Engagement Design

Daniel Fershtman[*] Alessandro Pavan[†]

April 2026

(Preliminary and Incomplete)

Abstract

We consider the problem of a designer maximizing a searcher's engagement, defined as the expected number of alternatives explored, with each exploration discounted by the time at which it occurs. We frame this novel design problem within Weitzman (1979)'s search model, where alternatives are Pandora's boxes whose opening is costly and reveals a prize drawn from a known distribution. We allow the designer to select not only the distributions from which the prizes are drawn but also the order the boxes are provided. Fixing the number of alternatives explored, the designer prefers the searcher to walk away with a lower prize. We show that there is no value in differentiating the boxes and making the most attractive ones available only in later periods: the designer cannot do better than homogenizing the boxes, with each prize drawn from a binary distribution. When there are no costs in supplying the boxes, it is optimal to make all of them available at the outset. When, instead, it takes time to supply the boxes, sequential provision dominates, even if all boxes are identical.

1 Introduction

Digital environments increasingly rely on sequential engagement: users browse products, swipe through profiles, or explore content streams, often deciding at each step whether to continue or stop. A central design problem in such settings is how to structure the flow of opportunities so as to sustain engagement over time. This paper studies a novel formulation of this problem: how a designer should construct and sequence "opportunities" in order to

---

[*] Tel Aviv University and Emory University. Email: danfershtman@gmail.com

[†] Department of Economics, Northwestern University. Email: apavan31@gmail.com

maximize the extent of exploration by a forward-looking agent who faces time and search frictions.

We formalize this question within a generalized version of the classic Pandora's box model of Weitzman (1979), in which a searcher sequentially inspects alternatives whose values are initially unknown and must be discovered at a cost. In contrast to the canonical formulation, we introduce an active designer who controls both the distribution of values associated with each alternative and the order in which these alternatives are presented. The designer's objective is not to maximize the searcher's payoff, but rather to maximize engagement, defined as the expected (discounted) number of alternatives the searcher explores, net of the (discounted) cost associated with the selected alternative.

This perspective captures a broad range of applications. In online platforms, firms curate menus of products, matches, or pieces of content to prolong user interaction. In labor or education markets, institutions design evaluation or screening processes that shape how many options are considered before a decision is made. More generally, any environment in which a principal controls both the supply and sequence of alternatives made available to an optimizing agent can be viewed through these lenses. Across these contexts, a fundamental tension arises: while the designer benefits from continued exploration, the searcher ultimately selects a single option and therefore values early access to high-quality alternatives.

The key challenge in this environment is that the designer influences behavior only indirectly, through the statistical properties and timing of the alternatives. Because the searcher behaves optimally given his beliefs and incentives, the designer's problem is inherently one of incentive design. In particular, the designer must shape the trade-off faced by the searcher between exploiting currently available options (e.g., by irreversibly selecting one) and continuing to explore in the hope of finding better alternatives in the future.

A natural conjecture is that sustaining engagement requires dynamic differentiation: the designer might wish to "save" the most attractive opportunities for later periods, thereby giving the searcher a reason to continue exploring. Alternatively, one might expect that offering a rich and heterogeneous set of options upfront could stimulate exploration by increasing the dispersion of potential payoffs. More broadly, it is not a priori clear whether optimal design calls for concentration or dispersion of value, nor whether it should rely on dynamic sequencing or static provision of the alternatives to explore.

This paper provides a characterization of optimal design in this environment and delivers a set of results that run counter to these intuitions. Our first main contribution is to show that, without loss of optimality, the designer can restrict attention to binary distributions for each alternative: each box yields either the maximal prize or the lowest possible value. That is, any richer distribution can be replaced by a two-point distribution that preserves the searcher's incentives while weakly increasing the designer's payoff.

The intuition for this result is rooted in the structure of optimal search. The searcher's behavior is governed by an index rule: each alternative is evaluated through a sufficient statistic – the *inspection index* – that summarizes its ex-ante attractiveness. Because the alternatives are made available sequentially to the searcher, a second statistic – the *expansion index* – summarizes the value of pausing the exploration of the alternatives already available to request new ones. While each alternative's inspection index is a function only of the distribution from which its prize is drawn, the expansion index is a function of the distributions of all the alternatives that the designer has yet to make available to the searcher.

We show that, fixing the collection of the inspection indexes, the designer benefits from minimizing the probability that opening a box leads to immediate stopping. Binary distributions achieve precisely this by concentrating probability mass at the extremes, thereby preserving the ex-ante appeal of each alternative while reducing the likelihood that its exploration yields a sufficiently high realization to terminate the search. The key step in the proof shows that index-preserving binarizations of the original sequence of prize distributions reduce the expansion indexes. Such a reduction is beneficial to the designer because it (a) induces the searcher to swap expansion with exploration, with the effect that alternatives with a smaller inspection index are explored earlier, (b) reduces the delay in the explorations, and (c) lowers the expected cost to the designer of providing a high prize.

Our second main result characterizes the optimal binary design. We show that all alternatives should be identical and should deliver the highest prize with the lowest probability consistent with being explored. This probability is pinned down by the condition that the searcher's inspection index is exactly zero. Intuitively, each alternative is made just attractive enough to be worth exploring, in comparison with the outside option, instead of the value of exploring other alternatives. Any increase in the probability of a high prize (1) raises the likelihood that the searcher stops after exploring the alternative, and (2) raises the expansion index, making it more difficult to persuade the searcher to explore the alternative immediately after receiving it.

This characterization highlights a central trade-off in engagement design. On the one hand, alternatives must be sufficiently appealing to induce immediate exploration; on the other hand, they should not be so appealing as to prematurely terminate the search. The optimal design balances this trade-off by making each alternative marginally worth exploring, thereby sustaining exploration for as long as possible, while also reducing the expected cost of providing a high prize.

Our third set of results concerns the timing of provision. We compare environments in which alternatives are supplied sequentially with those in which they are all made available to the searcher at the outset. When supplying alternatives is costless in time, we show that the designer cannot improve upon simultaneous provision: because there is no value in differentiating the alternatives, there is also no value in providing them sequentially. In contrast, when provision itself is time-consuming, sequential supply strictly dominates simultaneous supply, even if the optimal sequence does not involve any differentiation across the alternatives. This is because sequential provision allows the designer to avoid the time cost of supplying additional alternatives in case the exploration of the earlier ones yields a favorable outcome that induces the searcher to stop.

Related literature. The paper contributes to several strands of the literature. First, it extends the theory of optimal search (e.g., Gittins and Jones (1974), Gittins (1979), and Weitzman (1979)) by endogenizing the distribution from which each alternative's payoff is drawn and by relaxing the assumption that all alternatives are available to the searcher at the outset of the exploration. The optimality of the searcher's strategy follows from Fershtman and Pavan (2025); that paper identifies conditions under which the optimal policy takes the form of an index rule, with a novel index for the expansion of the searcher's consideration set (the pool of available options) characterized in recursive form. The present paper builds on that result to endogenize the distributions from which the alternatives' payoffs are drawn and the order in which the alternatives are made available to the searcher. Related is also Auster and Che (2025), who study a variant of the original Pandora's boxes problem in which the searcher is uncertain about the distribution from which each box's prize is drawn.[1]

Second, the paper relates to the growing literature on information design and persuasion, where a principal shapes the information environment to influence an agent's behavior (see Bergemann and Morris (2019) and Kamenica (2019) for overviews). Here, instead of designing signals about a fixed state, the designer controls the stochastic properties of the opportunities themselves as well as the order they are made available to the searcher.

Third, the paper connects to recent work on platform design and attention management, where the objective is to maximize user engagement rather than welfare (see, for example, Gossner, Steiner and Stewart (2021) and Hebert and Zhong (2025)). A common theme in this literature is that the designer influences the agent's stopping by shaping the likelihood that each exploration leads to termination. Contrary to these papers, our analysis is conducted within a classical multi-armed bandit model and involves the design of the alternatives' payoff distributions instead of their information.

To the best of our knowledge, ours is the first paper comparing the benefit of providing (optimally designed) alternatives simultaneously versus sequentially, which is one of the key

---

[1] In that paper, an adversarial Nature plays a role analogous to our designer. The designer's objective is different in the two papers and so is the searcher's optimal strategy (that paper considers min-max regret as a criterion of optimality whereas, in our analysis, because the distributions are endogenous but unambiguous, the searcher's strategy is an index policy with a special index for box expansion).

contributions of this work.

Organization. The remainder of the paper is organized as follows. Section 2 introduces the model and formalizes the interaction between the designer and the searcher. Section 3 characterizes the searcher's optimal behavior and shows it takes the form of an index policy with a special index for the expansion of the set of boxes. Section 4 derives the optimal box design and establishes the optimality of binary distributions. Section 5 compares simultaneous and sequential box provision. The Appendix collects proofs not in the main text.

2 Model

A designer (she) interacts with a searcher (he) who sequentially inspects a set of "Pandora's boxes." The designer controls the order in which $m \in \mathbb{N}$ boxes are provided to the searcher as well as the distributions from which each box's prize is drawn. In particular, she commits to a sequence of distributions $\mathbf{F} = (F_1, ..., F_m)$ such that each box's prize v_j is drawn from the distribution F_j with support contained in $[0, \bar{v}]$, where $\bar{v} > 0$, with the draws independent across boxes. Hereafter, we refer to a generic sequence $\mathbf{F} = (F_1, ..., F_m)$ of distributions as a *box design.*[2]

Given $\mathbf{F}$, at $t = 0$, the searcher can either request the first box or opt out and enjoy the outside option whose payoff is zero. At $t = 1$, conditional on having requested the first box at $t = 0$, the searcher can either open the received box, request a new one, or opt out by selecting the outside option. At the beginning of each subsequent period $t \geq 2$, the searcher decides among three actions: (i) selecting a box he opened already, bringing an end to the search, with payoff $\delta^t v_j$, where $\delta \in (0, 1)$ denotes the searcher's discount factor and v_j the prize of the selected box, (ii) open one of the unopened boxes among those received

---

[2] Depending on the application, the bound $\bar{v} > 0$ on the prize may reflect the maximal attractiveness of a possible partner, the maximal discount for the purchase of a good of given quality, or, in case the box design problem is interpreted literally, the designer's limited budget (in this later case, note that, because the searcher cannot walk away with two boxes, the designer can afford to put $\bar{v}$ in more than one box, while respecting her budget constraint).

in previous periods, incurring a cost of $\delta^t c$, with $c \in (0, \bar{v})$, and learning the box's content, v_j, or (iii) request a new box, incurring the time cost of waiting a period. Requesting a box makes the box available at the start of the next period. This time friction captures the cost to the searcher to familiarize with the box; it may also reflect the cost to the designer of providing the box.

The designer receives a gross payoff equal to 1 for each box the searcher opens. Such payoffs are discounted by the times at which the openings occur. In particular, a box opened after t periods brings a payoff equal to ρ^t to the designer, where $\rho \in (0, 1)$. If the searcher ultimately chooses a box with realized value v_j after t periods, the designer incurs a cost $\rho^t a v_j$, with $a \in [0, \delta/c\rho)$. As it will become clear in due time the assumption is both necessary and sufficient to guarantee that, under the optimal sequence, the designer receives more than her outside option, which we assume is also equal to zero.

3 Searcher's optimal strategy

We start by describing the searcher's optimal policy. The result in this section (summarized in Lemma 1) follows from Fershtman and Pavan (2025).

Given any box design $\mathbf{F} = (F_1, ..., F_m)$, the sequence of inspection and expansion decisions that maximizes the searcher's expected discounted payoff, net of all costs, is an index rule. For each box j, let I_j be the box's *inspection index*. Furthermore, for each expansion j, let E_j denote the *expansion index* associated with the request of box j. Each I_j depends on $\mathbf{F}$ only through F_j, whereas each E_j depends on $\mathbf{F}$ only through $(F_j, ..., F_m)$. The following lemma characterizes the searcher's optimal behavior.

Lemma 1. *For any $\mathbf{F} = (F_1, ..., F_m)$, the searcher's optimal policy is the following. Take any history at which the searcher has received boxes $l \in \{1, ..., m-1\}$. The searcher:*

1. *asks for box $l+1$ if the expansion index E_{l+1} is greater than the inspection index I_j of every box $j \leq l$ not opened yet, and greater than the flow value $(1-\delta)v_j$ of every*

opened box $j \leq l$;

2. *opens box $j \leq l$ if its inspection index I_j is greater than the maximum between (a) the inspection index I_s of any other unopened box $s \leq l$, (b) the expansion index E_{l+1}, and (c) the flow value $(1-\delta)v_i$ of every opened box $i \leq l$;*

3. *chooses the opened box $j \leq l$ if its flow value $(1-\delta)v_j$ is greater than the maximum between (a) the expansion index E_{l+1}, (b) the inspection index I_i of every unopened box $i \leq l$, and (c) the flow value $(1-\delta)v_s$ of any other opened box $s \leq l$.*

The inspection indices solve the following equations:

$$I_j = \frac{-c + \delta \int_{\frac{I_j}{1-\delta}}^{\overline{v}} v \, dF_j(v)}{1 + \frac{\delta}{1-\delta}\left(1 - F_j\left(\frac{I_j}{1-\delta}\right)\right)}. \tag{1}$$

The expansion indices are defined recursively as follows:

$$E_j = \frac{\mathbb{E}^{\chi^*}\left[\sum_{s=0}^{\tau^*-1} \delta^s U_s\right]}{\mathbb{E}^{\chi^*}\left[\sum_{s=0}^{\tau^*-1} \delta^s\right]}, \tag{2}$$

where (a) χ^* is the index policy described above, (b) U_s is the s-th flow payoff (cost or reward) under the policy χ^* when the latter is conducted in a fictitious environment starting with an empty set of boxes and where the available boxes are only those encountered after the j-th expansion, and (c) τ^* is the first time at which, under such a policy, the expansion index E_l, for $l > j$, as well as all the inspection indices of the boxes I_l, for $l \geq j$, received after the j-th expansion and the flow values $(1-\delta)v_l$ of the boxes $l \geq j$ opened under the policy χ^* fall weakly below E_j.

The lemma follows from Theorem 1 in Fershtman and Pavan (2025). Note that the optimality of an index policy cannot be established by interpreting the request of a box as an additional cost and then applying the same arguments as in standard bandit problems where all boxes are available from the outset. The reason is that the searcher cannot request box j if he has not requested all boxes $l < j$. As a result of this friction, the expansion index E_j typically depends not only on the box's distribution F_j but also on the distributions F_l of boxes $l > j$.

For each box j, the inspection index I_j is the familiar Gittins index (Gittins (1979)), which, in the problem under consideration, coincides with Weitzman (1979) reservation price.[3] The expansion index E_j, instead, is equal to the maximal expected payoff, per expected discounted unit of time, with the optimization restricted to the boxes encountered after the j-th expansion. That is,

$$E_j = \sup_{\tau,\chi} \frac{\mathbb{E}^{\chi}\Big[\sum_{s=0}^{\tau-1} \delta^s \, U_s\Big]}{\mathbb{E}^{\chi}\Big[\sum_{s=0}^{\tau-1} \delta^s\Big]},$$

where χ is a generic policy specifying, for each period, whether (a) to open one of the boxes received after the j-th expansion, (b) request an additional box, (c) select one of the boxes received after the j-th expansion that have been opened already, or (d) stop the search by selecting the outside option. The terms U_s in the numerator are the generic flow payoff under such a policy. The term τ is a stopping time, specifying when to interrupt the entire search process in this fictitious problem where all boxes $r = 1, ..., j-1$ are not available and the search starts with the request of the j-th box. Note that, by stopping immediately after requesting the j-th box (without opening it), the searcher can guarantee himself a payoff of zero. It follows that, for all boxes and all designs $\mathbf{F}$, $E_j \geq 0$.

4 Optimal box design

We now characterize the sequence of prize distributions that maximizes the designer's ex-ante expected payoff. For any design $\mathbf{F}$, we assume that, at any history, in case of indifference, the searcher follows the designer's recommended action.

---

[3] Weitzman (1979) defines the "reservation prize" as the solution to $\hat{I} = \delta \int_{\hat{I}}^{\infty} v dF(v) + \delta F(\hat{I})\hat{I} - c$. The reservation prize is thus equal to the Gittins index (equivalently, our inspection index) divided by $1 - \delta$.

Below we first establish a few lemmas that jointly will permit us to arrive at our first main result, the optimality of binary distributions.

Let $h : [0, 1] \to \mathbb{R}$ be the function defined, for all $p \in [0, 1]$, by

$$h(p) = \frac{-c + \delta p \overline{v}}{1 + \frac{\delta}{1-\delta} p}. \tag{3}$$

Lemma 2. *Take any distribution F with support contained in $[0, \overline{v}]$ such that supp$[F] \neq \{0, \overline{v}\}$. Let I be the inspection index associated with such a distribution, and assume $I \geq 0$. There exists a binary distribution F^B with support $\{0, \overline{v}\}$ such that the inspection index I^B under F^B is the same as under the original distribution F. The distribution F^B is characterized by the unique value of $p^B = \Pr(\tilde{v} = \overline{v})$ that solves $h(p^B) = I$. The result follows from these properties along with the observation that, for any binary distribution F^B with support $\{0, \overline{v}\}$ and $\Pr(\tilde{v} = \overline{v}) = p$, the inspection index is $h(p)$.*

Proof of Lemma 2. Note that h is continuous and strictly increasing, with $h(0) = -c$, and $h(1) = (-c + \delta \overline{v})(1 - \delta) \geq I$, with the inequality strict unless F assigns probability one to $\overline{v}$. Hence, there exists a unique $p^B \in (0, 1]$ such that $h(p^B) = I$. $\square$

Hereafter, we focus on designs $\mathbf{F}$ such that $I_j \geq 0$ all j; because any box with a negative inspection index is never opened, it is without loss of optimality for the designer to restrict attention to designs such that $I_j \geq 0$ all j.[4]

Definition 1. For any box design $\mathbf{F} = (F_1, ..., F_m)$ such that $I_j \geq 0$ all j, its *binarization* is the design $\mathbf{F}^B = (F_1^B, ..., F_m^B)$ such that, for each $j = 1, ..., m$,

1. F_j^B has support $\{0, \overline{v}\}$,
2. $I_j^B = I_j$, with I_j^B denoting box j's inspection index under the distribution F_j^B.

The following lemma compares the probability the prize exceeds the inspection index under the original distribution and its binarization.

---

[4] If the designer does not want the searcher to open one of the boxes, she can always make its inspection index equal to zero and recommend that the searcher does not open it.

Lemma 3. *Take any distribution F with inspection index $I \geq 0$ and let p^B be the unique solution to $h(p^B) = I$. Then $p^B \leq 1 - F(I/(1-\delta))$.*

Proof of Lemma 3. Note that

$$I = \frac{-c(1-\delta) + \delta\left[1 - F\left(\frac{I}{1-\delta}\right)\right](1-\delta)\int_{\frac{I}{1-\delta}}^{\overline{v}} v \frac{\mathrm{d}F(v)}{1-F\left(\frac{I}{1-\delta}\right)}}{1 - \delta + \delta\left[1 - F\left(\frac{I}{1-\delta}\right)\right]},$$

is a convex combination between $-c$ and

$$(1-\delta)\int_{\frac{I}{1-\delta}}^{\overline{v}} v \frac{\mathrm{d}F(v)}{1-F\left(\frac{I}{1-\delta}\right)} = (1-\delta)\mathbb{E}^F\left[\tilde{v} | \tilde{v} \geq I/(1-\delta)\right]$$

with weights given by

$$\frac{1-\delta}{1 - \delta + \delta\left[1 - F\left(\frac{I}{1-\delta}\right)\right]}$$

and

$$\frac{\delta\left[1 - F\left(\frac{I}{1-\delta}\right)\right]}{1 - \delta + \delta\left[1 - F\left(\frac{I}{1-\delta}\right)\right]}$$

respectively. Likewise, $h(p^B)$ is a convex combination of $-c$ and $(1-\delta)\overline{v}$ with weights equal to

$$\frac{1-\delta}{1-\delta+\delta p^B}$$

and

$$\frac{\delta p^B}{1-\delta+\delta p^B}$$

respectively. Because

$$\mathbb{E}^F\left[\tilde{v} | \tilde{v} \geq I/(1-\delta)\right] \leq \overline{v}$$

the fact that $h(p^B) = I$ implies $p^B \leq 1 - F\left(\frac{I}{1-\delta}\right)$. □

Intuitively, Lemma 3 must hold because, if it did not, that would mean that not only does F^B assign a higher probability to prizes above $I = I^B$ than F, but the probability is concentrated at the highest possible value, $\overline{v}$. Since realizations below the inspection index

$I = I^B$ are irrelevant for the calculation of the index, the index associated with F^B would be higher than that associated with F.

Next, we show that a similar property holds for any pair of cutoffs below I, irrespectively of their order.

Lemma 4. *Take any distribution F with inspection index $I \geq 0$ and let F^B be the unique binary distribution with support $\{0, \overline{v}\}$ such that $p^B = \Pr(\tilde{v} = \overline{v})$ solves $h(p^B) = I$. For any pair of cut-offs $a, b \in [0, I/(1-\delta)]$, with $a \leq b$, $F^B(a) \geq F(b)$.*

Proof of Lemma 4. This follows from the fact that

$$F^B(a) = 1 - p^B = \underbrace{F(I/(1-\delta))}_{\geq F(b)} + \underbrace{1 - F(I/(1-\delta)) - p^B}_{\geq 0} \geq F(b),$$

where the inequality $F(I/(1-\delta)) \geq F(b)$ follows from the fact that $b \leq I/(1-\delta)$, whereas the inequality $1 - F(I/(1-\delta) \geq p^B$ follows from Lemma 3. $\square$

The next lemma shows that, for any box design, its binarization induces lower expansion indices, which we denote by $\mathbf{E}^B = (E_1^B, ..., E_m^B)$.

Lemma 5. *For any box design $\mathbf{F} = (F_1, \ldots, F_m)$ such that $I_j \geq 0$ for all j, $E_j^B \leq E_j$, all j.*

Proof of Lemma 5. The proof is relegated to the Appendix as it is a bit long. The difficulty stems from the fact that each expansion index depends on the inspection and expansion indices of boxes encountered after the expansion under consideration. The result is thus established through an induction argument that leverages the property that binarizations, by shifting the probability assigned by the original distributions to values between 0 and the inspection indices to the lowest possible prize, reduce the expansion indices.

So far we have shown that the binarization $\mathbf{F}^B$ preserves each box's inspection index I_j, increases the probability that each box j, once opened, yields a realization below any cut-off

$\kappa \leq I_j/(1-\delta)$, and reduces all expansion indices E_l, $l = 1, ..., m$. Equipped with these results, we are now ready to establish that, to maximize her payoff, the designer can confine attention to sequences of binary distributions with support $\{0, \overline{v}\}$.

Proposition 1. *Let* $\mathbf{F}$ *be any design such that* $I_j \geq 0$ *all* j*. The designer's expected payoff under* $\mathbf{F}$*'s binarization* $\mathbf{F}^B$ *is at least as high as under* $\mathbf{F}$*.*

Proof of Proposition 1. The proof is in three steps which build on the lemmas established above.

Step 1. Under the index policy of Lemma 1, whenever the searcher opens box j and observes a realization $v_j > I_j/(1-\delta)$, he stops and chooses box j, ending the problem. For cut-offs $\kappa \in (I_j/(1-\delta), \overline{v})$, it is not guaranteed that, under the binarization of $\mathbf{F}$, $F_j^B(\kappa) > F_j(\kappa)$. However, note that at any history of past actions and prize discoveries at which the searcher opens box j under the original sequence $\mathbf{F}$, it must be that box j has the highest index (i.e., $E_\ell \leq I_j$, where E_ℓ is the expansion index at that history and, for any other box $l \neq j$, $I_l \leq I_j$, if box l has not been opened yet, and $(1-\delta)v_l \leq I_j$ if box l has been opened already). Hence, the relevant cutoff determining whether the searcher stops after opening box j is some $\kappa \leq I_j/(1-\delta)$: Any prize value v_j strictly above $I_j/(1-\delta)$ leads to immediate stopping. By virtue of Lemma 3, for any box j, the probability the searcher stops immediately after opening box j is smaller under $\mathbf{F}^B$ than under $\mathbf{F}$.

Step 2. Next observe that (a) each box's inspection index is the same under $\mathbf{F}$ and $\mathbf{F}^B$ (by definition of $\mathbf{F}^B$), (b) all expansion indices are non-negative under both designs (by virtue of the fact that the only cost of each expansion is the delay of any other decision by one period), and (c) $E_j^B \leq E_j$ for all j (by virtue of Lemma 5).

Because of (c), at each period prior to stopping, the set of boxes available to the searcher under $\mathbf{F}$ is a superset of those under $\mathbf{F}^B$. Hence, prior to stopping, the searcher's behavior may differ under the two designs either because (1) he opens a box under $\mathbf{F}^B$ whereas he expands under $\mathbf{F}$, or (2) he opens a box under both designs but the two boxes are not the same. For any $s = 1, ..., m$, let $I(s)$ and $I^B(s)$ denote the inspection indices of the s-th box opened respectively under the original design $\mathbf{F}$ and its binarization $\mathbf{F}^B$. That is, $I(1)$ is the inspection index of the first box opened under $\mathbf{F}$, whereas $I^B(1)$ that of the first box opened under $\mathbf{F}^B$, and likewise for all $s > 1$.[5] Next, for any $s = 1, ..., m$, let $t(s)$ and $t^B(s)$ denote the time at which the s-th opening occurs, respectively under $\mathbf{F}$ and $\mathbf{F}^B$, with these times equal to ∞ in case the opening never occurs. Properties (a)-(c) above, along with the fact that the searcher follows an index policy, imply that, for any $s = 1, ..., m$, $I(s) \geq I^B(s)$. This last property, together with Step 1 above and the fact that, under the binarization $\mathbf{F}^B$, the probability that a box's opening yields the prize $\bar{v}$ is increasing in the box's inspection index, imply the following are true for all $s = 1, ..., m$: (A) the probability the searcher stops immediately after the s-th opening is smaller under $\mathbf{F}^B$ than under $\mathbf{F}$, and (B) $t^B(s) \leq t(s)$.[6]

The above properties imply that the expected discounted gross payoff from the opening of the boxes is larger under $\mathbf{F^B}$ than under $\mathbf{F}$.

Step 3. To complete the proof, it suffices to show that the designer's expected cost is also weakly smaller under $\mathbf{F}^B$ than under $\mathbf{F}$. To see this, recall that, by the previous step, for any $s = 1, ..., m$, the inspection index of the s-th box opened under $\mathbf{F}^B$ is smaller than the inspection index of the s-th box opened under $\mathbf{F}$ (i.e., $I^B(s) \leq I(s)$ for all s). Consider the designer's expected cost when, at the s-th opening, the box's prize exceeds the box's inspection index (recall that, in this case, the searcher stops immediately after opening the box and walks away with the box's prize). Let l denote the expansion at which the searcher obtained the s-th box opened under $\mathbf{F}$ and r the expansion at which the searcher obtained the s-th box opened under $\mathbf{F}^B$. Observe that $r \leq l$ by virtue of the fact that, as explained above, the set of boxes available at any period prior to stopping to the searcher under $\mathbf{F}$ is

---

[5] For each $s = 1, ..., m$, the inspection index of the s-th box supplied under $\mathbf{F}$ is the same as under $\mathbf{F}^B$. However, when s denotes the order of the openings, the indices may differ due to the fact that, in some periods, the searcher may expand under one design while inspecting under the other. Also observe that, given the prize realizations, some boxes may never be opened. However, because there is no uncertainty over the sequence of distributions corresponding to each design, for any $s = 1, ..., m$, conditional on the s-th opening occurring, the index of the s-th opened box is uniquely pinned down by the selected design.

[6] That the probability the searcher stops immediately after the s-th opening is smaller under $\mathbf{F}^B$ than under $\mathbf{F}$ follows from the fact that (a) the solution to $h(p) = I$ is increasing in I, (b) $I(s) \geq I^B(s)$, and (c) for any $j = 1, ..., m$ and $\kappa \leq I_j/(1-\delta)$, $F_j^B(\kappa) \geq F_j(\kappa)$, as shown in Step 1 above.

a superset of those under $\mathbf{F}^B$. As established in Step 2 above,

$$I_l = \frac{-c + \delta \int_{\frac{I_l}{1-\delta}}^{\overline{v}} v \mathrm{d}F_l(v)}{1 + \frac{\delta}{1-\delta}\left(1 - F_l\left(\frac{I_l}{1-\delta}\right)\right)} \geq \frac{-c + \delta p_r^B \overline{v}}{1 + \frac{\delta}{1-\delta} p_r^B} = I_r^B,$$

with $p_r^B \leq p_l^B$ as $I_r^B \leq I_l^B = I_l$. By Lemma 3, $p_l^B \leq 1 - F_l\left(\frac{I_l}{1-\delta}\right)$ and hence $p_r^B \leq 1 - F_l\left(\frac{I_l}{1-\delta}\right)$, which implies that

$$\int_{\frac{I_l}{1-\delta}}^{\overline{v}} v \mathrm{d}F_l(v) \geq p_r^B \overline{v}.$$

Hence, conditional on the s-th opening revealing a flow payoff $(1-\delta)v$ exceeding the box's inspection index (which triggers the selection of the s-th opened box), the expected cost to the designer is smaller under $\mathbf{F}^B$ than under $\mathbf{F}$. When, instead, the s-th opening reveals a flow payoff $(1-\delta)v$ smaller than the box's index, the searcher may either stop and walk away with that box immediately after learning its prize, or return to it (to walk away with its prize) in future periods. In either case, the expected cost to the designer is higher under $\mathbf{F}$ than under $\mathbf{F}^B$. This is because, under $\mathbf{F}^B$, if the searcher ever selects a box whose prize is below its inspection index, this is because all opened boxes revealed a prize of zero, implying that the cost to the designer is the lowest possible one. Combining the two cases, we conclude that the expected cost at each opening is smaller under $\mathbf{F}^B$ than under $\mathbf{F}$. This completes the proof of the proposition. Q.E.D.

Without loss of optimality, all of the boxes can thus be assumed to have a binary distribution. The intuition is the following. One way to think of binarizations is that the inspection index captures the relevant "local attractiveness" statistic of a box. Once a box is about to be opened, the searcher cares about it only through its index and through the chance that its prize is high enough to trigger stopping, relative to the relevant cutoff. Lemma 2 shows that any distribution for which $I_j \geq 0$ can be replaced by a binary that preserves the inspection index. Lemmas 3-4 in turn show that the binary distribution puts less probability mass on prize realizations that trigger stopping than the original distribution. Intuitively, the probability one needs to put on the largest prize $\overline{v}$ to preserve the index is lower than the probability the original distribution puts above the index itself. Hence, the binarization of each box preserves the box's ex-ante appeal while decreasing the probability that, once opened, the box triggers stopping. Another advantage of binarizations is that, when the prize is below the index it is equal to the lowest possible value (zero), which also guarantees that the searcher will not prefer coming back to the box instead of opening other boxes in future periods.

We now proceed to characterize the optimal binary design. Let $\underline{p} \in (0,1)$ be the solution to $h(p) = 0$, i.e.,

$$\underline{p} = \frac{c}{\delta \overline{v}},$$

and note that any distribution with support $\{0, \overline{v}\}$ such that $\Pr(\tilde{v} = \overline{v}) = p < \underline{p}$ has a negative inspection index. Now let $\mathbf{F}^*$ be the box design such that each F_j^* is a binary distribution with support $\{0, \overline{v}\}$, with $\Pr(\tilde{v}_j = \overline{v}) = \underline{p}$.

Theorem 1. *The design* $\mathbf{F}^*$ *is optimal.*

Proof of Theorem 1. The result follows directly from the previous proposition along with the following observations: (a) any box with binary distribution F_j with support $\{0, \overline{v}\}$ such that $p_j = \Pr(\tilde{v}_j = \overline{v}) < \underline{p}$ is never opened; (b) the designer's payoff under any sequence of binary distributions $\mathbf{F}$, each with support $\{0, \overline{v}\}$, such that $p_j > \underline{p}$ for some j is lower than under the sequence $\mathbf{F}^*$. This is because, under $\mathbf{F}^*$, at each opening, the probability the searcher obtains $\overline{v}$ and thus stops is smaller than under $\mathbf{F}$. Furthermore, under $\mathbf{F}^*$, the searcher can be asked to open each box in the period immediately following its request, a property that is not guaranteed under $\mathbf{F}$. Finally, the expected cost of each opened box is smaller under $\mathbf{F}^*$ than under $\mathbf{F}$. Along with Proposition 1, jointly the above properties imply that the designer's expected payoff under $\mathbf{F}^*$ is higher than under any other design. Q.E.D.

The intuition for Theorem 1 is simple. Because of Proposition 1, the designer can restrict attention to sequences of binary distributions, each with support $\{0, \overline{v}\}$. Any such distribution is uniquely identified by the probability p_j it delivers the highest prize $\overline{v}$. The threshold $\underline{p} = c/(\delta\overline{v})$ is the lowest value of p_j for which the index of any such binary distribution is nonnegative (it also coincides with the lowest value of p_j for which the box's expected discounted prize, net of its opening cost, is positive). Any box with $p_j < \underline{p}$ is so unattractive that is never opened. Relative to $\mathbf{F}^*$ any sequence of binary distributions for which $p_j > \underline{p}$ for some j yields a lower payoff to the designer for three reasons. First, it may induce the searcher to pass on the opening of some of the received boxes with the intention of starting the exploration with more attractive boxes delivered after future expansions. Second, it may induce the searcher to stop with a higher probability conditional on stopping. Third, it may result in a higher cost to the designer. Because of these properties, the designer's payoff under $\mathbf{F}^*$ is as high as under any other design.

Fixing $\overline{v}$, the designer's ex-ante payoff is higher the lower c/δ. This is because a lower inspection cost c or a larger discount factor for the searcher imply a lower probability $\underline{p}$ each box delivers the prize $\overline{v}$. More interestingly, the designer's payoff is also increasing in the upper bound on the prize $\overline{v}$. This is because the probability $\underline{p}$ that each opening leads to stopping is lower the higher $\overline{v}$ is, whereas the expected cost at each opening $a\rho\underline{p}\overline{v} = a\rho c/\delta$ is invariant in $\overline{v}$. Also note that the condition $a < \delta/(\rho c)$ is both necessary and sufficient to guarantee that the designer's ex-ante payoff under $\mathbf{F}^*$ exceeds her outside option, as it implies that, at each opening, the expected net payoff is $1 - a\rho\underline{p}\overline{v} > 0$.

5 Simultaneous versus sequential box provision

Consider now a situation where all boxes are available to the searcher from the outset. In this environment, the searcher does not need to request boxes over time, but rather at each period, he simply chooses whether to stop or open one of the remaining unopened boxes. Thus, from the searcher's perspective, this is the classic Weitzman (1979) Pandora's boxes problem, except for the fact that the designer chooses the distributions of the prizes that maximize her net payoff.

The following example shows that, for a fixed heterogeneous pair of boxes, sequential provision can strictly dominate simultaneous provision, even when both boxes are already available to the designer at the outset. Suppose $m = 2$, $\bar{v} = 1$, $\delta = 1/2$, and $c = 1/20$, and let $a = 0$ in order to isolate the benefit of maximizing engagement. Let F_1 and F_2 be two binary distributions supported on $\{0, \bar{v}\}$, with $p_1 = 1/2$ and $p_2 = 9/10$, so that box 2 is more attractive than box 1. Using the formulas for the inspection indices in (3) we have that $I_1 = h\,(1/2) = 2/15 \approx 0.1333$ and $I_2 = h\,(9/10) = 4/19 \approx 0.2105$. Suppose the designer supplies the most attractive box 2 later as a way of keeping the searcher exploring. In this case, the expansion for box 2 is equal to $E_2 = [-\delta c + \delta^2 p_2]/[1 + \delta + \delta^2 p_2/(1-\delta)] = 4/39 \approx 0.1026$. Therefore, $I_2 > I_1 > E_2$. Under sequential provision, the searcher immediately opens box 1 upon receiving it. The expected number of openings is therefore $1 + \Pr(v_1 = 0) = 3/2$. Under simultaneous provision (i.e., when both boxes are made available to the searcher at the outset), the searcher opens box 2 first. The expected number of openings is thus $1 + \Pr(v_2 = 0) = 11/10$. Hence sequential provision yields a larger (undiscounted) number of openings. Hence, provided the designer is patient (namely, $\rho > 1/5$) sequential provision is preferable. This conclusion, however, is reversed under the optimal design, no matter the designer's discount factor:

Proposition 2. *The designer's payoff under simultaneous box provision is no smaller than under sequential box provision.*

Proof of Proposition 2. The result follows from the fact that the collection $\mathbf{F}^*$ of binary distributions with support $\{0, \bar{v}\}$ such that $p_j = \Pr(\tilde{v}_j = \bar{v}) = \underline{p}$ for all j is also optimal under simultaneous provision. To see this, note that, given any $\mathbf{F}$, its binarization $\mathbf{F}^B$ yields the designer a weakly higher payoff. The arguments are the same as under sequential provision: the binarization reduces the probability that the searcher stops after each opening and also reduces the expected cost to the designer in case a box is selected. In fact, the arguments are significantly simpler under simultaneous provision due to the fact that the binarization also preserves the order under which the boxes are opened. That the designer's payoff under $\mathbf{F}^*$ is higher than under any other sequence of binary distributions with supports $\{0, \overline{v}\}$ in turn follows from the fact that the probability the opening of each box induces stopping as well as the designer's expected cost are increasing in each p_j.

Hence, $\mathbf{F}^*$ is also optimal under simultaneous provision. That the designer's payoff is higher when all boxes are provided at the outset (i.e., in period 1) then follows from the fact that, under simultaneous provision, the opening of each box occurs earlier given that the searcher does not need to pause the exploration to request additional boxes. Q.E.D.

The next proposition considers the case where it takes the designer one period to provide each box, no matter whether all the boxes are provided before the searcher starts the exploration or each box is provided after the searcher requests it.

Proposition 3. *Suppose it takes the designer one period to provide each box. The designer's payoff is higher under sequential provision than under simultaneous provision (strictly if $\rho < 1$ and $m > 1$), despite the fact that, in both cases, all boxes are identical under the optimal design.*

Proof of Proposition 3. Without loss of optimality, the box design is $\mathbf{F}^*$ irrespective of whether the boxes are provided simultaneously or sequentially. Furthermore, because all boxes are identical to each other, without loss of optimality assume the searcher opens them in the order they are provided (i.e., box 1 is opened first, box 2 is opened second, etc.). Clearly, each box opened under simultaneous provision is also opened under sequential provision and vice versa (for any collection of prizes drawn from $\mathbf{F}^*$). However, the same sequence of openings happens earlier under sequential provision. To see this, note that under simultaneous provision the searcher can start opening boxes only in period m, after all boxes have been provided. So, if the search reaches the s-th opening, that opening occurs in period $m + s - 1$. Under sequential provision, instead, the first opening occurs in period 1, and each additional opening occurs after one period spent requesting the next box and one period waiting for it to arrive. Hence the s-th opening occurs in period $2s-1$. Since $2s-1 \leq m+s-1$ for every $s \leq m$, each opening occurs weakly earlier under sequential provision, and strictly earlier for every $s < m$.

That the designer is better off under sequential provision then follows from the above observations, along with the fact that, under $\mathbf{F}^*$, each opened box yields prize v with probability $\underline{p} = c/(\delta\overline{v})$. Thus, if a box is opened in period s, the designer gets gross payoff ρ^s from that opening, while the cost $\rho^{s+1}a\overline{v}$ is incurred only with probability $\underline{p}$. Therefore, the designer's expected payoff, net of the cost, from an opening in period s is $\rho^s - \rho^{s+1}a\underline{p}\overline{v} = \rho^s(1-\rho a\underline{p}\overline{v})$. Because $1-\rho a\underline{p}\overline{v} = 1-\rho ac/\delta > 0$ when $a < \delta/(\rho c)$, and because each opening occurs earlier under sequential provision, the designer's ex-ante payoff is higher under sequential provision (strictly so if $\rho < 1$ and $m > 1$). ■

6 Conclusions

The paper considers the design of an experimentation environment for a searcher who sequentially explores multiple alternatives of unknown quality before selecting one of them. The designer controls the distributions from which the searcher's payoffs are drawn as well as the order the alternatives are provided. Because expanding the set of alternatives takes time, the designer can induce the searcher to explore alternatives that are provided earlier even if those supplied later are more attractive. The designer seeks to maximize the searcher's engagement (i.e., the number of alternatives explored, discounted by the time at which the explorations occur) net of the cost of the selected option.

Perhaps surprisingly, we show that the designer does not benefit from making the alternatives heterogeneous and supplying the most attractive ones in later periods as a way of keeping the searcher engaged. The optimal design features a collection of binary distributions each yielding either the lowest or the highest feasible payoff. At any period prior to receiving the highest prize, the searcher is indifferent between (a) opting for the outside option, (b) opening the last box received, (c) asking for an additional box, and (d) walking away with any of the previously opened boxes.

When there are no costs in providing the boxes, the designer cannot do better than providing them all at the outset, as in Weitzman (1979)'s original problem. When, instead, supplying the boxes takes time, sequential provision is superior, for any non-trivial discount factor.

In future work, it would be interesting to extend the analysis to richer settings in which the cost of requesting additional boxes combines delay with a monetary payment or other physical costs. It would also be interesting to extend the analysis to problems in which the designer can make the searcher pay for the alternatives he explores, with the payment possibly depending on the order of the explorations.

References

Auster, S. and Y-K. Che (2025). "Pandora's Box Reopened: Robust Search and Choice Overload," working paper, Columbia University. **Bergemann, D. and S. Morris** (2019). "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57 (1), 44–95.

Fershtman, D. and A. Pavan (2025). "Searching for Arms: Experimentation with Endogenous Consideration Sets," working paper, Northwestern University and Tel Aviv University.

Gittins, J. C. (1979). "Bandit processes and dynamic allocation indexes," *Journal of the Royal Statistical Society. Series B (Methodological)* 41 (2), 148–177.

Gittins, J. and D. Jones (1974). "A dynamic allocation index for the sequential design of experiments. In J. Gani (Ed.)," *Progress in Statistics*, pp. 241-266. Amsterdam, NL: North- Holland.

Gossner, O., Steiner, J. and C. Stewart (2021). "Attention, please!," *Econometrica*, 89(4), 1717-1751.

Hebert, B. and W. Zhong (2025). "Engagement Maximization," working paper, Stanford

University.

Kamenica, E. (2019). “Bayesian Persuasion and Information Design,” *Annual Review of Economics*, 11, 249-272.

Weitzman, M. (1979). “Optimal search for the best alternative,” *Econometrica* 47 (3), 641–654.

7 Appendix

Proof of Lemma 5. We prove the lemma by backward induction on j.

Step 1 - “Induction base”: Consider the expansion index when $j = m$, that is, when there is only one box left to bring in. Below we show that $E_m^B \leq E_m$.

That $I_m \geq 0$ implies that I_m is given by the unique solution to

$$I_m = \frac{-c + \delta \int_{\frac{I_m}{1-\delta}}^{\overline{v}} v \mathrm{d}F_m(v)}{1 + \frac{\delta}{1-\delta}\left(1 - F_m\left(\frac{I_m}{1-\delta}\right)\right)}.$$

Since $I_m^B = I_m$, p_m^B is such that $h(p_m^B) = I_m$, which, by Lemma 3, implies that $p_m^B \leq 1 - F_m\left(I_m/(1-\delta)\right)$. That $I_m \geq 0$ in turn implies that E_m is given by the unique solution to

$$E_m = \frac{-\delta c + \delta^2 \int_{\frac{E_m}{1-\delta}}^{\overline{v}} v \mathrm{d}F_m(v)}{1 + \delta + \frac{\delta^2}{1-\delta}\left(1 - F_m\left(\frac{E_m}{1-\delta}\right)\right)}. \tag{4}$$

Similarly, the expansion index E_m^B is given by the unique solution to

$$E_m^B = \frac{-\delta c + \delta^2 p_m^B \overline{v}}{1 + \delta + \frac{\delta^2}{1-\delta} p_m^B}. \tag{5}$$

For any $x \in [0, \overline{v}]$, and any $G \in \Delta([0, \overline{v}])$, let

$$\kappa_x \equiv \frac{x}{1-\delta}$$

$$N_m^G(x) \equiv -\delta c + \delta^2 \int_{\kappa_x}^{\bar{v}} v \mathrm{d}G(v),$$

$$D_m^G(x) \equiv 1 + \delta + \frac{\delta^2}{1-\delta}\Big(1 - G(\kappa_x)\Big),$$

and

$$R_m^G(x) \equiv \frac{N_m^G(x)}{D_m^G(x)}.$$

Observe that E_m is given by the unique solution to $R_m^{F_m}(x) = x$, whereas E_m^B is given by the unique solution to $R_m^{F_m^B}(x) = x$. Also note that, for any $G \in \Delta([0, \overline{v}])$, any $x \in [0, \bar{v}]$,

$$R_m^G(x) \geq x \quad \Longleftrightarrow \quad \Gamma_m^G(x) \equiv N_m^G(x) - xD_m^G(x) \geq 0. \tag{6}$$

Next observe that

$$\begin{aligned}
\Gamma_m^G(x) &= \Big(-\delta c + \delta^2 \int_{\kappa_x}^{\bar{v}} v \mathrm{d}G(v)\Big) - x\Big(1 + \delta + \frac{\delta^2}{1-\delta}(1 - G(\kappa_x))\Big) \\
&= -\delta c - x(1+\delta) + \delta^2 \left[\int_{\kappa_x}^{\bar{v}} v \mathrm{d}G(v) - \kappa_x(1 - G(\kappa_x))\right] \\
&= -\delta c - x(1+\delta) + \delta^2 \int_{\kappa_x}^{\bar{v}} (v - \kappa_x)\, \mathrm{d}G(v).
\end{aligned}$$

Now, for any $\kappa \in [0, \bar{v}]$, any $G \in \Delta([0, \overline{v}])$, any $\ell \in \{1, \ldots, m\}$, let

$$S_\ell^G(\kappa) \equiv \int_{\kappa}^{\bar{v}} (v - \kappa)\, \mathrm{d}G(v). \tag{7}$$

Claim 1. For any $\ell \in \{1, \ldots, m\}$ and every $\kappa \leq \kappa_{I_\ell}$, $S_\ell^{F_\ell}(\kappa) \geq S_\ell^{F_\ell^B}(\kappa)$.

Proof. Recall that I_ℓ is the unique solution to

$$I_\ell = \frac{-c + \delta \int_{\kappa_{I_\ell}}^{\bar{v}} v \mathrm{d}F_\ell(v)}{1 + \frac{\delta}{1-\delta}\big(1 - F_\ell(\kappa_{I_\ell})\big)}.$$

Multiply both sides of the above expression by its denominator to obtain

$$I_\ell\Big(1+\tfrac{\delta}{1-\delta}(1-F_\ell(\kappa_{I_\ell}))\Big)=-c+\delta\int_{\kappa_{I_\ell}}^{\bar{v}} v\mathrm{d}F_\ell(v),$$

or, equivalently,

$$0=-c-I_\ell+\delta\int_{\kappa_{I_\ell}}^{\bar{v}} v\mathrm{d}F_\ell(v)-\delta\kappa_{I_\ell}\big(1-F_\ell(\kappa_{I_\ell})\big). \tag{8}$$

Next note the identity

$$\int_{\kappa_{I_\ell}}^{\bar{v}} v\mathrm{d}F_\ell(v)=\int_{\kappa_{I_\ell}}^{\bar{v}}\big(\kappa_{I_\ell}+(v-\kappa_{I_\ell})\big)\mathrm{d}F_\ell(v)=\kappa_{I_\ell}\big(1-F_\ell(\kappa_{I_\ell})\big)+\int_{\kappa_{I_\ell}}^{\bar{v}}(v-\kappa_{I_\ell})\mathrm{d}F_\ell(v).$$

Substituting this into (8), and using the definition of $S_\ell^{F_\ell}(\kappa_{I_\ell})$, we have that

$$I_\ell=-c+\delta S_\ell^{F_\ell}(\kappa_{I_\ell}), \tag{9}$$

or, equivalently,

$$S_\ell^{F_\ell}(\kappa_{I_\ell})=\frac{I_\ell+c}{\delta}.$$

Under F_ℓ^B, the only value strictly above κ_{I_ℓ} is $\bar{v}$, and it occurs with probability p_ℓ^B. Hence, from the definition of $S_\ell^{F_\ell^B}(\kappa_{I_\ell})$, we have that

$$S_\ell^{F_\ell^B}(\kappa_{I_\ell})=\int_{\kappa_{I_\ell}}^{\bar{v}}(v-\kappa_{I_\ell})\,\mathrm{d}F_\ell^B(v)=p_\ell^B(\bar{v}-\kappa_{I_\ell}).$$

Because $I_\ell=I_\ell^B$, $S_\ell^{F_\ell}(\kappa_{I_\ell})=S_\ell^{F_\ell^B}(\kappa_{I_\ell})$ and hence

$$S_\ell^{F_\ell}(\kappa_{I_\ell})=p_\ell^B(\bar{v}-\kappa_{I_\ell}). \tag{10}$$

Now pick any cutoff $\kappa \leq \kappa_{I_\ell}$ and observe that

$$S_\ell^{F_\ell}(\kappa) = \int_\kappa^{\kappa_{I_\ell}} (v-\kappa)\mathrm{d}F_\ell(v) + \int_{\kappa_{I_\ell}}^{\bar{v}} (v-\kappa)\mathrm{d}F_\ell(v). \tag{11}$$

The first integral in (11) is nonnegative. For the second integral, use the fact that $v - \kappa = (v - \kappa_{I_\ell}) + (\kappa_{I_\ell} - \kappa)$ to obtain

$$\begin{aligned}\int_{\kappa_{I_\ell}}^{\bar{v}} (v-\kappa)\mathrm{d}F_\ell(v) &= \int_{\kappa_{I_\ell}}^{\bar{v}} (v-\kappa_{I_\ell})\mathrm{d}F_\ell(v) + (\kappa_{I_\ell} - \kappa)\big(1 - F_\ell(\kappa_{I_\ell})\big) \\ &= S_\ell^{F_\ell}(\kappa_{I_\ell}) + (\kappa_{I_\ell} - \kappa)\big(1 - F_\ell(\kappa_{I_\ell})\big).\end{aligned}$$

Hence

$$S_\ell^{F_\ell}(\kappa) \geq S_\ell^{F_\ell}(\kappa_{I_\ell}) + (\kappa_{I_\ell} - \kappa)\big(1 - F_\ell(\kappa_{I_\ell})\big). \tag{12}$$

Because $p_\ell^B \leq 1 - F_\ell(\kappa_{I_\ell})$, as established in Lemma 3, conditions (12) and (10) jointly imply that

$$S_\ell^{F_\ell}(\kappa) \geq p_\ell^B(\bar{v} - \kappa_{I_\ell}) + (\kappa_{I_\ell} - \kappa)p_\ell^B = p_\ell^B(\bar{v} - \kappa) = S_\ell^{F_\ell^B}(\kappa),$$

as claimed. □

The above claim, when applied to $\ell = m$, implies that, for all $x \leq I_m$ (and hence for all $\kappa_x \leq \kappa_{I_m}$),

$$\Gamma_m^{F_m}(x) - \Gamma_m^{F_m^B}(x) = \delta^2 \left(S_m^{F_m}(\kappa_x) - S_m^{F_m^B}(\kappa_x) \right) \geq 0. \tag{13}$$

In turn, this implies that, for all $x \leq I_m$,

$$N_m^{F_m}(x) - xD_m^{F_m}(x) \;\geq\; N_m^{F_m^B}(x) - xD_m^{F_m^B}(x).$$

Recall that E_m^B solves $R_m^{F_m^B}(x) = x$, or, equivalently,

$$N_m^{F_m^B}(E_m^B) - E_m^B D_m^{F_m^B}(E_m^B) = 0.$$

Therefore, $N_m^{F_m}(E_m^B) - E_m^B D_m^{F_m}(E_m^B) \geq 0$, or, equivalently,

$$R_m^{F_m}(E_m^B) \geq E_m^B.$$

Because, for any G, $R_m^G(\cdot)$ is decreasing in x, we conclude that $E_m \geq E_m^B$, which completes the proof of the induction base.

Step 2–"Induction step": Fix $j < m$ and assume that $E_\ell^B \leq E_\ell$ for all $\ell = j+1, \ldots, m$. Below we show that the same inequality holds for j, i.e., that $E_j^B \leq E_j$.

Fix $x = E_j^B$ and let $\kappa_x = x/(1-\delta)$. For each design $\mathbf{G} \in \{\mathbf{F}, \mathbf{F}^B\}$, consider a fictitious search problem identical to the original one but where the boxes $k = 1, ..., j-1$ are not available to the searcher (i.e., the search is restricted to the boxes $k = j, ..., m$ and it starts with the searcher requesting box j). Let $\tau_j^{\mathbf{G}}(x)$ be the first time at which, in such a fictitious problem, when the searcher follows the index policy χ^* of Lemma 1, all of the following are weakly below x: (a) the expansion index, (b) the inspection indices of all boxes $\ell \geq j$ already requested, and (c) the flow values $(1-\delta)v_\ell$ of all boxes $\ell \geq j$ already opened. If the flow value $(1-\delta)v_\ell$ of an opened box exceeds x, let $\tau_j^{\mathbf{G}}(x) = \infty$.

Define

$$\begin{aligned}
N_j^{\mathbf{G}}(x) &\equiv \mathbb{E}_{\chi^*}^{\mathbf{G}}\left[\sum_{s=0}^{\tau_j^{\mathbf{G}}(x)-1} \delta^s U_s\right] \\
D_j^{\mathbf{G}}(x) &\equiv \mathbb{E}_{\chi^*}^{\mathbf{G}}\left[\sum_{s=0}^{\tau_j^{\mathbf{G}}(x)-1} \delta^s\right] \\
R_j^{\mathbf{G}}(x) &\equiv \frac{N_j^{\mathbf{G}}(x)}{D_j^{\mathbf{G}}(x)},
\end{aligned}$$

where the expectations are all under the index rule χ^*, with U_s denoting the generic s-th payoff under such a rule, in the fictitious problem described above, under the box design $\mathbf{G}$.

As in the base case, $R_j^{\mathbf{G}}(x) \geq x \iff N_j^{\mathbf{G}}(x) - xD_j^{\mathbf{G}}(x) \geq 0$. Define the "benchmark-$x$ net gain" under $\mathbf{G}$ by

$$\sum_{s=0}^{\tau_j^{\mathbf{G}}(x)-1} \delta^s \big(U_s - x\big)$$

and let

$$\Gamma_j^{\mathbf{G}}(x) \equiv N_j^{\mathbf{G}}(x) - xD_j^{\mathbf{G}}(x) = \mathbb{E}_{\chi^*}^{\mathbf{G}} \left[\sum_{s=0}^{\tau_j^{\mathbf{G}}(x)-1} \delta^s \big(U_s - x\big) \right].$$

Because $x = E_j^B$, $\Gamma_j^{\mathbf{F}^B}(x) = 0$. Thus, to prove that $E_j \geq x$ (and hence that $E_j^B \leq E_j$), it suffices to prove that

$$\Gamma_j^{\mathbf{F}}(x) \geq \Gamma_j^{\mathbf{F}^B}(x) = 0.$$

We now compare $\Gamma_j^{\mathbf{F}}(x)$ and $\Gamma_j^{\mathbf{F}^B}(x)$ at $x = E_j^B$. Consider any history before $\tau_j^{\mathbf{G}}(x)$ in the fictitious problem where the search is restricted to the boxes $\{j, ..., m\}$. If, under the index policy, at that history, the searcher opens some box $\ell \geq j$, then necessarily $I_\ell > x$, and hence $\kappa_x = x/(1-\delta) < I_\ell/(1-\delta) = \kappa_{I_\ell}$.[7] Using Lemma 1, we thus have that $S_\ell^{F_\ell}(\kappa_x) \geq S_\ell^{F_\ell^B}(\kappa_x)$.

Step 2a. In this step, we show that the contribution of the realizations $v_\ell > \kappa_x$ to the benchmark-x net gain

$$\sum_{s=0}^{\tau_j^{\mathbf{G}}(x)-1} \delta^s \big(U_s - x\big)$$

is weakly greater under $\mathbf{G} = \mathbf{F}$ than under $\mathbf{G} = \mathbf{F}^B$. Fix a history h at which box ℓ is

---

[7] Recall that $\tau_j^{\mathbf{G}}(x)$ is the first period at which, in the fictitious problem, all the following are weakly below x: (a) the expansion index for the next expansion, (b) the inspection indices I_ℓ of all boxes $\ell \geq j$ requested already but unopened, and (c) the flow values $(1-\delta)v_\ell$ of all boxes $\ell \geq j$ opened already. Hence, at the history under consideration, because the searcher follows the index policy χ^*, if box ℓ is opened, it must be that both the expansion index, the inspection indices of all unopened boxes and the flow values of all opened boxes are weakly below I_ℓ. By the definition of $\tau_j^{\mathbf{G}}(x)$, I_ℓ cannot be smaller or equal to x. Hence, $x < I_\ell$.

opened, and let $t = t(h)$ denote the date at which this opening occurs. Consider the event $H \equiv \{v_\ell > \kappa_x\}$. On this event, we have that $(1-\delta)v_\ell > x$, and hence $\tau_j^{\mathbf{G}}(x) = \infty$ for both $\mathbf{G} = \mathbf{F}$ and $\mathbf{G} = \mathbf{F}^B$. Moreover, starting from the next period, if the decision maker were to choose box ℓ forever, the per-period flow would be $(1-\delta)v_\ell$ and the benchmarked per-period excess would be $(1-\delta)v_\ell - x$. The discounted present value (from date t) of this benchmarked excess stream equals

$$\sum_{r=1}^{\infty} \delta^{t+r}\big((1-\delta)v_\ell - x\big) = \delta^{t+1}\frac{(1-\delta)v_\ell - x}{1-\delta} = \delta^{t+1}\big(v_\ell - \kappa_x\big).$$

Letting $\mathbf{1}\{W\}$ be the indication function (taking value 1 if W is true and 0 otherwise), and taking conditional expectations given the pre-opening history h, we have that

$$\mathbb{E}^{\mathbf{F}}\left[\delta^{t+1}(v_\ell - \kappa_x)\mathbf{1}\{v_\ell > \kappa_x\} \mid h\right] = \delta^{t+1}\int_{\kappa_x}^{\bar{v}} (v - \kappa_x)\,\mathrm{d}F_\ell(v) = \delta^{t+1} S_\ell^{F_\ell}(\kappa_x)$$

when $\mathbf{G} = \mathbf{F}$, and

$$\mathbb{E}^{\mathbf{F}^B}\left[\delta^{t+1}(v_\ell - \kappa_x)\mathbf{1}\{v_\ell > \kappa_x\} \mid h\right] = \delta^{t+1} S_\ell^{F_\ell^B}(\kappa_x)$$

when $\mathbf{G} = \mathbf{F}^B$. Because $S_\ell^{F_\ell}(\kappa_x) \geq S_\ell^{F_\ell^B}(\kappa_x)$, we conclude that, for every pre-opening history h at which ℓ is opened, the expected excess from the high-realization event $H \equiv \{v_\ell > \kappa_x\}$ is weakly larger under $\mathbf{F}$ than under $\mathbf{F}^B$:

$$\mathbb{E}^{\mathbf{F}}\left[\delta^{t+1}(v_\ell - \kappa_x)\mathbf{1}\{v_\ell > \kappa_x\} \mid h\right] \geq \mathbb{E}^{\mathbf{F}^B}\left[\delta^{t+1}(v_\ell - \kappa_x)\mathbf{1}\{v_\ell > \kappa_x\} \mid h\right].$$

Step 2b. We now show that the continuation after realizations weakly below x is also weakly more valuable, in terms of the future expected benchmark-x net gain,

$$\sum_{s=0}^{\tau_j^{\mathbf{G}}(x)-1} \delta^s\big(U_s - x\big)$$

under $\mathbf{F}$ than under $\mathbf{F}^B$. Consider the event $L \equiv \{v_\ell \leq \kappa_x\}$. Under the induction hypothesis, all future expansion indices are weakly larger under $\mathbf{F}$ than under $\mathbf{F}^B$, i.e., $E_\ell \geq E_\ell^B$ for all $\ell > j$. Furthermore, all inspection indices are identical under $\mathbf{F}$ and $\mathbf{F}^B$. Therefore, conditional on any post-opening history in L, the continuation problem under $\mathbf{F}$ weakly dominates the one under $\mathbf{F}^B$: the inspection indices are the same, box ℓ yields a weakly higher flow value under $\mathbf{F}$ than under $\mathbf{F}^B$ (under $\mathbf{F}^B$, the only low realization is 0), and, by the induction hypothesis, the expansion indices of all remaining boxes are weakly larger under $\mathbf{F}$. By Lemma 1 and the principle of optimality, the continuation payoff under the index rule from that history onward is exactly the value of this continuation problem. Hence the continuation benchmark-x net gain from that history onward is weakly larger under $\mathbf{F}$ than under $\mathbf{F}^B$.

Combining Steps 2a and 2b. Fix any pre-opening history h at which box ℓ is opened at date t, and let

$$Z_{t+1}^{\mathbf{G}} \equiv \sum_{s=t+1}^{\tau_j^{\mathbf{G}}(x)-1} \delta^s (U_s - x)$$

denote the remaining benchmark-x net gain from $t+1$ onward. Let $H \equiv \{v_\ell > \kappa_x\}$ and $L \equiv \{v_\ell \leq \kappa_x\}$. Then for each box design $\mathbf{G} \in \{\mathbf{F}, \mathbf{F}^B\}$,

$$\mathbb{E}_{\chi^*}^{\mathbf{G}}[Z_{t+1}^{\mathbf{G}} \mid h] = \mathbb{E}_{\chi^*}^{\mathbf{G}}[Z_{t+1}^{\mathbf{G}} \mathbf{1}_H \mid h] + \mathbb{E}_{\chi^*}^{\mathbf{G}}[Z_{t+1}^{\mathbf{G}} \mathbf{1}_L \mid h].$$

On event H, under F^B, one necessarily has that $v_\ell = \overline{v}$, so, by Lemma 1, the searcher stops immediately and chooses box ℓ. Hence,

$$Z_{t+1}^{F^B} \mathbf{1}_H = \delta^{t+1}(v_\ell - \kappa_x)\mathbf{1}_H. \tag{14}$$

Under F, the expression on the right-hand side of (14) is a lower bound on $Z_{t+1}^{F} \mathbf{1}_H$ because, after observing $v_\ell > \kappa_x$, the searcher can always choose box ℓ at $t+1$. Using Step 2a we then conclude that

$$\mathbb{E}^{F}_{\chi^*}[Z^{F}_{t+1}\mathbf{1}_H|h] \geq \mathbb{E}^{F^B}_{\chi^*}[Z^{F^B}_{t+1}\mathbf{1}_H|h].$$

Step 2b in turn implies that $\mathbb{E}^{\mathbf{F}}_{\chi^*}[Z^{\mathbf{F}}_{t+1}\mathbf{1}_L \mid h] \geq \mathbb{E}^{\mathbf{F}^B}_{\chi^*}[Z^{\mathbf{F}^B}_{t+1}\mathbf{1}_L \mid h]$. We thus have that, for every pre-opening history h,

$$\mathbb{E}^{\mathbf{F}}_{\chi^*}[Z^{\mathbf{F}}_{t+1} \mid h] \geq \mathbb{E}^{\mathbf{F}^B}_{\chi^*}[Z^{\mathbf{F}^B}_{t+1} \mid h].$$

Taking expectations over h (law of iterated expectations) and using the fact that the above property holds for all possible opening dates, we conclude that $\Gamma^{\mathbf{F}}_j(x) \geq \Gamma^{\mathbf{F}^B}_j(x)$. Because $\Gamma^{\mathbf{F}^B}_j(x) = 0$, we thus have that $\Gamma^{\mathbf{F}}_j(x) \geq 0$. Because $R^{\mathbf{F}}_j(x) \geq x \iff \Gamma^{\mathbf{F}}_j(x) \geq 0$, we have that $R^{\mathbf{F}}_j(x) \geq x$. Because $R^{\mathbf{F}}_j(\cdot)$ is decreasing in x and E_j is the unique solution to $R^{\mathbf{F}}_j(x) = x$, we conclude that $E_j \geq x = E^B_j$, which completes the induction step.

This completes the proof of the lemma. ☐