta}{1-\beta} - [x_i + z_i]\frac{\beta}{1-\beta} = \frac{\beta}{1-\beta}(y_i - z_i).$$

Hence, $\Delta(p_i) \geq \Delta(p_{i-1})$ implies that

$$y_i - z_i \geq y_{i-1} - z_{i-1}.$$

Recall the upward adjacent IC constraint:

$$IC_{i-1,i} : \lambda_h(i-1) + \lambda_\ell(i-1) \geq \frac{p_{i-1}}{p_i}\lambda_h(i) + \frac{1-p_{i-1}}{1-p_i}\lambda_\ell(i)$$

This becomes

$$[x_{i-1} + y_{i-1}]\frac{p_{i-1}\beta}{1-\beta} + [x_{i-1} + z_{i-1}]\frac{(1-p_{i-1})\beta}{1-\beta} \geq p_{i-1}[x_i + y_i]\frac{\beta}{1-\beta} + (1-p_{i-1})[x_i + z_i]\frac{\beta}{1-\beta}$$

$$\Rightarrow p_{i-1}[x_{i-1} + y_{i-1}] + (1-p_{i-1})[x_{i-1} + z_{i-1}] \geq p_{i-1}[x_i + y_i] + (1-p_{i-1})[x_i + z_i]$$

$$\Rightarrow x_{i-1} + p_{i-1}y_{i-1} + (1-p_{i-1})z_{i-1} \geq x_i + p_{i-1}y_i + (1-p_{i-1})z_i$$

Notice, we can still get away with setting $w_i = 0$, as the variable is absent from both constraints. Let π_i be the probability that the agent is of type p_i. The grand (or primal) problem is

$$V = \max \sum_i \pi_i[(m_i\beta - c)x_i + (p_i\beta - c)y_i + (-(1-p_i)\beta - c)z_i]$$

$$\text{s.t. } x_i + (1-(1-p_i)\beta)y_i + (1-p_i\beta)z_i + (1-\beta)\lambda_0(i) = 1 \ \forall i \ (a_i)$$

$$x_i + y_i + z_i \leq 1 \ \forall i \ (b_i)$$

$$y_{i-1} - z_{i-1} \leq y_i - z_i \ \forall i > 1 \ (d(i-1,i))$$

$$x_i + p_{i-1}y_i + (1-p_{i-1})z_i \leq x_{i-1} + p_{i-1}y_{i-1} + (1-p_{i-1})z_{i-1} \ \forall i > 1 \ (e(i-1,i))$$

$$x_i, y_i, z_i, \lambda_0(i) \geq 0 \ \forall i$$

The dual variables appear in red brackets next to each constraint. The dual is

$$W = \min \sum_i a_i + b_i$$

$$\text{s.t. } a_i + b_i + e(i-1,i) - e(i,i+1) \geq \pi_i(m_i\beta - c) \ \forall i \ (x_i)$$

$$(1-(1-p_i)\beta)a_i + b_i - d(i-1,i) + d(i,i+1) + p_{i-1}e(i-1,i) - p_ie(i,i+1) \geq \pi_i(p_i\beta - c) \ \forall i \ (y_i)$$

$$(1-p_i\beta)a_i + b_i + d(i-1,i) - d(i,i+1) + (1-p_{i-1})e(i-1,i) - (1-p_i)e(i,i+1) \geq \pi_i(-(1-p_i)\beta - c) \ \forall i \ (z_i)$$

$$(1-\beta)a_i \geq 0 \ \forall i \ (\lambda_0(i))$$

$$a_i, b_i, d(i-1,i), e(i-1,i) \geq 0 \ \forall i \ \text{ with } \ d(0,1) = e(0,1) = d(n,n+1) = e(n,n+1) = 0$$

The primal variables appear in blue brackets next to each constraint.

Note that, for all i, the right-hand sides of the constraints are ordered:

$$p_i\beta - c \geq m_i\beta - c \geq -(1-p_i)\beta - c$$

where the first inequality is strict if $p_i < 1$ and the second if $p_i > 0$.

We will now use duality to show there are only three types of solution to our problem:

- *Full ideal inspection*: Open any box i and keep only if i has a high draw,

$$x_i = z_i = 0 \text{ and } y_i = 1 \text{ for all } i$$

- *Partial separation*: Given a threshold t, open any box i that is above t and keep only if i has a high draw, and open box any i below t with probability p_t and unconditionally allocate,

$$x_i = z_i = 0 \text{ and } y_i = 1 \text{ for all } i > t$$

$$x_i = p_t \text{ and } z_i = y_i = 0 \text{ for all } i \leq t$$

- *No allocation*: Never open any box i,

$$x_i = z_i = y_i = 0 \text{ for all } i$$

As a demonstration, we have the following immediate result.

Proposition 1 *Full ideal inspection is optimal if all types are sufficiently high, and no allocation is optimal if all types are sufficiently low. That is,*

- *if $p_i\beta - c \geq 0$ for all i, then $x_i = z_i = 0$ and $y_i = 1$ for all i is optimal, and*
- *if $p_i\beta - c \leq 0$ for all i, then $x_i = z_i = y_i = 0$ for all i is optimal.*

Proof: By duality, if the value of a feasible solution to the primal is equal to the value of a feasible solution to the dual, these solutions are optimal. As such, if we can find a feasible solution to the dual with the same value as the conjectured solution, we've proved optimality.

Begin with $x_i = z_i = 0$ and $y_i = 1$ for all i. Clearly this is feasible - making sure to set $\lambda_0(i) = \frac{(1-p_i)\beta}{1-\beta}$ for all i) - with monotonicity and incentive compatibility binding for each i. The primal value for this solution is:

$$V = \sum_i \pi_i(p_i\beta - c)$$

For the dual, complementary slackness gives us $a_i = 0$ and that the second constraint should bind. Set $b_i = \pi_i(p_i\beta - c)$ and $e(i, i+1) = 0$ and $d(i, i+1) = 0$ for all i. Then, clearly the first and third constraints are satisfied by the ordering of the right-hand sides, and $b_i \geq 0$ if $p_i\beta - c \geq 0$. The dual value for this solution is:

$$W = \sum_i \pi_i(p_i\beta - c)$$

Now consider $x_i = z_i = y_i = 0$ for all i. Clearly this is feasible - making sure to set $\lambda_0(i) = \frac{1}{1-\beta}$ for all i - with monotonicity and incentive compatibility binding for each i. The primal value for this solution is:

$$V = 0$$

For the dual, complementary slackness gives us $a_i = 0$. Set $b_i = 0$, $e(i, i+1) = 0$ and $d(i, i+1) = 0$ for all i. Then, clearly the first, second, and third constraints are satisfied if $0 \geq \pi_i(p_i\beta - c)$ and by the ordering of the right-hand sides. The dual value for this solution is:

$$W = 0$$

Then these two are indeed optimal solutions to the primal given their associated condition. ■

This result should be unsurprising: if the first best is incentive compatible, then the first best is achievable. The interesting case is when types fall on either side of this threshold and the first best isn't achievable, which is described by the following result.

Proposition 2 *Partial separation is optimal if the first best treats different types differently and if the value of partial separation is positive. That is, if $p_1\beta - c < 0$ and $p_n\beta - c > 0$, and $V \geq 0$, then the optimal solution is given by*

- *$x_i = z_i = 0$ and $y_i = 1$ if $i > t$, and*
- *$x_i = p_t$ and $z_i = y_i = 0$ if $i \leq t$,*

where t is set to maximise

$$V_t = \sum_{i \leq t} \pi_i p_t (m_i\beta - c) + \sum_{i > t} \pi_i (p_i\beta - c)$$

If $V \leq 0$, no allocation is optimal. That is, $x_i = z_i = y_i = 0$ for all i.

Proof: To prove this, we will first argue that partial separation is feasible and highlight a few facts about the threshold, t. Then we will find a feasible dual solution whose value coincides with partial separation proving the initial result. This dual solution will have the feature that, if the value of partial separation is negative, then it coincides with no allocation instead and no allocation has already been shown to be feasible, proving the final result.

Clearly partial separation is feasible - making sure to set $\lambda_0(i) = \frac{1-p_i\beta}{1-\beta}$ if $i \leq t$ and $\lambda_0(i) = \frac{1-p_t}{1-\beta}$ if $i > t$ - with incentive compatibility binding for each i and monotonicity binding for all $i \neq t$.

The primal value for this solution is:

$$V = \sum_{i \leq t} \pi_i p_t (m_i \beta - c) + \sum_{i > t} \pi_i (p_i \beta - c)$$

Observe that $p_t < \frac{c}{\beta}$. That is, under the first best, type t is not worth inspecting. To show this, note that:

- if $p_i \leq \frac{c}{\beta}$ then $0 \geq p_i\beta - c > m_i\beta - c$,
- if $\frac{c}{\beta} < p_i \leq \frac{1}{2} + \frac{c}{2\beta}$ then $p_i\beta - c > 0 \geq m_i\beta - c$, and
- if $\frac{1}{2} + \frac{c}{2\beta} < p_i$ then $p_i\beta - c \geq m_i\beta - c > 0$.

The proposition assumes $c < p_n\beta \leq 1$ so $\frac{c}{\beta} < \frac{1}{2} + \frac{c}{2\beta}$. Now suppose, $p_t \geq \frac{c}{\beta}$, and consider lowering t to s, the highest type such that $p_s < \frac{c}{\beta}$. The net change in V is given by

$$\sum_{i \leq s} \pi_i p_s (m_i\beta - c) + \sum_{i > s} \pi_i (p_i\beta - c) - \sum_{i \leq t} \pi_i p_t (m_i\beta - c) - \sum_{i > t} \pi_i (p_i\beta - c)$$

$$\Rightarrow -\sum_{i \leq s} \pi_i (p_t - p_s)(m_i\beta - c) + \sum_{i > s}^{t} \pi_i [(p_i\beta - c) - p_t(m_i\beta - c)] > 0$$

The inequality comes from the first term being negative as $m_i\beta - c < 0$ for all $i \leq s$, and the second term being positive as, for any fraction α, $p_i\beta - c > 0 \geq \alpha(m_i\beta - c)$ if $p_i \in (\frac{c}{\beta}, \frac{1}{2} + \frac{c}{2\beta}]$ and $p_i\beta - c > \alpha(m_i\beta - c) > 0$ if $p_i > \frac{1}{2} + \frac{c}{2\beta}$. This is a contradiction to the optimality of t.

As t is chosen to maximise V, it must be that,

$$\sum_{i \leq t} \pi_i p_t (m_i\beta - c) + \sum_{i > t} \pi_i (p_i\beta - c) \geq \sum_{i \leq t+1} \pi_i p_{t+1} (m_i\beta - c) + \sum_{i > t+1} \pi_i (p_i\beta - c)$$

$$\Rightarrow \pi_{t+1}(p_{t+1}\beta - c) \geq \pi_{t+1} p_{t+1} (m_{t+1}\beta - c) + \sum_{i \leq t} \pi_i (p_{t+1} - p_t)(m_i\beta - c)$$

$$\Rightarrow \pi_{t+1}[(p_{t+1}\beta - c) - p_{t+1}(m_{t+1}\beta - c)] \geq \sum_{i \leq t} \pi_i (p_{t+1} - p_t)(m_i\beta - c)$$

and similarly,

$$\sum_{i \leq t} \pi_i p_t (m_i\beta - c) + \sum_{i > t} \pi_i (p_i\beta - c) \geq \sum_{i \leq t-1} \pi_i p_{t-1} (m_i\beta - c) + \sum_{i > t-1} \pi_i (p_i\beta - c)$$

$$\Rightarrow \pi_t p_t (m_t\beta - c) + \sum_{i \leq t-1} \pi_i (p_t - p_{t-1})(m_i\beta - c) \geq \pi_t (p_t\beta - c)$$

$$\Rightarrow \sum_{i \leq t-1} \pi_i (p_t - p_{t-1})(m_i\beta - c) \geq \pi_t [(p_t\beta - c) - p_t(m_t\beta - c)]$$

Let,

$$\varphi(i) := \pi_i[(p_i\beta - c) - p_i(m_i\beta - c)] - \sum_{j \le i-1} \pi_j(p_i - p_{i-1})(m_j\beta - c)$$

Then these two conditions can be written succinctly as: $\varphi(t+1) \ge 0 \ge \varphi(t)$.

Finally, note that $p_t\beta - c > p_t(m_t\beta - c)$. This follows from $0 \ge \varphi(t)$; if $p_t\beta - c \le p_t(m_t\beta - c)$ then lowering t weakly increases the value of treating type t and lowers the probability of assigning to all types below t, contradicting the optimality of t.

To find the corresponding dual solution, note that complementary slackness gives us $a_i = 0$ for all i, and that the first constraint should bind for $i \le t$ and the second for $i > t$. We are looking for weakly positive values of b_i, $e(i, i+1)$ and $d(i, i+1)$ that are feasible, satisfy complementary slackness and for which the value function coincides. Ignoring the third constraint and confirming it's satisfied at the end, the two constraints of importance are:

$$b_i + e(i-1, i) - e(i, i+1) \ge \pi_i(m_i\beta - c)$$

$$b_i - d(i-1, i) + d(i, i+1) + p_{i-1}e(i-1, i) - p_i e(i, i+1) \ge \pi_i(p_i\beta - c)$$

For $i \le t$, set $b_i = 0$. As the first constraint must bind,

$$e(i, i+1) - e(i-1, i) = -\pi_i(m_i\beta - c)$$

This is a difference equation and an initial value, $e(0, 1)$, set at 0, so:

$$\begin{aligned} e(i, i+1) &= -\pi_i(m_i\beta - c) + e(i-1, i) \\ &= -\pi_i(m_i\beta - c) + -\pi_{i-1}(m_{i-1}\beta - c) + e(i-2, i-1) \\ &= \dots \\ &= -\sum_{j \le i} \pi_j(m_j\beta - c) \end{aligned}$$

Recall that $m_i\beta - c < 0$ for all $p_i \le t$, and so $e(i, i+1) > 0$.

Rearranging the second constraint,

$$d(i, i+1) - d(i-1, i) \ge \pi_i(p_i\beta - c) + p_i e(i, i+1) - p_{i-1}e(i-1, i)$$

and substituting for $e(i,i+1)$,

$$\begin{aligned} d(i,i+1)-d(i-1,i) &\geq \pi_i(p_i\beta-c)-p_i\sum_{j\leq i}\pi_j(m_j\beta-c)+p_{i-1}\sum_{j\leq i-1}\pi_j(m_j\beta-c) \\ &= \pi_i(p_i\beta-c)-p_i\pi_i(m_i\beta-c)-\sum_{j\leq i-1}\pi_j(p_i-p_{i-1})(m_j\beta-c) \\ &= \pi_i[(p_i\beta-c)-p_i(m_i\beta-c)]-\sum_{j\leq i-1}\pi_j(p_i-p_{i-1})(m_j\beta-c)=\varphi(i) \end{aligned}$$

As we know $\varphi(i)\leq 0$ by the optimality of t, then the right hand side is negative, and as such setting $d(i,i+1)-d(i-1,i)=0$ satisfies this inequality.

Now consider the largest type, s, such that the value of partial separating would be negative with types only below s is negative. That is, the largest s such that,

$$\sum_{i\leq t}\pi_i p_t(m_i\beta-c)+\sum_{i=t+1}^{s}\pi_i(p_i\beta-c)\leq 0$$

Note that $p_{s+1}>\frac{c}{\beta}$ as for the left hand side to be greater than 0 when we increase the index there must be at least one type such that $p_i\beta-c>0$.

For $i>t$ but $i\leq s$, once again set $b_i=0$ and $e(i,i+1)-e(i-1,i)=-\pi_i(m_i\beta-c)$. Even though this may no longer be positive, we know that as $i\leq s$,

$$\sum_{j\leq t}\pi_j p_t(m_j\beta-c)+\sum_{j=t+1}^{i}\pi_j(p_j\beta-c)\leq 0$$

$$-\sum_{j\leq i}\pi_j(m_j\beta-c)+\sum_{j\leq t}\pi_j p_t(m_j\beta-c)+\sum_{j=t+1}^{i}\pi_j(p_j\beta-c)\leq-\sum_{j\leq i}\pi_j(m_j\beta-c)$$

$$-\sum_{j\leq t}\pi_j(1-p_t)(m_j\beta-c)+\sum_{j=t+1}^{i}\pi_j[(p_j\beta-c)-(m_j\beta-c)]\leq e(i,i+1)$$

And this left hand side is positive as $m_j\beta-c<0$ when j and $p_j\beta-c>m_j\beta-c$ for all j, so $e(i,i+1)$ must be positive too.

This means, the second, and now binding, constraint simplifies to:

$$d(i,i+1)-d(i-1,i)=\varphi(i)$$

And solving the difference equation:

$$\begin{aligned}
d(i,i+1) &= \varphi(i) + d(i-1,i) \\
&= \varphi(i) + \varphi(i-1) + d(i-2,i-1) \\
&= \varphi(i) + \varphi(i-1) + \ldots + d(t,t+1) \\
&= \sum_{j=t+1}^{i} \varphi(j)
\end{aligned}$$

We need to show this too is positive. Consider the following:

$$\varphi_i = \pi_i[(p_i\beta - c) - p_i(m_i\beta - c)] - \sum_{j \le i-1} \pi_j(p_i - p_{i-1})(m_j\beta - c)$$

$$\varphi_{i-1} = \pi_{i-1}[(p_{i-1}\beta - c) - p_{i-1}(m_{i-1}\beta - c)] - \sum_{j \le i-2} \pi_j(p_{i-1} - p_{i-2})(m_j\beta - c)$$

$$\Rightarrow \varphi_i + \varphi_{i-1} = \pi_i[(p_i\beta - c) - p_i(m_i\beta - c)] + \pi_{i-1}[(p_{i-1}\beta - c) - p_i(m_{i-1}\beta - c)] - \sum_{j \le i-2} \pi_j(p_i - p_{i-2})(m_j\beta - c)$$

Repeating this sum gives us:

$$\sum_{j=t+1}^{i} \varphi(j) = \sum_{j=t+1}^{i} \pi_j[(p_j\beta - c) - p_i(m_j\beta - c)] - \sum_{j \le t} \pi_j(p_i - p_t)(m_j\beta - c) \ge 0$$

where the inequality comes from: $m_j\beta - c < 0$ when $j \le t$ and $p_j\beta - c > m_j\beta - c$ for all j.

For $i = s+1$, set $b_{s+1} = \pi_{s+1}(p_{s+1}\beta - c) + \sum_{i \le t} \pi_i p_t(m_i\beta - c) + \sum_{i=t+1}^{s} \pi_i(p_i\beta - c)$, and for $i > s+1$, set $b_i = \pi_i(p_i\beta - c)$. Notice that even with $e(i,i+1) - e(i-1,i) = 0$ the first constraint holds,

$$b_i + e(i-1,i) - e(i,i+1) = \pi_i(p_i\beta - c) \ge \pi_i(m_i\beta - c)$$

and with $d(i,i+1) - d(i-1,i) = 0$ the second constraint holds with equality,

$$b_i - d(i-1,i) + d(i,i+1) + p_{i-1}e(i-1,i) - p_i e(i,i+1) = \pi_i(p_i\beta - c)$$

Further, note that $b_i > 0$ as $i \ge s+1$ and thus $p_i > \frac{c}{\beta}$.

As such, we have a feasible solution to the dual. The value of this solution is

$$W = \sum_{i \le t} \pi_i p_t(m_i\beta - c) + \sum_{i > t} \pi_i(p_i\beta - c)$$

if $s < n$, and 0 otherwise. In the later case, we already know of a feasible solution that has $V = 0$:

no allocation where $x_i = y_i = z_i = 0$ for all i.

As such, the values correspond, and these are indeed optimal solutions. ∎