

Doval, Laura

Working Paper

Whether or not to open Pandora's box

Discussion Paper, No. 1574

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Doval, Laura (2013) : Whether or not to open Pandora's box, Discussion Paper, No. 1574, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/97253>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



CMS-EMS
Center for Mathematical Studies in Economics
And Management Science

Discussion Paper #1574

Whether or not to open Pandora's box

Laura Doval

Northwestern University

November 2013

JEL Classification: D83

Keywords: search, information acquisition



NORTHWESTERN
UNIVERSITY



Whether or not to open Pandora's box

Laura Doval^{*†}

First draft: November 2012. This draft: November 2013

Abstract

Weitzman's [15] search model requires that, conditional on stopping, the agent only takes boxes which have already been inspected. We relax this assumption and allow the agent to take any uninspected box without inspecting its contents when stopping. Thus, each uninspected box is now a potential outside option. This introduces a new trade-off: every time the agent inspects a box, he loses the value of the option to take it without inspection. Nevertheless, we find that, under conditions common in the search and information acquisition literature, boxes are inspected following the same order as in Weitzman's rule; however, the stopping rule is different, and we characterize it. Moreover, we provide additional results that partially characterize the optimal policy when these conditions fail.

KEYWORDS: search, information acquisition

JEL CLASSIFICATION: D83

^{*}I am very thankful to Eddie Dekel for useful discussions and his detailed comments. I am particularly grateful to Jeff Ely for his comments and encouragement when this project was in its formative stages, and for his continuing guidance and support. I also wish to thank Wojciech Olszewski, Mallesh Pai, and Pablo Schenone for useful discussions. All errors are, of course, my own.

[†]Ph.D. Candidate in Economics, Department of Economics, Northwestern University, Evanston, IL 60208. (e-mail:laura.doval@u.northwestern.edu)

1 Introduction

Weitzman’s [15] model has been used in a wide variety of situations (see Olszewski [11], and the references therein) that fit the following framework: an agent possesses N boxes, each of which contains an unknown prize, he can search for prizes sequentially at a cost, and search is with recall. Weitzman characterizes the optimal search rule, which is defined by an *order* in which boxes are inspected, and a *stopping rule*: boxes are assigned reservation values, they are inspected in descending order of their reservation values, and search stops when the maximum sampled prize is greater than the maximum reservation value amongst uninspected boxes. An assumption in Weitzman [15] is that the agent cannot take a box without first inspecting its contents. This assumption is responsible for the simplicity of the rule, and it restricts the applications of the model (see Section 2 for examples). Our paper addresses Weitzman’s search problem without this assumption. Within this framework, we show that, under conditions commonly used in the search and information acquisition literature, the optimal *order* coincides with Weitzman’s. However, the optimal *stopping rule* is different and we characterize it (see Section 1.1, and Propositions 1, 2 and 3 in Section 5).

Before discussing our results in detail, consider the following example. Say the agent is a student who has to make a choice between schools A and B , to which he has been admitted to, or not going to school (call this option Z). The student has the option of attending the visit day at each institution and finding out how suitable a match the school is. This requires effort and time, which are costly to the agent. We interpret each school as a box, how good a match the school is as the prize in the box, attending the visit day as inspecting a box, and the effort and time invested as the box’s inspection cost. Weitzman’s assumption implies the agent can only choose from programs whose visit day he has attended.

To make the example more concrete assume that each school’s distribution over prizes is given by the following table, based on an example by Postl [13]:

A	Prize	1	2	5	Inspection cost
	Probability	0.25	0.50	0.25	0.25
B	Prize	0		3	Inspection cost
	Probability	0.50		0.50	0.25
Z	Prize		x_Z		Inspection cost
	Probability		1		0

Table 1: PRIZE DISTRIBUTION FOR EACH SCHOOL

Assume, for now, that $x_Z = 0$. Under Weitzman’s assumption, it is easy to show that the following is the optimal policy: school A should be visited first; if the prize is $x_A = 5$ search stops, while if the prize is $x_A \in \{1, 2\}$, the agent visits school B and chooses the school with the highest realized prize. To see this, first, note that if search starts with school A , it should stop if the realized prize is $x_A = 5$: visiting school B is costly, and the highest feasible prize is $x_B = 3$. Second, if the realized prize at school A is $x_A \in \{1, 2\}$, it is worth visiting school B as well: visiting foregoes $x_A + 0.25$ which is outgained by the expected gain $0.5 \times x_A + 0.5 \times 3$. Third, notice it is optimal to visit at least one school: visiting school A alone yields expected utility $2.5 - 0.25 = 1.75 > 0$. Fourth, it is optimal to visit school A first: since the agent may obtain a prize $x_A = 5$, and prizes can’t be collected without inspection, it is optimal to start search with A . Moreover, it can be shown that, for any x_Z , if it is optimal to visit any school, the agent should visit school A first; the inspection order is independent of the value of x_Z .

However, a student always has the option to accept admission to a school without attending its visit day, and this may be optimal. In this example, taking into account this option changes the optimal search rule when $x_Z = 0$. In this case, school B should be visited first; if the prize is $x_B = 0$ search stops, and school A is selected without inspection, while if the prize is $x_B = 3$, the agent visits school A and chooses the school with the highest realized prize. We show this in five steps. First, conditional on visiting school B first, if $x_B = 0$, it is optimal to accept school A without visiting it: since $x_A > 0$, school A dominates school B with probability 1. Second, by a similar argument as before, conditional on $x_B = 3$, visiting school A and selecting the best school is optimal. Third, if the agent visits school A first, the optimal continuation policy is the one in the previous paragraph: (i) if $x_A = 5$, it is clear that school A dominates school B , (ii) if $x_A = 2$, visiting school B and

choosing the best school dominates stopping and accepting school A (which is the best alternative conditional on stopping), and (iii) if $x_A = 1$, visiting school B and choosing the best school dominates accepting school B without visiting it. Fourth, it is optimal to visit at least one school: visiting school A first, and following the optimal continuation, dominates accepting school A without visiting it. Finally, it is optimal to visit school B first: visiting school A has a benefit (he might discover $x_A = 5$ and can save the cost of visiting school B) and a cost (he could have visited school B , learned that $x_B = 0$, and saved the inspection costs of A). The parameters in the example imply that the cost of visiting A first outweighs the benefit. Since at least one school should be visited (by step 4), the agent must start search with B . Steps 1 and 2 complete the description of the optimal search rule.

We conclude the example with two remarks. First, we could have tried an alternative method to prove that the above rule is optimal. Namely, create, for each school, a *virtual* school with zero inspection cost, and a prize equal to the mean prize in that school. Then, treat this problem as a search problem with four schools (boxes), and use similar arguments as in the first part of the example. Section 1.3 discusses why this approach does not work in our setting. Second, note that, contrary to Weitzman's model, in which the value of x_Z only determines whether it is optimal to search or not, the optimal search order may also depend on x_Z in our setting. In particular, if $x_Z = 2.1$, the optimal policy coincides with Weitzman's. When $x_Z = 0$, and the agent finds out that $x_B = 0$, visiting school A has no value. By contrast, when $x_Z = 2.1$, and the agent finds out that $x_B = 0$, then, by not committing to accepting school A without visiting it first (which he values at 2.5), the agent has the option of choosing whatever is best between going to school A , and not going to school at all (exercising this option yields utility $\frac{1}{4} \times 5 + \frac{3}{4} \times 2.1 = 2.575$). Since the agent visits school A regardless of x_B , the same arguments as before show that Weitzman's policy is optimal.

1.1 Discussion of results

In our model, each box is characterized by *two* cutoff values, defined formally in Section 3.2 (see equations (1)-(2)), and represented in Figure 1 below. The first, the reservation value (denoted by x^R), is the value of the maximum previously sampled prize that would make the agent indifferent between inspecting the box and taking the sampled prize. It reflects the trade-off between exploration and exploitation: by inspecting the box, the agent may obtain a higher reward (exploration), but this comes at a cost since inspection is not free (exploitation). It is not optimal to inspect the box if the maximum sampled prize is above the reservation value. The second cutoff, which we call the box's *backup value* (denoted by x_B), is the value of the maximum previously sampled prize that would make the agent indifferent between inspecting the box and taking it without inspection. It reflects the trade-off between insurance and exploration: by taking the box without inspection, the agent receives a certain expected payoff without paying the inspection costs (insurance), but by inspecting the box, he learns its contents (exploration). The agent takes the box without inspecting it first if the maximum sampled prize is below the box's backup value.

Figure 1 illustrates the search/stopping regions as a function of the maximum sampled prize, \bar{z} , when there is only one box left to inspect with mean μ :




Figure 1: OPTIMAL POLICY FOR ONE BOX

In general, the order of x_B, x^R is not determined; though, it must be the case that $\mu \in [\min\{x_B, x^R\}, \max\{x^R, x_B\}]$. Assumption 1 in Section 3.2 implies that $x_B < x^R$. This assumption rules out uninteresting cases (see Appendix A.4 for details).

Without further assumptions, the solution for one uninspected box cannot be applied inductively to the case of N boxes. While Weitzman [15] shows that boxes

should be inspected in decreasing order of their reservation values, this is not necessarily the case in our model. Whenever the agent decides which box he inspects next, he foregoes the option of taking that box without inspection. The agent would like to inspect boxes in increasing order of their backup values, since boxes with higher backup values are the ones that, without inspection, he values the most. Since the boxes with high reservation values might also be the boxes with high backup values, this introduces a trade-off when deciding which box to sample next; how this trade-off is resolved depends on the available boxes and the already sampled prizes, and determines in which order boxes are inspected. Hence, the problem of finding the optimal search policy without imposing further assumptions is intractable.¹

Our contribution lies in identifying conditions under which the optimal *order* policy coincides with Weitzman’s. These conditions are common in the search and information acquisition literature. However, the optimal *stopping rule* is different, and we characterize it (see Section 5, Propositions 1-3). We also provide a number of additional results that partially characterize the optimal policy when the conditions of Propositions 1-3 fail (see Section 4, Lemmas 1-3).

Proposition 1 in Section 5 states that if, given any two boxes i, j , box i has a higher reservation value than box j if, and only if, box i has a lower backup value than box j , then the optimal order policy coincides with Weitzman’s [15], and his stopping rule applies to all but the last box. As explained, Weitzman’s order requires inspecting first boxes with high reservation values, while boxes with high backup values are the best to take without inspection. Therefore, if the box with the highest reservation value is the box with the lowest backup value, then, when inspecting this box, the agent never foregoes the option of taking without inspection his best backup. Proposition 1 holds if, for example, given any two boxes the prize distribution of one box is obtained by a *mean preserving spread* of the prize distribution of the other, and all boxes share the same inspection cost. On the one hand, boxes with higher “variance” are better for inspection since the agent

¹Indeed, the search problem can be cast as a *multi-armed restless bandit*, i.e. a multi-armed bandit where, conditional on pulling one arm, the states of all arms may change. This class of bandit problems is shown to be P-SPACE hard in [12]. Indeed, our paper may be seen as a contribution to the restless bandit literature by providing instances of a restless bandit problem where an exact solution can be computed.

can get better draws; on the other, these are the boxes that are not good backups: they can also contain worse draws. The mean preserving spread assumption is used in Vishwanath [14] to obtain the reservation value rule in her parallel search model, while Chade and Smith [4] apply it in their simultaneous search model.

Under the conditions for Propositions 2 and 3, given any two boxes i, j , box i has a higher reservation value than box j if, and only if, box i has a higher backup value than box j . The conditions are: (i) all boxes have the same binary prizes, same inspection cost, and differ on the probability of obtaining the highest prize (Section 5, Proposition 2), or (ii) there are only two boxes which share the same inspection cost, and prizes normalized by their mean are distributed according to the same symmetric distribution (Section 5, Proposition 3). The conditions in (i) allows us to extend Chade and Smith’s [4] simultaneous search model with binary prizes to our setting. Indeed, while their model is well suited to analyze the decision of which colleges to apply to, our model can be used to determine how to sequentially acquire information on the schools the agent has been admitted to.

Unlike Proposition 1, where the optimal stopping rule coincides with Weitzman’s for all but the last box, the agent may choose to stop, and take a box without inspection, before reaching the last box under the conditions of Propositions 2 and 3. Stopping and taking a box without inspection is optimal in (i) whenever all realized prizes coincide with the lowest possible prize, the maximum realized prize is less than the highest backup value amongst uninspected boxes, and the value of inspecting the box with the highest reservation value and proceeding according to our rule is less than this box’s backup value. The rule’s complexity, compared to Weitzman’s, is because continuation values depend on the combination of realized prizes and backup values of uninspected boxes.

Our results, which analyze cases in which Weitzman’s order remains optimal, are valuable for two reasons. The first is tractability: when Weitzman’s order is not optimal, the problem becomes intractable and not useful for applications. Second, the conditions on primitives that are sufficient for our results to hold coincide with assumptions that are common in the search and information acquisition literature. In those settings, Weitzman’s order is optimal without loss of generality.

Besides our results on when Weitzman’s order policy is still optimal, we provide a number of additional results that partially characterize the optimal policy when

the conditions of Propositions 1-3 fail, and are presented as lemmas (see Section 4). Lemmas 1 and 2 in Section 4 illustrate the role of the backup value in this search problem. Indeed, if the maximum sampled prize is higher than the backup value (not the mean) of the remaining uninspected boxes, then Weitzman’s sampling policy is optimal in the continuation search problem, and if it is optimal to stop and take a box without inspection, it has to be that this box is the one with the highest backup value. Lemma 3 illustrates the nature of deviations from Weitzman’s ordering of boxes.

1.2 Related Literature

To the best of our knowledge, this is the first attempt to provide a solution to this search problem in its general form, without restricting the number of boxes or the inspection costs. Postl [13] postulates this search problem explicitly within the context of a principal-agent model. He focuses on the two-box version of our search problem, and assumes boxes have equal inspection costs. He discusses an analogue of Proposition 1 in this simplified setting. Proposition 1 in our paper generalizes this result, showing that it is not necessary to assume two boxes nor that the boxes have equal costs.

Regarding the search literature, the most relevant paper for us is Klabjan, Olszewski and Wolinsky [7], who study optimal search for attributes. Contrary to our setting, the agent’s utility function is given by the sum of the prizes (attributes) in their model. The agent has to choose whether or not he acquires an object that is characterized by a series of attributes (boxes), which he can inspect by paying a cost. The agent does not have to inspect all attributes in order to make his decision: he may decide to accept the object, taking the rest of the attributes without inspection. Under sequential search, the authors characterize the optimal solution when there are only two attributes whose distributions are symmetric around 0. The rule coincides with inspecting attributes in decreasing order of their reservation value (see Section 5, Proposition 3, for a similar result in our setup).

The double-sided stopping rule in Figure 1 has appeared in previous work in the mechanism design literature (Chade and Kovrijnykh [3], Krahmer and Strausz [8]),

though none of these papers provide a solution for the search problem analyzed here. It also appears in Klabjan, Olszewski and Wolinsky [7] up to a relabelling of what the outside option (\bar{z}) is in their paper.

1.3 A “solution”

To further explain the difficulties involved, it is useful to present what might appear to be a natural approach to solving our search problem. We then explain why this approach is flawed. Assume there are two boxes. For each box $i \leq 2$, create a *virtual box* \hat{i} which has zero inspection cost, and contains a prize equal to the mean prize in box i . This construction leads to a set of 4 boxes. Then one could apply Weitzman’s rule to the new set of boxes. This “solution” is intuitive: virtual box \hat{i} represents the option of taking box i without inspection. Indeed, if the agent does so, he does not pay the inspection cost of box i , and he receives a prize, in expectation, equal to the mean prize in box i .

Assume, without loss of generality, that the four boxes’ reservation values are such that Weitzman’s rule indicates inspecting box 1 first, then box $\hat{1}$, and then the remaining boxes.² Moreover, assume that after inspecting box 1, the prize is below the reservation value of box $\hat{1}$. Weitzman’s rule indicates the agent should inspect box $\hat{1}$ next. However, inspecting box $\hat{1}$, that is, taking box 1 without inspection, is no longer an option for the agent.

There is another reason why this approach is flawed. Since the virtual boxes are free to take, and there is no uncertainty regarding their contents, the reservation values of these boxes coincide with the prize inside the box. This is saying that the value of the option of taking the box without inspection coincides with the mean value of the box’s prize distribution. Part of our contribution is to show that this option value is given by the backup value of the box, and not the mean of the prize inside the box (Section 4, Lemmas 1 and 2).

The rest of the paper is organized as follows. Section 2 illustrates several applications where this relaxation of Weitzman’s model is natural and relevant; it may be skipped without loss of continuity. Section 3 describes the model, provides a

²Indeed, this is the case in the school choice example outlined before.

formal definition of the cutoffs and intuition of their role in the search problem. Section 4 provides a series of properties the optimal sampling policy must satisfy regardless of the environment. Section 5 focuses on the optimal order and stopping policies, and contains our main results. The statements of Propositions 1-3 are presented informally to streamline notation; the Appendix contains the formal statements. Section 6 concludes. Proofs are relegated to the Appendix.

2 Examples

Besides the school choice example already outlined, this extension of Weitzman’s search problem is useful in other settings. We consider three examples in detail, though it is clear that our relaxation is useful in other settings. The examples are based on existing papers in the search and information acquisition literature. This highlights that our relaxation is useful for settings that are already being studied, and may provide new insights to these. It is worth noting that none of these papers solve for our search problem within their setup. The authors either (i) consider a one-box setup (first example), (ii) analyze a mechanism design setting where incentive compatibility implies the mechanism designer learns the unknown values of the prizes without having to inspect the boxes (second example), or (iii) they maintain Weitzman’s setting (third example).

The first example is pre-project planning (see Krahmer and Strausz [8], and the references therein). An agent wants to complete a project, which he values v . He can complete it with expected cost μ . The agent can invest in pre-project planning to observe the true production cost x . Pre-project planning costs $k > 0$. If the agent has the option to complete the project without further investigations, then he would only acquire information for intermediate values of the mean cost. If the mean cost is either too low or too high, then costly information acquisition is deemed unnecessary. If, on the contrary, he can only complete the project upon acquiring information, he would choose to acquire information more often. In particular, whenever he was willing to complete the project without acquiring information, he now acquires information first. In this sense, Weitzman’s model makes the agent acquire too much information.

The second example is the mechanism design model with costly state verification of Ben Porath, Dekel and Lipman [2]. A dean has one job slot to allocate to one of N departments. Each department has private information about the quality of its candidate, and wants his candidate to obtain the job slot regardless of the candidate's quality. The dean can verify at a cost the quality of any department's candidate. In terms of our setting, the dean is the agent, the departments are the boxes, the candidate's quality is the prize inside the box, and the verification cost is the cost of inspecting the box. As long as the dean has the option of allocating the job slot to a department without verifying its candidate's quality, the dean faces the same decision problem as the one presented in our paper. It is worth noting that what we define as a backup value plays a key role in the optimal mechanism characterized in [2], despite the differences in the setup. Indeed, if all departments have the same inspection cost, then the optimal mechanism assigns the slot to the department with the highest backup value, if all announced values are below this cutoff; it requires the principal to inspect any department that announces a value above this cutoff, and assign it the slot if the report is truthful ³.

The third example is the social search model of Mueller-Frank and Pai [10]. In their model, each member of a group selects sequentially and irreversibly an action from a finite set. Each agent freely observes the actions already chosen by others before making his choice. Actions are ex-ante identical, and their utilities are unknown. Each agent, in order to select an action, must pay a cost and find out its value. That is, each action represents a box in terms of our model, and each agent faces a search problem that resembles Weitzman's when choosing which action to select. If we consider the case in which the different actions correspond to different restaurants, agents observe other agents' choices by observing their check-ins at Yelp, and learning a restaurant's quality corresponds to reading the online reviews, then it is natural to assume that an agent can choose to go to a restaurant without reading the reviews.

By similar relabellings the search deterrence model of Armstrong and Zhou [1], and the limited attention model of De Clippel, Eliaz and Rozen [5] can be cast in terms of our setup, and represent other applications where the assumption of

³A similar intuition applies if the departments have different inspection costs, but now the cutoff depends on both the backup value and the inspection cost.

being able to choose an option without inspection is natural.

3 Model and preliminary lemmas

An agent possesses a set $\mathcal{N} = \{1, \dots, N\}$ of boxes, each containing a prize, x_i , distributed according to distribution function F_i , with mean μ_i . Box i has inspection cost k_i . F_i and k_i are known, however x_i is not. Prizes are distributed independently of one another, and we assume that, for all $i \in \mathcal{N}$, $\int |x_i| dF_i(x_i) < +\infty$. The agent has an outside option, x_0 , which is normalized to 0. We assume the agent is risk neutral, and given a vector of realized prizes $z = (z_1, \dots, z_n)$, his utility function is given by $u(z) = \max_{j \in \{1, \dots, n\}} z_j$.

3.1 Sampling Policy

The agent inspects boxes sequentially, and search is with recall. Given a set of boxes \mathcal{U} (“U” stands for uninspected), and a vector of already sampled prizes z , the agent decides which box to inspect first, if any. Let $\sigma(\mathcal{U}, z)$ denote such box. If he inspects box σ , he pays k_σ , and observes its prize, x_σ . Having observed x_σ , the agent decides whether to stop search, which we denote by $\varphi(\mathcal{U}, z \cup \{x_\sigma\}) = 1$, or to continue search ($\varphi = 0$). If $\varphi = 1$, the agent chooses between x_σ , any prize in z , and any uninspected box in $\mathcal{U} \setminus \{\sigma\}$. If $\varphi = 0$, the agent chooses which box to inspect next out of the boxes in $\mathcal{U}' \equiv \mathcal{U} \setminus \{\sigma\}$, taking into account that the vector of realized prizes is $z \cup \{x_\sigma\}$. Given a vector z , we denote by \bar{z} , its highest coordinate. Given (\mathcal{U}, z) , the strategy σ, φ , together with the distributions $\{F_i\}_{i \in \mathcal{U}}$ determine a probability distribution over continuation paths in the natural way. Throughout, we denote with stars the optimal strategies.

Given a decision node described by (\mathcal{U}, z) , each strategy σ, φ determines the agent’s expected payoff at that decision node, which we denote $V(\mathcal{U}, z)$. We use stars to denote the payoff V when it results from using the optimal policy in (\mathcal{U}, z) .

3.2 Cutoffs

Suppose the agent has only one box left to inspect, i , with expected value μ_i . Recall \bar{z} is the maximum sampled prize, i.e., his outside option. Suppose $\bar{z} > \mu_i$; hence, conditional on stopping the agent would take \bar{z} . The agent inspects the last box if, and only if, the following holds:

$$\bar{z} \leq \int_{-\infty}^{\bar{z}} \bar{z} dF_i(x_i) + \int_{\bar{z}}^{+\infty} x_i dF_i(x_i) - k_i \Leftrightarrow k_i \leq \int_{\bar{z}}^{+\infty} (x_i - \bar{z}) dF_i(x_i)$$

Define the box's *reservation value* to be the number x_i^R such that:

$$k_i = \int_{x_i^R}^{+\infty} (x_i - x_i^R) dF_i(x_i) \quad (1)$$

i.e., x_i^R is the value of the outside option that leaves the agent indifferent between stopping and taking prize x_i^R , and inspecting box i . The agent inspects the last box whenever $\bar{z} \leq x_i^R$.

Consider now the case in which $\bar{z} \leq \mu_i$. Thus, if the agent stops, he takes the box without inspection. Therefore, the agent inspects the last box if, and only if the following holds:

$$\begin{aligned} \mu_i &\leq -k_i + \int_{-\infty}^{\bar{z}} \bar{z} dF_i(x_i) + \int_{\bar{z}}^{+\infty} x_i dF_i(x_i) \\ \Leftrightarrow \int_{-\infty}^{+\infty} x_i dF_i(x_i) &\leq -k_i + \int_{-\infty}^{\bar{z}} \bar{z} dF_i(x_i) + \int_{\bar{z}}^{+\infty} x_i dF_i(x_i) \\ \Leftrightarrow k_i &\leq \int_{-\infty}^{\bar{z}} (\bar{z} - x_i) dF_i(x_i) \end{aligned}$$

The first line of the above expression makes clear that, conditional on inspecting the box, the agent loses the option of getting a payoff equal to μ_i . Indeed, the possible prizes are either \bar{z} , or the new sampled prize in the expression on the right hand side. Define the *backup value* to be the value x_{iB} such that:

$$k_i = \int_{-\infty}^{x_{iB}} (x_{iB} - x_i) dF_i(x_i) \quad (2)$$

i.e., x_{iB} is the value of the outside option that leaves the agent indifferent between inspecting box i and taking it without inspection. The agent inspects box i if $x_{iB} \leq \bar{z}$; otherwise, he takes it without inspection. Throughout, we make the following assumption to ensure that $x_{iB} \leq \mu_i \leq x_i^R$ always holds:

Assumption 1. $(\forall i \in \mathcal{N}) : k_i \leq \int_{-\infty}^{\mu_i} (\mu_i - x_i) dF_i(x_i)$

If the set of boxes \mathcal{N} contains a (at least one) box that violates Assumption 1, then said box is (boxes are) never inspected (see Appendix A.4 for a proof, and for a discussion of why Assumption 1 rules out uninteresting cases).

Notice that, when there is one box left, the optimal policy is determined by comparing the maximum sampled prize, \bar{z} , with the cutoffs, x_B, x^R . This is summarized in Figure 1 in Section 1.1, and is recorded in Lemma 0 below.

Lemma 0. Assume that $N = 1$, and let \bar{z} be the agent's outside option. The optimal policy is as follows:

1. If $\bar{z} < x_B$, the agent takes the box without inspection.
2. If $x_B \leq \bar{z} \leq x^R$, the agent inspects the box, and keeps the maximum prize between \bar{z} and the sampled prize x .
3. If $x^R < \bar{z}$, the agent does not inspect the box, and keeps his outside option.

Results similar to Lemma 0 have appeared in the one-box-settings of Chade and Kovrijnykh [3], and Krahmer and Strausz [8], in the two-box setting of Postl [13], as well as in the attributes model of Klabjan, Olszewski and Wolinsky [7].

We can use Lemma 0 to provide an intuition for why we refer to x_{iB} as box i 's backup value. The next table represents, for different values of the outside option \bar{z} , the payoff the agent obtains by applying the optimal policy:

	$\bar{z} < x_{iB}$	$x_{iB} \leq \bar{z} \leq x_i^R$	$x_i^R < \bar{z}$
$x_i^R < x_i$	μ_i	$x_i - k_i$	\bar{z}
$x_{iB} \leq x_i \leq x_i^R$	μ_i	$\max\{x_i, \bar{z}\} - k_i$	\bar{z}
$x_i < x_{iB}$	μ_i	$\max\{x_i, \bar{z}\} - k_i$	\bar{z}

Table 2: Payoff obtained by applying Lemma 0

Table 3 below is obtained from Table 2 by applying equation (1) to the middle column, which allows us, for each value of \bar{z} , to substitute k_i with $x_i - x_i^R$ when $x_i > x_i^R$ and with 0 otherwise. The subsequent table, Table 4, is obtained from Table 3 by rewriting equation (2) in the following way:

$$\mu_i = -k_i + \int_{x_{iB}}^{+\infty} x_i dF_i(x_i) + \int_{-\infty}^{x_{iB}} x_{iB} dF_i(x_i) \quad (3)$$

Equation (3) then allows us to replace μ_i in the first column by $x_i - k_i$ when $x_i \geq x_{iB}$ and by $x_{iB} - k_i$ otherwise. Finally, we also replace k_i as we did in creating Table 3.

	$\bar{z} < x_{iB}$	$x_{iB} \leq \bar{z} \leq x_i^R$	$x_i^R < \bar{z}$
$x_i^R < x_i$	μ_i	x_i^R	\bar{z}
$x_{iB} \leq x_i \leq x_i^R$	μ_i	$\max\{x_i, \bar{z}\}$	\bar{z}
$x_i < x_{iB}$	μ_i	$\max\{x_i, \bar{z}\}$	\bar{z}

Table 3: Table 2 replacing k_i in column 2

	$z < x_{iB}$	$x_{iB} \leq \bar{z} \leq x_i^R$	$x_i^R < \bar{z}$
$x_i^R < x_i$	x_i^R	x_i^R	\bar{z}
$x_{iB} \leq x_i \leq x_i^R$	x_i	$\max\{x_i, \bar{z}\}$	\bar{z}
$x_i < x_{iB}$	x_{iB}	$\max\{x_i, \bar{z}\}$	\bar{z}

Table 4: Table 3 replacing μ_i and, then, k_i in column 1

We remark on two features of Table 4. First, the reservation value represents the highest prize the agent expects to get from inspecting box i , after internalizing inspection costs. Second, if the agent didn't have the option of taking box i without inspection, the first column of Table 2 would have $\max\{x_i, \bar{z}\} - k_i$ in each row since $\bar{z} < x_i^R$. Hence, by similar steps, the first and second columns of Table 4 would be equal; in particular, when \bar{z} and x_i are both less than x_{iB} , the payoff would be $\max\{x_i, \bar{z}\} < x_{iB}$ (after internalizing inspection costs). However, when the agent has the option of taking box i without inspection, the payoff he obtains when \bar{z} and x_i are both less than x_{iB} is x_{iB} , after internalizing that he did not pay box i 's inspection cost. This is the reason why we refer to x_{iB} as box i 's *backup value*: when the agent takes box i without inspection, it is *as if* his payoff is bounded

below by x_{iB} .

The next subsection provides a different interpretation of the cutoffs. It can be skipped without loss of continuity, but is useful for intuition.

3.3 A different interpretation of x^R and x_B

We conclude Section 3 by providing a different interpretation for the cutoffs. Consider again the case in which $N = 1$. When $\bar{z} > \mu$, if the agent stops, he chooses to take z . However, if he inspects the box, he may discover that it contains a prize better than \bar{z} , increasing his *ex-post* payoff by $x - \bar{z}$. Thus, by inspecting the box, the agent avoids rejecting a box that contains a better prize than the outside option (type I error). *Ex-ante*, the value of inspecting the box is then given by: $V_I(\bar{z}) = \int_{\bar{z}}^{+\infty} (x - \bar{z})dF(x)$. When $\bar{z} \leq \mu$, if the agent stops, he chooses to take the uninspected box. However, if he inspects the box, he may discover that the box contains something worse than \bar{z} , which yields an *ex-post* loss of $\bar{z} - x$. By inspecting the box, the agent avoids taking boxes that are worse than what he has (type II error). *Ex-ante*, this is worth $V_{II}(\bar{z}) = \int_{-\infty}^{\bar{z}} (\bar{z} - x)dF(x)$ to the agent. Thus, given the outside option, the value of the new information for the agent is given by:

$$V(\bar{z}) = \begin{cases} V_{II}(\bar{z}) & \text{if } \bar{z} \leq \mu \\ V_I(\bar{z}) & \text{if } \bar{z} > \mu \end{cases}$$

Hence, the decision whether to acquire information or not is determined by whether $V(\bar{z}) \geq (\leq)k$, as illustrated by the following figure:




Figure 2: The value of information

We have that $V_{II}(\mu) = V_I(\mu)$ since $\mu : \int_{-\infty}^{\mu} (\mu - x) dF(x) = \int_{\mu}^{+\infty} (x - \mu) dF(x)$.

$V_I(\bar{z})$ represents the value of information in Weitzman's problem. The higher the outside option, the lower the value of finding out that the uninspected box contains a better option. $V_{II}(\bar{z})$ represents the value of information when the agent is considering taking the box without inspecting it first: it decreases with the difference between the outside option and μ .

The existence of two cutoffs introduces two different priorities for inspecting boxes. The first is given by the ordering of the reservation values: when the agent is considering his choice between stopping and taking the highest sampled prize, or inspecting one more box, it should be the box with the highest x^R . The second is given by the ordering of the backup values: when deciding which box to take without inspection, in case he finishes search, the agent prefers boxes with higher backup values (see Lemma 2). Thus, he prefers inspecting boxes with higher backup values last. Therefore, whenever he decides to search, the agent must decide which box he inspects next, and which he leaves to take without inspection, knowing that the first, once inspected, can never be taken without inspection.

4 Preliminary Lemmas

Lemmas 1 and 2 below formalize the idea that the backup value of box i represents the value of the option of taking box i without inspection. To see this, recall that \mathcal{U} is the set of uninspected boxes, and that $\mu_i = \mathbb{E}_{F_i} x$, for each $i \in \mathcal{U}$. If, for all $i \in \mathcal{U}$, the maximum sampled prize, \bar{z} , is greater than μ_i , from then onwards, the optimal sampling policy is given by applying Weitzman's rule to the boxes in \mathcal{U} . Lemma 1 shows that this is not necessary for Weitzman's rule to be optimal. Indeed, it states that whenever the maximum sampled prize exceeds the highest backup value amongst uninspected boxes, the option of taking a box without inspection has no value to the agent. Hence, Weitzman's rule is optimal from that moment on. Lemma 2 shows that, given (\mathcal{U}, z) , if the optimal policy implies stopping and taking a box without inspection, then the chosen box, which is the box with the highest μ_i in the set \mathcal{U} , is also the box with the highest backup value. Moreover, this is also the box with the highest reservation value. This last part follows from the following. Lemma 1 implies that, since it is optimal to stop and take a box without inspection, it has to be that $\bar{z} < \max_{i \in \mathcal{U}} x_{iB}$, and assumption 1 implies that $\max_{i \in \mathcal{U}} x_{iB} < \max_{i \in \mathcal{U}} x_i^R$. Therefore, if the box with the highest reservation value is not the box with the highest backup value, the agent always prefers inspecting the first box, and choosing whatever is best between keeping the prize and taking the second box without inspection to stopping and taking the second box without inspection.

Lemma 1. Let (\mathcal{U}, z) denote the set of boxes, and the vector of realized prizes, respectively. If $(\forall i \in \mathcal{U}) : \bar{z} \geq x_{iB}$, then Weitzman's optimal sampling policy is optimal in all continuation histories.

Lemma 2. Let (\mathcal{U}, z) denote the set of boxes, and the vector of realized prizes. Assume $\bar{z} < \max_{i \in \mathcal{U}} x_{iB}$. If $\varphi^*(\mathcal{U}, z) = 1$, then it has to be the case that:

$$\arg \max_{i \in \mathcal{U}} x_{iB} = \arg \max_{i \in \mathcal{U}} x_i^R \subseteq \arg \max_{i \in \mathcal{U}} \mu_i$$

Moreover, the sets $\arg \max_{i \in \mathcal{U}} x_{iB}$ and $\arg \max_{i \in \mathcal{U}} x_i^R$ are singletons.

Observe that Lemma 1 also shows that the differences between our search problem and Weitzman's arise when the maximum sampled prize falls below the highest backup value.

Our next result, Lemma 3, shows that there are two reasons why, given a set of boxes \mathcal{U} and a maximum sampled prize \bar{z} , the agent may deviate from Weitzman's order when selecting which box to inspect next. Let l be the box with the maximum reservation value, and let $j \neq l$ be the box which is inspected now according to the optimal policy. Then, he expects that after inspecting j , he might either (i) take box l without inspection in $\mathcal{U} \setminus \{j\}$, or (ii) continue search in $\mathcal{U} \setminus \{j\}$ whenever \bar{z} and x_j are below x_l^R , but deviate yet again from Weitzman's order.

To understand (i), consider the school choice example in the introduction. There, the agent inspects school B (box j) first, which is the one with the lowest reservation value. Notice that, if after inspecting school B , the agent observes $x_B = 0$, then he accepts school A (box l) without inspection. That is, the agent deviates from Weitzman's order since he assigns positive probability to accepting school A without inspection: had he visited school A first, he would have lost the option to do so.

To understand (ii), consider the following example. Let $\mathcal{U} = \{l, j, i\}$, let $l = \arg \max_{k \in \mathcal{U}} x_k^R$, and let j be the box which is inspected next according to the optimal strategy. Suppose further that the maximum sampled prize, $\bar{z} < x_{kB}$, for some $k \in \mathcal{U}$.⁴ Lemma 3 states that there exists a sufficiently low value of $\max\{x_j, \bar{z}\}$ such that the agent finds it optimal, conditional on continuing search, to violate Weitzman's order once more, and open box i .⁵

Lemma 3. Let (\mathcal{U}, z) denote the set of boxes, and the vector of realized prizes, respectively. Assume that $\sigma^*(\mathcal{U}, z) = j$, where $x_j^R < \max_{i \in \mathcal{U}} x_i^R \equiv x_l^R$. Then, it cannot be the case that $\sigma^*(\mathcal{U} \setminus \{j\}, z \cup \{x_j\}) = l$, and $\varphi^*(\mathcal{U} \setminus \{j\}, z \cup \{x_j\}) = 0$, whenever $\max\{x_j, \bar{z}\} \leq x_l^R$.

⁴Otherwise, by Lemma 1, Weitzman's policy is optimal.

⁵Moreover, by (i), it has to be that, conditional on deviating and inspecting box i , the agent expects he might take box l without inspection in the continuation.

5 Optimal Policy: Order and Stopping

Applying Lemma 0 and the results in Section 4 inductively to the case of N boxes is hard since, as the example in the introduction illustrates, the order in which boxes are inspected may be history dependent. Nevertheless, we identify conditions under which the optimal order policy coincides with Weitzman's [15] after every history. However, the optimal stopping rule is different, and we characterize it.

Our first result, Proposition 1, requires that, given any two boxes i, j , $x_j^R \leq x_i^R$ if, and only if, $x_{iB} \leq x_{jB}$. Since the box with the highest x^R is the box with the lowest x_B , the agent never foregoes taking without inspection a good backup. This implies that Weitzman's order is optimal. Moreover, by Lemma 2, we know that stopping and taking a box without inspection is never optimal when there is more than one box to be inspected: the box with the highest backup value is never the box with the highest reservation value. Hence, Weitzman's stopping rule applies to all, but the last box. Proposition 1 states the result formally, and Corollary 1 provides conditions under which the conditions in Proposition 1 are satisfied.⁶

Proposition 1. *Fix a set $\mathcal{N} = \{1, \dots, N\}$ of boxes. Assume that boxes can be labelled so that $[x_{iB}, x_i^R]$ forms a monotone sequence in the set inclusion order. The following is the optimal policy:*

Order *If a box is to be inspected next, it should be the box with the highest reservation value.*

Stopping

1. *If there is more than one box remaining, stop only if the maximum sampled prize is higher than the highest reservation value amongst uninspected boxes, and take the maximum sampled prize.*
2. *If only one box remains, stop if the maximum sampled prize is less than x_B or is higher than x^R . In the first case, take the remaining box without inspection; otherwise, take the maximum sampled prize.*

⁶An analogue of Proposition 1 is discussed in Postl [13] for a two-boxes-equal-inspection-costs setup. We show that the restriction to two boxes or equal inspection costs is not necessary, and provide meaningful properties of the distributions under which Proposition 1 holds.

Corollary 1 shows conditions on the primitives such that the ordering of the cutoffs is the one in Proposition 1.

Corollary 1. *Assume $\{F_i\}_{i \in \mathcal{N}}$ is such that if $i < i'$, then $F_{i'}$ is a mean-preserving spread of F_i . Moreover, assume $\forall i \in \mathcal{N} \quad k_i = k$. Then, the optimal policy is given by Proposition 1.*

Corollary 1 has an easy interpretation. On the one hand, boxes with higher dispersion are better for inspection since the agent can get better draws; on the other, these are the boxes that are not good backups: they can also contain worse draws. As discussed in Section 1.1, the same assumptions as in Corollary 1 are used in Vishwanath's [14] to obtain the reservation value rule in her parallel search model, and in the working paper version of Chade and Smith [4] to extend their binary-prize simultaneous search model to one with a continuum of possible prizes.

Remark 1. It is worth noting that something weaker than mean-preserving spreads is needed for Proposition 1 to hold when all boxes share the same inspection cost. Indeed, it suffices that if $i < i'$, then, for all convex functions with non-negative range $\phi : \mathbb{R} \mapsto \mathbb{R}_+$, $\int \phi(x) dF_i(x) \leq \int \phi(x) dF_{i'}(x)$.⁷

The two-box case serves to illustrate the reasoning behind Proposition 1, and serves as a proof sketch for the N -box case.⁸ To see this, assume that $N = 2$, and denote the two boxes $i = 1, 2$. Consider the case in which $x_2^R \leq x_1^R$. Weitzman's rule states the agent should inspect box 1 first. We want to see when we can have the same property. Once the agent inspects box i , the optimal policy is to take x_i whenever $x_i \geq x_j^R$, inspect box $j \neq i$ if $x_{jB} \leq x_i \leq x_j^R$, and take box j without inspection otherwise (see Lemma 0). Denote by Π_{ij} this payoff. The difference

⁷Mean preserving spreads, or the convex-order as it is defined in Ganuza and Penalva [6], and Li and Shi [9], requires the condition to hold for all convex functions.

⁸A similar calculation appears in Postl [13].

$\Pi_{12} - \Pi_{21}$ determines which box should be inspected first, and is given by:

$$\begin{aligned} \Pi_{12} - \Pi_{21} = & \underbrace{\int_{x_2^R}^{+\infty} \int_{x_2^R}^{+\infty} (\min\{x_1, x_2, x_1^R\} - x_2^R) dF_1 dF_2}_{\Pi_{12}^W - \Pi_{21}^W} \\ & + \underbrace{\int_{-\infty}^{\max\{x_{1B}, 0\}} \int_{-\infty}^{\max\{x_{1B}, 0\}} (\max\{x_1, x_2, \max\{x_{2B}, 0\}\} - \max\{x_{1B}, 0\}) dF_1 dF_2}_{(\Pi_{12} - \Pi_{12}^W) - (\Pi_{21} - \Pi_{21}^W)} \end{aligned} \quad (4)$$

where the terms $\max\{x_{jB}, 0\}$, $j = 1, 2$ cover the cases in which the optimal policy after inspecting box $i \neq j$ involves not taking box j without inspection.

The two terms in (4) correspond to the two trade-offs the agent has to balance when deciding which box to inspect first. The first term is identical to what appears in Weitzman's [15] model, and it favors inspecting box 1 first since it is the box with the highest reservation value. It is the difference between inspecting box 1 first and following Weitzman's policy (denote this by Π_{12}^W), and inspecting box 2 first and following Weitzman's policy (denote this by Π_{21}^W).

The second term corresponds to the trade-off between taking box 1 or box 2 without inspection, and it favors the box with the highest backup value. It is the difference between two terms. The first one is $\Pi_{12} - \Pi_{12}^W$: it measures the improvement in the agent's payoff from having the option to take box 2 without inspection. The second one is $\Pi_{21} - \Pi_{21}^W$. The difference between these two terms yields the second term in (4). It shows that, if $x_{1B} < x_{2B}$, then adding the option of taking a box without inspection improves more the value of starting search with box 1 than the value of starting search with box 2 since, by starting search with box 1, the agent retains the option of taking the better backup box (box 2) without inspection.

Equation (4) makes clear that if box 1 is more efficient for search (i.e. $x_2^R < x_1^R$), and box 2 is a better backup than box 1 (i.e. $x_{1B} < x_{2B}$), the difference is unambiguously non-negative. Proposition 1 shows this reasoning can be extended to an arbitrary number of boxes N .

Proposition 1 is not enough to characterize the optimal sampling policy in every environment: we need to consider the case in which $x_2^R < x_1^R$ and $x_{2B} < x_{1B}$.

When $x_{2B} < x_{1B}$, inspecting box 1 first implies the agent has to forego his best backup box. There are two cases of interest in which, despite this trade-off being present, the optimal policy still involves inspecting boxes according to Weitzman's order. The first case considers boxes with only two prizes (Proposition 2). The second case is when the agent possesses 2 boxes, both boxes share the same inspection cost, and prizes normalized by their mean are distributed according to a symmetric distribution (Proposition 3). Similar conditions have been used before in search models: Chade and Smith [4] use binary prizes in their simultaneous search model, while Klabjan, Olszewski and Wolinsky [7] consider two boxes with symmetric distributions. Under these conditions, given any two boxes i, j it holds that $x_j^R \leq x_i^R$ if, and only if, $x_{jB} \leq x_{iB}$. The propositions show that the trade-off between inspecting the box with the highest reservation value and taking it without inspection is resolved either by continuing search with this box, or stopping search and taking this box without inspection; the agent never finds it optimal to search boxes in a different order.

Proposition 2. *Fix a set $\mathcal{N} = \{1, \dots, N\}$ of boxes. Assume that boxes have binary prizes, i.e. $x_i \in \{y, x\}$, $0 < y < x$, $P(x_i = x) = p_i$, and all boxes have the same inspection cost. The following is the optimal policy:*

Order *If a box is to be inspected next, it should be the box with the highest reservation value.*

Stopping *Search stops when: (i) the maximum sampled prize is x , or (ii) the maximum sampled prize is below the highest backup value, and the continuation value of inspecting one more box according to the optimal policy is below the highest backup value amongst uninspected boxes. If (i), take the maximum sampled prize; if (ii), take the box with the highest backup value without inspection. Moreover, whenever the maximum sampled prize is below the highest backup value, if the second highest reservation value is smaller than the highest backup value amongst uninspected boxes, (ii) holds.*

Contrary to Weitzman's model, in which, given \bar{z} , the highest reservation value is the only number needed to determine whether to stop search or not, the rule in Proposition 2 requires calculating the value of continuing search to determine

whether to stop or not. The last statement in Proposition 2 provides a sufficient condition under which looking at the reservation and backup values alone suffices to know when to stop search.

Proposition 3. *Let $\mathcal{N} = \{1, 2\}$, and assume prizes $x_i \sim F(\cdot - \mu_i)$, where F has pdf f , which is symmetric around 0. Moreover, assume that $k_i = k > 0$ for all $i \in \mathcal{N}$. Then, if $x_2^R \leq x_1^R$, the following is the optimal policy:*

Order *Box 1 is inspected first.*

Stopping *If $x_{2B} \geq 0$, search stops when x_1 is either (i) less than x_{2B} , or (ii) higher than x_2^R . If (i), take box 2 without inspection; if (ii) take x_1 . If $x_{2B} < 0$, search stops only if (ii). Search starts if, and only if, $\mu_1 < \Pi_{12}$.*

The intuition for why Proposition 3 holds lies in equations (1),(2) and (4). Equation (1) shows that the reservation value depends only on the upper tail of the probability distribution over prizes, while equation (2) shows that the backup value depends only on the lower tail of the distribution. A distribution with a fat upper tail has a high reservation value, and a distribution with a fat lower tail has a low backup value. Equation (4) then shows that the choice to start search with box 1 or 2 depends on the weight each distribution puts on the upper and lower tails. When distributions coincide up to the mean, and are symmetric, both tails are treated “equally” by both distributions, and when $x_{2B} \geq 0$, both terms exactly offset each other. When $x_{2B} < 0$, and, hence, by Lemma 1, Weitzman’s rule applies to box 2, after inspecting box 1, the first term in (4) dominates the second.

Propositions 2 and 3 suggest that whenever $\mathcal{N} = \{1, 2\}$ such that F_1 first-order stochastically dominates F_2 , and $k_1 = k_2$, then box 1 should be inspected first, if any. The next example shows that this is not true:

Example. Suppose $\mathcal{N} = \{1, 2\}$, and $X_1 = X_2 = \{0, 2, 10\}$. Suppose $P(X_1 = 2) = P(X_2 = 2) = 0.2$, and $P(X_1 = 10) = 0.7, P(X_2 = 10) = 0.5$, so that $F_1 >_{FOSD} F_2$. Assume that $k_1 = k_2 = 1$. It can be shown that $x_{1B} = \frac{14}{3} > x_{2B} = 2.8$, and $x_1^R = \frac{60}{7} > x_2^R = 8$. Notice that after inspecting box i , search always stops: the agent takes the inspected box when $x_i = 10$, and takes box j without inspection whenever $x_i \leq 2$. Since $\mu_1 < x_2^R$, one can show that inspecting box 2 first dominates taking box 1 without inspection; moreover, inspecting box 2 first dominates inspecting

box 1 first since: $8.62 = 0.7 \times 10 + 0.3 \times \mu_2 < 0.5 \times 10 + 0.5 \times \mu_1 = 8.7$.

Contrast this with Weitzman's model where if F_i first-order stochastically dominates $F_{i'}$, and $k_i = k_{i'}$, then box i is inspected first. This is because it is more probable to obtain a higher prize under F_i than under $F_{i'}$, and the agent has to inspect boxes in order to obtain prizes in Weitzman's model. However, F_i also has a higher backup value than $F_{i'}$: box i has lower probability of yielding a low prize when taken without inspection. Therefore, a first-order stochastic dominance shift makes box i both more attractive to search and to take without inspection. The example shows that the trade-off is not always resolved in favor of either inspecting first box i , or taking box i without inspection and not searching at all, as is the case in Propositions 2 and 3. It also shows why Corollary 1 cannot be relaxed to second-order stochastic dominance.

6 Conclusions

We consider a relaxed version of Weitzman's search problem; namely, we allow the agent to take any uninspected box without inspecting its contents first upon stopping. We show that, under conditions common in the search and information acquisition literature, the optimal policy involves following Weitzman's inspection order, and characterize the optimal stopping rule in those cases. Moreover, we provide properties of the optimal policy that must hold across all environments.

A Appendix

We denote by $|\cdot|$ the cardinality of a set in what follows.

A.1 Proof of Propositions 1 and 2

Proposition1. *Fix a set \mathcal{U} of boxes, and let z be the vector of previously realized prizes. Assume that boxes are labelled so that $[x_{iB}, x_i^R]$ forms a monotone sequence*

in the set inclusion order, that is $i < i'$ implies $[x_{iB}, x_i^R] \subset [x_{i'B}, x_{i'}^R]$. The following is the optimal policy:

Order $\sigma^*(\mathcal{U}, z) = \arg \max\{i | i \in \mathcal{U}\}$

Stopping If $|\mathcal{U}| > 1$, $\varphi^*(\mathcal{U}, z) = 1$ if, and only if, $\bar{z} > \arg \max_{i \in \mathcal{U}} x_i^R$. If $|\mathcal{U}| = 1$, $\varphi^*(\mathcal{U}, z) = 1$ if, and only if, (i) $\bar{z} > \arg \max_{i \in \mathcal{U}} x_i^R$, or (ii) $\bar{z} < \arg \max_{i \in \mathcal{U}} x_{iB}$

Proof. The argument is the same as the one in Weitzman's proof of Pandora's rule. We include it for completeness and since it is used in the proof of Lemma 1. We proceed by induction on $U = |\mathcal{U}|$. Let $P(U)$ denote the following predicate:

P(U): $(\forall z)(\forall \mathcal{U}) : |\mathcal{U}| = U$, and \mathcal{U} is enumerated as in Proposition 1, the order and stopping rules in Proposition 1 are optimal.

Lemma 0 shows that $P(1) = 1$. We show that Proposition 1 is valid for $U = 2$, and, then, prove the inductive step.

Step 1: $P(2) = 1$

Recall 2 is the box with the highest label in \mathcal{U} . We start by showing that the stopping rule is optimal. We do so by considering two cases:

$\bar{z} \geq x_2^R$ Note that if some box $i \in \{1, 2\}$ is inspected, then, by Lemma 0, since $\max\{\bar{z}, x_i\} > x_j^R, j \neq i$, stopping is optimal. Moreover, the payoff from inspecting i and stopping is less than \bar{z} since:

$$-k_i + \int \max\{x_i, \bar{z}\} dF_i(x_i) < \bar{z}$$

by equation (1). Therefore, when $x_2^R \leq \bar{z}$ it is optimal to stop search.

$\bar{z} < x_2^R$ If $\max\{\bar{z}, \max_i \mu_i\} \neq \mu_2$, then inspecting box 2 alone, and stopping dominates stopping and obtaining payoff $\max\{\bar{z}, \max_{i \in \mathcal{U}} \mu_i\}$ by (1). If $\max\{\bar{z}, \max_i \mu_i\} = \mu_2$, since $\max\{\bar{z}, \mu_1\} > x_{1B} \geq x_{2B}$, by equation (2) we have that inspecting box 2 and stopping dominates obtaining payoff μ_2 .

Finally, it remains to show that inspecting box 2 first is optimal whenever $\bar{z} < x_2^R$. This follows from equation (4), replacing 0 with \bar{z} .

Step 2: $P(U) = 1 \Rightarrow P(U + 1) = 1$

Assume $P(U)$ is true. Fix \mathcal{U} as in $P(U+1)$, and recall $U+1 = \arg \max\{i | i \in \mathcal{U}\}$. Note that, by assumption, $U+1$ is the box with the highest reservation value. Let $y = \max\{\bar{z}, \max_{i \in \mathcal{U}} \mu_i\}$ be the outside option. We show first that the stopping rule is optimal. We do so by considering two cases.

1. Assume $x_{U+1}^R < y$, and that box $U+1$ is inspected.⁹ Since $\mu_{U+1} < x_{U+1}^R$, and $x_{U+1}^R < y$ we have that $x_{U+1}^R < \max\{\bar{z}, \max_{i \in \mathcal{U} \setminus \{U+1\}} \mu_i\}$. Thus, $(\forall x_{U+1}) \max\{\bar{z}, x_{U+1}, \max_{i \in \mathcal{U} \setminus \{U+1\}} \mu_i\} > x_{U+1}^R$. Since $|\mathcal{U} \setminus \{U+1\}| = U$, it is optimal to stop by the inductive hypothesis. Therefore, the payoff from inspecting box $U+1$ is:

$$-k_{U+1} + \int \max\{\bar{z}, x_{U+1}\} dF_{U+1}(x_{U+1}) < \bar{z}$$

which follows from (1), since $\bar{z} > x_{U+1}$ (see footnote 9). Therefore, it is optimal to stop.

2. Assume that $y \leq x_{U+1}^R$. Since, by assumption 1, $x_{kB} < \mu_k$, and $(\forall k \in \mathcal{U} \setminus \{U+1\}) x_{U+1B} \leq x_{kB}$, it must be the case that $y > x_{U+1B}$, and $\max\{\bar{z}, \max_{i \in \mathcal{U} \setminus \{U+1\}} \mu_i\} > x_{U+1B}$. Then, inspecting box $U+1$ and stopping dominates the payoff the agent obtains by stopping. If $\mu_{U+1} = y$, this follows from $\max\{\bar{z}, \max_{i \in \mathcal{U} \setminus \{U+1\}} \mu_i\} > x_{U+1B}$, and equation(2); if $\mu_{U+1} < y$, this follows from 1. above.

Now we show that the order in Proposition 1 is optimal. Assume $y \leq x_{U+1}^R$ (otherwise, we just showed search stops). Let $j \in \mathcal{U}$ be a box such that $x_j^R < x_{U+1}^R$, and let $U = \arg \max_{i \in \mathcal{U} \setminus \{U+1\}} x_i^R$. Consider the following two policies:

P.J Open box j first. There are now U boxes left to be inspected, stop, or continue search according to the rule described in Proposition 1.

P.U+1 Open box $U+1$ first. If $x_U^* \leq x_{U+1}$, stop. Otherwise, open box j and, stop or continue search according to the rule described in Proposition 1.

Since the continuation policy is different if there are two or more boxes remaining after opening $U+1$, we prove that the ordering is optimal by considering the case in which $U+1 = 3$, and $U+1 > 3$. Moreover, when $U+1 = 3$, we focus on the

⁹Note that, since $U+1$ is the box with the highest reservation value, and, by Assumption 1, $(\forall k \in \mathcal{U}) \mu_k < x_k^* \leq x_{U+1}^R$, $y > x_{U+1}^R \Rightarrow \bar{z} > x_l^R$

case in which $x_{UB} > \bar{z}$ (if this is the case, when box h is the only one left to be inspected, it may be taken without inspection). When $x_{UB} \leq \bar{z}$, the proof is the same as in the case in which $U + 1 > 3$.

When $U + 1 = 3$, $x_{UB} \geq \bar{z}$, define the following quantities:

$\pi_{U+1} = P(x_{U+1} \geq x_{U+1}^R)$ $\pi_j = P(x_j \geq x_{U+1}^R)$ $\lambda_{U+1} = P(x_U^R \leq x_{U+1} \leq x_{U+1}^R)$ $\lambda_j = P(x_U^R \leq x_j \leq x_{U+1}^R)$ $\nu_{U+1} = P(x_{UB} \leq x_{U+1} \leq x_U^R)$ $\nu_j = P(x_{UB} \leq x_j \leq x_U^R)$	$w_{U+1} = E[x_{U+1} x_{U+1} \geq x_{U+1}^R]$ $w_j = E[x_j x_j \geq x_{U+1}^R]$ $\tilde{v}_j = E[\max\{x_j, \mu_U, \bar{z}\} x_j \in [x_U^R, x_{U+1}^R]]$ $\tilde{v}_{U+1} = E[\max\{x_{U+1}, \mu_U, \mu_j, \bar{z}\} x_{U+1} \in [x_U^R, x_{U+1}^R]]$ $\bar{v}_{U+1} = E[\max\{x_{U+1}, \mu_U, \bar{z}\} x_{U+1} \in [x_U^R, x_{U+1}^R]]$ $\underline{d} = E[\max\{x_{U+1}, x_j, \mu_U, \bar{z}\} x_{U+1}, x_j \in [x_U^R, x_{U+1}^R]]$
$\bar{d} = E[\max\{x_{U+1}, x_j, \mu_U, \bar{z}\} x_{U+1} \in [x_{UB}, x_U^R], x_j \in [x_U^R, x_{U+1}^R]]$ $\Phi_O = -k_U + E[\max\{x_{U+1}, x_U, x_j, \bar{z}\} x_{U+1}, x_j \in [x_{UB}, x_U^R]]$ $\Phi_O^m = -k_U + E[\max\{x_{U+1}, x_U, x_j, \bar{z}\} x_{-m} \leq x_{UB}, x_m \in [x_{UB}, x_U^R]]$ $\Phi_u^m = -k_U + E[\max\{x_{U+1}, x_U, x_j, \bar{z}\} x_m \leq x_{UB}, x_{-m} \in [x_{UB}, x_U^R]]$	

where $m, -m \in \{U+1, j\}$. Then, the payoffs from P.J and P.U+1 maybe written as:

$$\begin{aligned}
P.J &= -k_j + \pi_j w_j + \lambda_j [-k_{U+1} + \pi_{U+1} w_{U+1} + \lambda_{U+1} \underline{d} + (1 - \lambda_{U+1} - \pi_{U+1}) \tilde{v}_j] \\
&+ \nu_j [-k_{U+1} + \pi_{U+1} w_{U+1} + \lambda_{U+1} \bar{v}_{U+1} + \nu_{U+1} \Phi_O + (1 - \pi_{U+1} - \lambda_{U+1} - \nu_{U+1}) \Phi_O^j] \\
&+ (1 - \pi_j - \lambda_j - \nu_j) \left[\begin{aligned} &-k_{U+1} + \pi_{U+1} w_{U+1} + \lambda_{U+1} \bar{v}_{U+1} + \nu_{U+1} \Phi_u^j \\ &+ (1 - \pi_{U+1} - \lambda_{U+1} - \nu_{U+1}) \mu_U \end{aligned} \right] \\
P.U + 1 &= -k_{U+1} + \pi_{U+1} w_{U+1} + \lambda_{U+1} \tilde{v}_{U+1} \\
&+ \nu_{U+1} [-k_j + \pi_j w_j + \lambda_j \tilde{v}_j + \nu_j \Phi_O + (1 - \pi_j - \lambda_j - \nu_j) \Phi_O^{U+1}] \\
&+ (1 - \pi_{U+1} - \lambda_{U+1} - \nu_{U+1}) [-k_j + \pi_j w_j + \lambda_j \tilde{v}_j + \nu_j \Phi_u^{U+1} + (1 - \pi_j - \lambda_j - \nu_j) \mu_U]
\end{aligned}$$

The difference $P.U + 1 - P.J$, after canceling terms, is:

$$\begin{aligned}
P.U + 1 - P.J &= -k_{U+1} \pi_j + \pi_{U+1} \pi_j w_{U+1} + \lambda_{U+1} (\lambda_j + \pi_j) \bar{v}_{U+1} + k_j (\pi_{U+1} + \lambda_{U+1}) \\
&- \pi_j w_j (\lambda_{U+1} + \pi_{U+1}) - \lambda_j \lambda_{U+1} \underline{d}
\end{aligned} \tag{5}$$

Use the fact that $k_{U+1} = \pi_{U+1} [w_{U+1} - x_{U+1}^R]$, $k_j = \pi_j [w_j - x_j^R] + \lambda_j [v_j - x_j^R] +$

$\tau_j[u_j - x_j^R]$, where $u_j = E[x_j | x_j \in [x_j^R, x_U^R]]$, $\tau_j = P(x_j \in [x_j^R, x_U^R])$, $v_j = E[x_j | x_j \in [x_U^R, x_{U+1}^R]]$ to rewrite equation (5) as:

$$\begin{aligned} P.U + 1 - P.J &= \pi_j \pi_{U+1} (x_{U+1}^R - x_U^R) + \pi_j \lambda_{U+1} (\bar{v}_{U+1} - x_j^R) + \lambda_j (\lambda_{U+1} + \pi_{U+1}) (v_j - x_j^R) \\ &+ \tau_j (u_j - x_j^R) + \lambda_{U+1} \lambda_j \bar{v}_{U+1} - \lambda_{U+1} \lambda_j \underline{d} \end{aligned}$$

Note that \underline{d} can be written as:

$$\begin{aligned} \underline{d} &= E[\max\{\max\{x_{U+1}, \max\{\mu_U, \bar{z}\}\}, x_j\} | x_{U+1}, x_j \in [x_U^R, x_{U+1}^R]] \\ &= x_U^R + E[\max\{\max\{x_{U+1}, \max\{\mu_U, \bar{z}\}\} - x_U^R, x_j - x_U^R\} | x_{U+1}, x_j \in [x_U^R, x_{U+1}^R]] \\ &\leq x_U^R + E[\max\{x_{U+1}, \max\{\mu_U, \bar{z}\}\} - x_U^R + x_j - x_U^R | x_{U+1}, x_j \in [x_U^R, x_{U+1}^R]] \\ &= x_U^R + \bar{v}_{U+1} - x_U^R + v_j - x_U^R \\ &\leq \bar{v}_{U+1} - x_j^R + v_j \end{aligned}$$

Therefore,

$$\begin{aligned} P.U + 1 - P.J &\geq \pi_j \pi_{U+1} (x_{U+1}^R - x_j^R) + \pi_j \lambda_{U+1} (\bar{v}_{U+1} - x_j^R) \\ &+ \lambda_j \pi_{U+1} (v_j - x_j^R) + \tau_j (\lambda_{U+1} + \pi_{U+1}) (u_j - x_j^R) > 0 \end{aligned}$$

This shows, when $U + 1 = 3$, and $x_{UB} > \bar{z}$, that it is optimal to start search with box $U + 1$.

When either $U + 1 > 3$ or $x_{UB} < 0$, define $\Phi = E[V^*(\mathcal{U} \setminus \{U + 1, j\}, z \cup \{x_{U+1}\} \cup \{x_j\})]$, to be the expected (continuation) payoff the agent obtains by applying the rule in Proposition 1 when the set of boxes is $\mathcal{U} \setminus \{U + 1, j\}$ and the vector of realized prizes is $z \cup \{x_{U+1}\} \cup \{x_j\}$. Note that since $|\mathcal{U} \setminus \{U + 1, j\}| < U + 1$, and the boxes in $|\mathcal{U} \setminus \{U + 1, j\}| < U + 1$ can be enumerated as assumed in Proposition 1, the policy in Proposition 1 is optimal when applied to that set. Consider, again,

the payoffs obtained from following policies $P.U + 1$ and $P.J$:

$$\begin{aligned}
P.J &= -k_j + \pi_j w_j + \lambda_j [-k_{U+1} + \pi_{U+1} w_{U+1} + \lambda_{U+1} \underline{d} + (1 - \lambda_{U+1} - \pi_{U+1}) \tilde{v}_j] \\
&+ (1 - \pi_j - \lambda_j) [-k_{U+1} + \pi_{U+1} w_{U+1} + \lambda_{U+1} \bar{v}_{U+1} + (1 - \pi_{U+1} - \lambda_{U+1}) \Phi] \\
P.U + 1 &= -k_{U+1} + \pi_{U+1} w_{U+1} + \lambda_{U+1} \tilde{v}_{U+1} \\
&+ (1 - \pi_{U+1} - \lambda_{U+1}) [-k_j + \pi_j w_j + \lambda_j \tilde{v}_j + (1 - \pi_j - \lambda_j) \Phi]
\end{aligned}$$

Taking the difference $P.U + 1 - P.J$ yields the same expression as in (5), which shows that inspecting box $U + 1$ first is optimal. This completes our proof. \square

Corollary 1. *Assume $\{F_i\}_{i \in \mathcal{N}}$ is such that if $i < i'$, then $F_{i'}$ is a mean-preserving spread of F_i . Moreover, assume $\forall i \in \mathcal{N} \quad k_i = k$. Then, the optimal policy is given by Proposition 1.*

Proof. It suffices to show that if $i < i'$, then $[x_{iB}, x_i^R] \subseteq [x_{i'B}, x_{i'}^R]$. To see this, rewrite equation (1) for box i as:

$$k = \int_{x_i^R}^{+\infty} (x - x_i^R) dF_i(x) = \int_{-\infty}^{+\infty} \max\{x - x_i^R, 0\} dF_i(x)$$

and, note that, $F_{i'}$ is a mean-preserving spread of F_i , then we have that:

$$k = \int_{-\infty}^{+\infty} \max\{x - x_i^R, 0\} dF_i(x) \leq \int_{-\infty}^{+\infty} \max\{x - x_i^R, 0\} dF_{i'}(x)$$

Since $\int_{x_i^R}^{+\infty} (x - x_i^R) dF(x)$ is decreasing in x_i^R , we conclude that $x_i^R \leq x_{i'}^R$. Likewise, we may rewrite equation (2) as:

$$k = \int_{-\infty}^{x_{iB}} (x_{iB} - x) dF_i(x) = \int_{-\infty}^{+\infty} \max\{x_{iB} - x, 0\} dF_i(x)$$

Using the mean-preserving spread assumption again, we obtain that $i < i'$ implies that:

$$k = \int_{-\infty}^{+\infty} \max\{x_{iB} - x, 0\} dF_i(x) \leq \int_{-\infty}^{+\infty} \max\{x_{iB} - x, 0\} dF_{i'}(x)$$

Since $\int_{-\infty}^{x_{iB}} (x_{iB} - x) dF(x)$ is increasing in x_{iB} , we conclude that $x_{i'B} \leq x_{iB}$.

Therefore, we conclude that $[x_{iB}, x_i^R] \subseteq [x_{i'B}, x_{i'}^R]$. \square

Proposition2. Fix a set \mathcal{U} of boxes, and let $z \in \{x_0\} \times \{x, y\}^n, n \geq 0$ be the vector of previously realized prizes. Assume all boxes have $x_i \in \{y, x\}, y < x, 0 < x$, $P(x_i = x) = p_i$, and $k_i \equiv k$. The following is the optimal policy:

Order $\sigma^*(\mathcal{U}, z) \in \arg \max_{i \in \mathcal{U}} x_i^R$

Stopping $\varphi^*(\mathcal{U}, z) = 1$ if, and only, if $\bar{z} > x_{\sigma^*(\mathcal{U}, z)}^R$, or $\bar{z} = \max\{x_0, y\} < x_{\sigma^*(\mathcal{U}, z)B}$ and $V^*(\mathcal{U} \setminus \{\sigma^*(\mathcal{U}, z)\}, z \cup \{y\}) < x_{\sigma^*(\mathcal{U}, z)B}$

Define $l = \arg \max_{i \in \mathcal{U}} x_{iB}, h = \arg \max_{i \in \mathcal{U} \setminus \{l\}} x_i^R$. If $\max\{\bar{z}, x_h^R\} < x_{lB}$, then it is optimal to stop and take box l without inspection.

Proof. The proof is divided in two parts. First, we show that the order and stopping rule are optimal. Second, we show the sufficient condition for stopping. The proof of the first part is by induction on $U = |\mathcal{U}|$. Let $P(U)$ denote the following predicate:

P(U) $(\forall z)(\forall \mathcal{U}) : (|\mathcal{U}| = U)$, and \mathcal{U} satisfies the assumptions in Proposition 2, the order and stopping rules in Proposition 2 are optimal.

Step 1: $P(1) = 1$ The proof follows from Lemma 0.

Step 2: $P(U) = 1 \Rightarrow P(U + 1) = 1$

Let \mathcal{U} be such that $|\mathcal{U}| = U + 1$, and let $l = \arg \max_{i \in \mathcal{U}} x_i^R$. We first show that when $\bar{z} \geq x_l^R$ it is optimal to stop. Note that $\bar{z} > \mu_l = \max_{i \in \mathcal{U}} \mu_i$, hence, if the agent stops he selects \bar{z} as payoff. Moreover, if the agent inspects any box $k \in \mathcal{U}$, $|\mathcal{U} \setminus \{k\}| = U$, and then, by the inductive hypothesis, it is optimal to stop search. Moreover, by equation (1), it is not optimal to inspect a box in \mathcal{U} and stop.

Hence, assume that $\bar{z} < x_l^R$. We first show that, if a box is to be opened first, it has to be box l . Let j be any other box $j \neq l$. Let $V^*(\mathcal{U} \setminus \{l, j\}, z \cup \{x_l\} \cup \{x_j\})$ denote the value function in the continuation problem after inspecting boxes l, j .¹⁰

¹⁰That is, the payoff the agent obtains by following the policy stated in Proposition 2, which is optimal by the inductive step since $|\mathcal{U} \setminus \{l, j\}| < U + 1$

In a slight abuse of notation, define $V^*(\mathcal{U}, \max\{x_0, y\}) \equiv V^*(\mathcal{U}, z)$ whenever $\bar{z} < x$. If $V^*(\mathcal{U} \setminus \{l, j\}, \max\{x_0, y\}) \geq x_{lB}$ consider the following two policies:

P.L Open box l first. If $x_l = x$, stop. Otherwise, open box j and continue with the policy in the inductive hypothesis.

P.J Open box j first, and continue as indicated in the inductive hypothesis.

The payoff from applying **P.L** is:

$$p_l x + (1 - p_l)[p_j x + (1 - p_j)V^*(\mathcal{U} \setminus \{l, j\}, \max\{x_0, y\}) - k] - k$$

and the payoff from applying **P.J** is:

$$p_j x + (1 - p_j)[p_l x + (1 - p_l)V^*(\mathcal{U} \setminus \{l, j\}, \max\{x_0, y\}) - k] - k$$

The comparison of the payoffs yields the result, by noticing that it reduces to the case in which there are only two boxes l, j . When $V^*(\mathcal{U} \setminus \{l, j\}, \max\{x_0, y\}) < x_{lB}$, compare the following two policies:

P.L Open box l first. If $x_l = x$, stop. Otherwise, take box j without inspection.

P.J Open box j first, and continue as indicated in the inductive hypothesis.

The payoff from policy [P.L] is $p_l x + (1 - p_l)(p_j x + (1 - p_j)y) - k$, and that of policy [P.J] is $p_j x + (1 - p_j)(p_l x + (1 - p_l)y) - k$. The difference is null. Then, opening box l first dominates (weakly) opening any other box $j \neq l$. Now, to verify the rest of the stopping rule, note that if $\bar{z} = \max\{x_0, y\}$ and $\bar{z} < x_{lB}$,

$$\begin{aligned} p_l x + (1 - p_l)V^*(\mathcal{U} \setminus \{l\}, \max\{y, 0\}) - k &\geq p_l x + (1 - p_l)y \\ \Leftrightarrow V^*(\mathcal{U} \setminus \{l\}, \max\{y, 0\}) &\geq y + \frac{k}{1 - p_l} = x_{lB} \end{aligned}$$

Now we show that if $x_{lB} > \max\{\bar{z}, x_h^R\}$, $h = \arg \max_{i \in \mathcal{U} \setminus \{l\}} x_i^R$, it is optimal to stop. In order to do so, we prove the following claim:

Claim 1. Fix a set of boxes \mathcal{U} , and let z be the vector of previously realized prizes. Then, $V^*(\mathcal{U}, z) \leq \max\{\bar{z}, \max_{i \in \mathcal{U}} x_i^R\}$

Proof. The proof is by induction in $U = |\mathcal{U}|$. Let $P(U)$ denote the following predicate:

P(U) $(\forall z \in \{x_0\} \times \{x, y\}^n, n \geq 0)(\forall \mathcal{U}) : (|\mathcal{U}| = U)$, and \mathcal{U} satisfies the properties in Proposition 2, $V^*(\mathcal{U}, z) \leq \max\{\bar{z}, \max_{i \in \mathcal{U}} x_i^R\}$

Step 1: $P(1) = 1$ Given z and $\mathcal{U} = l$, by Lemma 0, we have that:

$$V^*(\mathcal{U}, z) = \max\{\bar{z}, \mu_l, -k + p_l x + (1 - p_l)\bar{z}\}$$

where we use the fact that if $y \geq \bar{z}$, then $\mu_l \geq \bar{z}$, and then it is not optimal to inspect the box to write the third term. Using that $p_l(x_l - x_l^R) = k$, we rewrite the above equation as:

$$\begin{aligned} V^*(\mathcal{U}, z) &= \max\{\bar{z}, \mu_l, -k + p_l x + (1 - p_l)\bar{z}\} \\ &= \max\{\bar{z}, \mu_l, p_l x_l^R + (1 - p_l)\bar{z}\} \leq \max\{\bar{z}, x_l^R\} \end{aligned}$$

which shows that $P(1) = 1$.

Step 2: $P(U) = 1 \Rightarrow P(U + 1) = 1$ Fix z and \mathcal{U} such that $|\mathcal{U}| = U + 1$. Consider first the case in which $\varphi^*(\mathcal{U}, z) = 1$. In that case, letting $l = \arg \max_{i \in \mathcal{U}} x_i^R$ we have:

$$V^*(\mathcal{U}, z) = \max\{\bar{z}, \mu_l\} \leq \max\{x_l^R, \bar{z}\}$$

Now, consider the case in which $\varphi^*(\mathcal{U}, z) = 0$; in particular, this means that $\bar{z} \neq x$. Then,

$$\begin{aligned} V^*(\mathcal{U}, z) &= -k + p_l x + (1 - p_l)V^*(\mathcal{U} \setminus \{l\}, z \cup \{y\}) \\ &= p_l x_l^R + (1 - p_l)V^*(\mathcal{U} \setminus \{l\}, z \cup \{y\}) \\ &\leq p_l x_l^R + (1 - p_l) \max\{\bar{z}, \max_{i \in \mathcal{U} \setminus \{l\}} x_i^R\} \\ &\leq \max\{\bar{z}, \max_{i \in \mathcal{U}} x_i^R\} \end{aligned}$$

where the second equality comes from $p_l(x - x_l^R) = k$, and the first inequality comes from applying the inductive hypothesis. This completes the proof.

□

To prove the sufficient condition, note that if, given \mathcal{U} and $\bar{z} = \max\{x_0, y\} < x_{lB}$, the agent inspects box l , then his expected payoff is bounded above by:

$$\begin{aligned} -k + p_l x + (1 - p_l) V^*(\mathcal{U} \setminus \{l\}, \max\{x_0, y\}) &\leq -k + p_l x + (1 - p_l) \max\{\bar{z}, x_h^R\} \\ &\leq p_l x + (1 - p_l) y \Leftrightarrow x_{lB} = y + \frac{k}{1 - p_l} \geq \max\{\bar{z}, x_h^R\} \end{aligned}$$

where the first inequality follows from the result in Claim 1. Therefore, it is optimal to stop and take box l without inspection. □

A.2 Proofs of Lemmas 1, 2, and 3

Lemma 1. Let \mathcal{U} be a set of boxes, and let z be the vector of realized prizes. If $(\forall i \in \mathcal{U}) : \bar{z} \geq x_{iB}$, Weitzman's sampling policy is optimal in all continuation histories.

Proof. The proof is by induction on $U = |\mathcal{U}|$. Let $P(U)$ denote the following predicate:

P(U) $(\forall \mathcal{U}) : (|\mathcal{U}| = U), (\forall z) : (\bar{z} \geq \max_{i \in \mathcal{U}} x_{iB})$, the order and stopping policy indicated in Lemma 1 is optimal.

Step 1: $P(1) = 1$ This follows from Lemma 0.

Step 2: $P(U) = 1 \Rightarrow P(U + 1) = 1$

Let $U + 1 = |\mathcal{U}|$, and let z be as in the statement of Lemma 1. Let $l \in \arg \max_{i \in \mathcal{U}} x_i^R$. First, we show that the stopping rule is optimal. We consider two cases:

$\bar{z} \geq x_l^R$ The argument in Step 2 of Proposition 1 implies that it is optimal to stop when $\bar{z} \geq x_l^R$.

$\bar{z} < x_l^R$ If $\max\{\bar{z}, \max_{i \in \mathcal{U}} \mu_i\} \neq \mu_l$, then, by equation (1), inspecting box l and stopping dominates stopping and obtaining payoff $\max\{\bar{z}, \max_{i \in \mathcal{U}} \mu_i\}$, since

$\max\{\bar{z}, \max_{i \in \mathcal{U}} \mu_i\} < x_l^R$. If $\max\{\bar{z}, \max_{i \in \mathcal{U}} \mu_i\} = \mu_l$, since $\bar{z} \geq x_{lB}$, we have that $\max\{\bar{z}, \max_{i \in \mathcal{U} \setminus \{l\}} \mu_i\} \geq x_{lB}$, and hence, by equation (2), inspecting box l and stopping dominates stopping and taking box l without inspection.

Finally, when $\bar{z} < x_l^R$, we need to show that inspecting box l first is optimal. This follows from the same argument as in Proposition 1. \square

Lemma 2. Let (\mathcal{U}, z) denote the set of boxes, and the vector of previously realized prizes. Assume $\bar{z} < \max_{i \in \mathcal{U}} x_{iB}$. If $\varphi^*(\mathcal{U}, z) = 1$, then it has to be the case that:

$$\arg \max_{i \in \mathcal{U}} x_{iB} = \arg \max_{i \in \mathcal{U}} x_i^R \subseteq \arg \max_{i \in \mathcal{U}} \mu_i$$

Moreover, the sets $\arg \max_{i \in \mathcal{U}} x_{iB}$ and $\arg \max_{i \in \mathcal{U}} x_i^R$ are singletons.

Proof. We will use l for boxes with the highest reservation value, k for boxes with the highest mean, and j for boxes with the highest backup value. We first show that $\arg \max_{i \in \mathcal{U}} x_{iB} \subseteq \arg \max_{i \in \mathcal{U}} \mu_i$. We do so by contradiction. Assume that $(\exists j, k \in \mathcal{U})(j \neq k) : \mu_j < \mu_k = \max_{i \in \mathcal{U}} \mu_i; x_{kB} < x_{jB} = \max_{i \in \mathcal{U}} x_{iB}$. Note that $\bar{z} < x_{jB}$ by assumption. We show that inspecting box k first, and then applying the policy in Lemma 0 to box j dominates stopping, and getting payoff μ_k . Therefore, it can't be optimal to stop, a contradiction.

If the agent inspects box k , and then applies Lemma 0 to inspect/take without inspection box j , his payoff is:

$$\Pi_{kj} = -k_k + \int_{x_j^R}^{+\infty} x_k dF_k + \int_{x_{jB}}^{x_j^R} (-k_j + \int \max\{x_j, x_k\} dF_j) dF_k + \int_{-\infty}^{x_{jB}} \mu_j dF_k$$

The payoff of stopping, and taking a box without inspection is given by μ_k . By definition:

$$\mu_k = -k_k + \int_{-\infty}^{x_{kB}} x_{kB} dF_k + \int_{x_{kB}}^{+\infty} x_k dF_k$$

Therefore, we can write:

$$\begin{aligned}
\Pi_{kj} - \mu_k &= \\
&- \int_{x_{kB}}^{x_j^R} x_k dF_k + \int_{x_{jB}}^{x_j^R} (-k_j + \int \max\{x_j, x_k\} dF_j) dF_k + \int_{-\infty}^{x_{kB}} (\mu_j - x_{kB}) dF_k \\
&+ \int_{x_{kB}}^{x_{jB}} \mu_j dF_k \\
&= \int_{x_{jB}}^{x_j^R} \left(\int_{x_j^R}^{+\infty} x_j^R dF_j + \int_{-\infty}^{x_j^R} \max\{x_j, x_k\} dF_j - x_k \right) dF_k + \int_{x_{kB}}^{x_{jB}} (\mu_j - x_k) dF_k \\
&+ \int_{-\infty}^{x_{kB}} (\mu_j - x_{kB}) dF_k
\end{aligned}$$

and, note the above is strictly positive: (i) the first integrand is non-negative because $\max\{\min\{x_j, x_j^R\}, x_k\} \geq x_k$ when $x_k < x_j^R$, (ii) the second integrand is positive because $\mu_j > x_{jB} > x_k$ by assumption 1, and (iii) the third integrand is positive because $\mu_j > x_{jB} > x_{kB}$. This shows that $\arg \max_{i \in \mathcal{U}} x_{iB} \subseteq \arg \max_{i \in \mathcal{U}} \mu_i$.

Now suppose that $\arg \max_{i \in \mathcal{U}} x_i^R \not\subseteq \arg \max_{i \in \mathcal{U}} x_{iB}$. Then, there exists $(\exists j, l \in \mathcal{U})(j \neq l) : x_j^R < x_l^R = \max_{i \in \mathcal{U}} x_i^R, x_{lB} < x_{jB} = \max_{i \in \mathcal{U}} x_{iB}$. Note that $\bar{z} < x_{jB} < x_j^R < x_l^R$. Consider the following policy: inspect box l first, and apply the policy in Lemma 0 to box j . By Proposition 1, and the fact that $x_{lB} < x_{jB} \leq x_j^R < x_l^R$, and $\bar{z} \leq x_{jB}$, we know that:

$$-k_l + \int_{x_j^R}^{+\infty} x_l dF_l + \int_{x_{jB}}^{x_j^R} (-k_j + \int \max\{x_l, x_j\} dF_j) dF_l + \int_{-\infty}^{x_{jB}} \mu_j dF_l > \mu_j$$

Therefore, $\arg \max_{i \in \mathcal{U}} x_{iB} \subseteq \arg \max_{i \in \mathcal{U}} x_i^R$, and note that we can actually conclude that both sets are equal.

Finally, we show that $\arg \max_{i \in \mathcal{U}} x_i^R, \arg \max_{i \in \mathcal{U}} x_{iB}$ are singletons. Suppose not. Then $(\exists l, l' \in \mathcal{U}) : x_l^R = x_{l'}^R, x_{lB} = x_{l'B}$. Moreover, by the previous step, we have that $\mu_l = \mu_{l'}$. Consider the following policy: inspect box l first, and then apply the policy in Lemma 0 for inspecting box l' . This improves upon stopping and

taking box l' , because:

$$\begin{aligned}
& -k_l + \int_{x_l^R}^{+\infty} x_l dF_l + \int_{x_{l'B}^R}^{x_{l'}^R} (-k_{l'} + \int \max\{x_l, x_{l'}\} dF_{l'}) dF_l + \int_{-\infty}^{x_{l'B}^R} \mu_{l'} dF_l \\
& = \int_{x_l^R}^{+\infty} (x_l^R - \mu_{l'}) dF_l + \int_{x_{l'B}^R}^{x_{l'}^R} \left(\int_{x_{l'}^R}^{+\infty} x_l^R + \int_{-\infty}^{x_{l'}^R} \max\{x_l, x_{l'}\} dF_{l'} - \mu_{l'} \right) dF_l > 0
\end{aligned}$$

where the first equality comes from using equation (1) for boxes l and l' , and the inequality comes from Assumption 1, and the fact that $\mu_{l'} = \int_{x_{l'}^R}^{+\infty} x_l^R + \int_{x_{l'B}^R}^{x_{l'}^R} x_{l'} dF_{l'} + \int_{-\infty}^{x_{l'B}^R} x_{l'B} dF_{l'}$ (note that the inequality is an equality only when $x_l^R = \mu_l = x_{lB}$).

Therefore, we conclude that $\arg \max_{i \in \mathcal{U}} x_i^R, \arg \max_{i \in \mathcal{U}} x_{iB}$ are both singletons. Moreover, the proof shows that the box with the highest backup value is taken without inspection when it is optimal to stop. \square

Lemma 3. Let \mathcal{U} be a set of boxes, and z be the vector of previously realized prizes. Assume that $\sigma^*(\mathcal{U}, z) = j$, where $x_j^R < \max_{i \in \mathcal{U}} x_i^* \equiv x_l^R$. Then, it cannot be the case that $\sigma^*(\mathcal{U} \setminus \{j\}, z \cup \{x_j\}) = l$, and $\varphi^*(\mathcal{U} \setminus \{j\}, z \cup \{x_j\}) = 0$, whenever $\max\{x_j, \bar{z}\} \leq x_l^R$.

Proof. Suppose $\sigma^*(\cdot) = \{j\}$ and the optimal continuation policy dictates inspecting box l whenever $\max\{x_j, \bar{z}\} \leq x_l^R$. The following policy improves on this, as shown by Proposition 1: inspect box l first. Whenever $x_l^R < x_l$, stop. Otherwise, open box j and then proceed by using the prescribed policy when $\mathcal{U} = \mathcal{U} \setminus \{l, j\}$. \square

A.3 Proof of Proposition 3

Proof.

Lemma 4 (Cutoffs are linear in means). Let x be a random variable such that $x \sim F(\cdot - \mu), E[x] = \mu$. Let k be the cost of inspecting the box with prizes distributed according to F . Then, $(\exists \underline{b}, \bar{b}) : x_B = \mu - \underline{b}, x^R = \bar{b}$.

Proof. We do the proof for x^R , the other one follows immediately. Recall that:

$$k = \int_{x^R}^{+\infty} (x - x^R) dF(x - \mu)$$

We guess and verify that $x^R = \mu + \bar{b}$, for some $\bar{b} > 0$.

$$k = \int_{\mu+\bar{b}}^{+\infty} (x - \mu - \bar{b}) dF(x - \mu)$$

Let $u = x - \mu$ and perform a change of variables in the above expression:

$$k = \int_{\bar{b}}^{+\infty} (u - \bar{b}) dF(u) \quad (6)$$

It remains to show that there is a solution to the above equation. Note that assumption 1 implies that if $\bar{b} = 0$, then $k < \int_0^{+\infty} u dF(u)$. On the other hand, as $\bar{b} \rightarrow \infty$, $\int_{\bar{b}}^{+\infty} (u - \bar{b}) dF(u) \rightarrow 0 < k$. Hence, since $g(b) = \int_b^{+\infty} (x - b) dF$ is continuous and decreasing in b , there exists $\bar{b} > 0$, such that the equality holds. This completes the proof. \square

Corollary 2. *Consider the same assumptions as before. If F is symmetric around 0 then $\bar{b} = \underline{b} = b > 0$*

Proof. The fact that $b > 0$ comes from the condition that $x_B < \mu < x^R$ for the problem to be well-defined.

Recall the definition of x_B :

$$k = \int_{-\infty}^{x_B} (x_B - x) dF(x - \mu)$$

Replacing our assumptions we get that the equation can be rewritten as:

$$k = \int_{-\infty}^{-\underline{b}} (-\underline{b} - u) dF(u)$$

where we changed variables by defining $u = x - \mu$. Also, we have that:

$$k = \int_{x^R}^{+\infty} (x - x^R) dF(x - \mu) = \int_{\bar{b}}^{+\infty} (u - \bar{b}) dF(u)$$

Now, symmetry of F implies that:

$$\int_{\bar{b}}^{+\infty} u dF(u) = - \int_{-\infty}^{-\bar{b}} u dF(u)$$

Hence, $(1 - F(\bar{b}))E[u|u \geq \bar{b}] = -F(-\bar{b})E[u|u \leq -\bar{b}]$ and $-(1 - F(\bar{b}))\bar{b} = -F(-\bar{b})\bar{b}$. Hence, $\bar{b} = \underline{b}$. \square

Now we are ready to prove Proposition 3. We start with the case in which $0 \leq x_{jB} \leq x_{iB}$, $x_j^R \leq x_i^R$. Equation (4) in Section 5 established that the difference between opening box i first and opening box j first is given by:

$$\begin{aligned}
\Pi_{ij} - \Pi_{ji} &= \int_{-\infty}^{x_{iB}} \int_{-\infty}^{x_{iB}} (\max\{x_i, x_j, x_{jB}\} - x_{iB}) dF_i dF_j \\
&+ \int_{x_j^R}^{+\infty} \int_{x_j^R}^{+\infty} (\min\{x_i, x_j, x_i^R\} - x_j^R) dF_i dF_j \\
&= (1 - F_i(x_i^R))(1 - F_j(x_i^R))(x_i^R - x_j^R) + \int_{x_j^R}^{x_i^R} \int_{x_i}^{+\infty} (x_i - x_j^R) dF_j dF_i \\
&+ \int_{x_j^R}^{x_i^R} \int_{x_j^R}^{x_i} (x_j - x_j^R) dF_j dF_i + (1 - F_i(x_i^R)) \int_{x_j^R}^{x_i^R} (x_j - x_j^R) dF_j \\
&+ F_i(x_{jB})F_j(x_{jB})(x_{jB} - x_{iB}) + F_i(x_{jB}) \int_{x_{jB}}^{x_{iB}} (x_j - x_{iB}) dF_j \\
&+ \int_{x_{jB}}^{x_{iB}} \int_{x_i}^{x_{iB}} (x_j - x_{iB}) dF_j dF_i + \int_{x_{jB}}^{x_{iB}} \int_{-\infty}^{x_i} (x_i - x_{iB}) dF_j dF_i
\end{aligned}$$

Replacing our assumptions, $u = x_i - \mu_i$, $\hat{u} = x_j - \mu_j$ and writing $a = \mu_i - \mu_j \geq 0$, we have that:

$$\begin{aligned}
G(a) &= \int_{b-a}^b \int_{u+a}^{+\infty} (u + a - b) dF(\hat{u}) dF(u) + \int_{b-a}^b \int_b^{u+a} (\hat{u} - b) dF(\hat{u}) dF(u) \\
&+ F(-b) \int_b^{b+a} (\hat{u} - b) dF(\hat{u}) + F(-b-a) \int_{-b}^{-b+a} (\hat{u} + b - a) dF(\hat{u}) \\
&+ \int_{-b-a}^{-b} \int_{u+a}^{-b+a} (\hat{u} + b - a) dF(\hat{u}) dF(u) + \int_{-b-a}^{-b} \int_{-\infty}^{u+a} (u + b) dF(\hat{u}) dF(u)
\end{aligned}$$

Note that $G(0) = 0$. We will show that $G'(0) = 0$, $G''(a) = 0(\forall a)$. All of these

together imply that $G(a) \equiv 0$.

$$\begin{aligned} G'(a) &= -\left[\int_{b-a}^b F(-b-a)dF(u) + \int_{-b-a}^{-b} (F(-b+a) - F(u+a))dF(u) \right. \\ &\quad \left. - \int_{b-a}^b F(-u-a)dF(u)\right] \end{aligned}$$

Note that $G'(0) = 0$. Moreover,

$$\begin{aligned} G''(a) &= F(-b-a)f(b-a) - \int_{b-a}^b f(-b-a)dF(u) + (F(-b-a) - F(-b))f(-b-a) \\ &\quad + \int_{-b-a}^{-b} (f(-b+a) - f(u+a))dF(u) - F(-b)f(b-a) + \int_{b-a}^b f(-u-a)dF(u) = 0 \end{aligned}$$

where we used that $f(x) = f(-x)$, $F(-x) = 1 - F(x)$ several times to cancel terms. This shows that $G(a) \equiv 0$. When $x_{jB} \leq 0 \leq x_{iB}$, we have that:

$$\begin{aligned} \Pi_{ij} - \Pi_{ji} &= \int_{-\infty}^{x_{iB}} \int_{-\infty}^{x_{iB}} (\max\{x_i, x_j, 0\} - x_{iB})dF_i dF_j \\ &\quad + \int_{x_j^R}^{+\infty} \int_{x_j^R}^{+\infty} (\min\{x_i, x_j, x_i^R\} - x_j^R)dF_i dF_j \end{aligned}$$

Since the previous proof never used the fact that $x_{jB} \geq 0$, and $x_{jB} < 0$ in this case, this shows that the previous difference is positive. Finally, when $x_{jB} \leq x_{iB} \leq 0$, the problem is exactly as Weitzman's, hence we know that the difference is strictly positive. This completes the proof. \square

A.4 Boxes for which $x^R < x_B$ are never inspected in the optimal policy

This last subsection shows that, if we allow for boxes $i \in \mathcal{N}$ such that $x_i^R < x_{iB}$, then box i is never inspected in the optimal policy. Therefore, for any such box $i \in \mathcal{N}$, it is either taken without inspection upon stopping search, or it is never used in the optimal policy. Moreover, note that only one such box can be taken

without inspection conditional on stopping search. Then, by redefining x_0 to be whatever is best between the agent's initial outside option and the best of the boxes for which $x_i^R < x_{iB}$, our analysis carries through by focusing on the boxes for which $x_{iB} < x_i^R$.

Given a set of boxes \mathcal{U} , define:

$$\begin{aligned}\mathcal{U}^{B < R} &= \{i \in \mathcal{U} : x_{iB} \leq x_i^R\} \\ \mathcal{U}^{R < B} &= \{i \in \mathcal{U} : x_i^R < x_{iB}\}\end{aligned}$$

Given a decision node (\mathcal{U}, z) , we use $(\mathcal{U}', z'), \mathcal{U}' \subset \mathcal{U}, z' = (z, z_{\mathcal{U} \setminus \mathcal{U}'})$ to denote a generic continuation history in which boxes in $\mathcal{U} \setminus \mathcal{U}'$ have been inspected, and prizes $z_{\mathcal{U} \setminus \mathcal{U}'}$ have been sampled.

Proposition 4. *Let \mathcal{U} be the set of boxes, and let z be a vector of realized prizes. Assume that $\mathcal{U}^{R < B} \neq \emptyset$. Then, $\forall i \in \mathcal{U}^{R < B}$, $\nexists (\mathcal{U}', z') : i \in \mathcal{U}' \subset \mathcal{U}, z' = (z, \tilde{z}_{\mathcal{U} \setminus \mathcal{U}'})$ such that $\varphi^*(\mathcal{U}', z') = 0, \sigma^*(\mathcal{U}', z') = i$.*

Proof. The proof is by double induction in the cardinality of \mathcal{U} and $\mathcal{U}^{R < B}$. Since $\mathcal{U}^{R < B} \subset \mathcal{U}$, we know that $|\mathcal{U}^{R < B}| \leq |\mathcal{U}|$. Induction will be done in $U = |\mathcal{U}|$, and n , where $|\mathcal{U}^{R < B}| = \max\{U, n\}$. Let $P(U, n)$ denote the following predicate:

P(U,n): $(\forall z)(\forall \mathcal{U}) : |\mathcal{U}| = U, \mathcal{U}^{R < B} \neq \emptyset, |\mathcal{U}^{R < B}| = \max\{n, U\}$, the optimal policy satisfies the property in Proposition 4.

We proceed by showing that $P(1, 1) = 1$, and that if $P(U', n') = 1$ holds for $U' \leq U$, and $n' \leq n$, not both with equality, then $P(U, n) = 1$ holds.

P(1,1)=1: Let $\mathcal{U} = \{i\}$ and let z denote the vector of already realized prizes.

Since $U = n = 1$, we have that $\mathcal{U}^{R < B} = \{i\}$. We show that:

$$-k_i + \int \max\{x_i, \bar{z}\} dF_i \leq \max\{\mu_i, \bar{z}\}$$

Suppose that $\bar{z} \geq \mu_i$. Then, since $i \in \mathcal{U}^{R < B}$, $x_i^R < \mu_i \leq \bar{z}$. Then,

$$-k_i + \int \max\{x_i, \bar{z}\} dF_i - \bar{z} = -k_i + \int_{\bar{z}} (x_i - \bar{z}) dF_i(x_i) < 0$$

since $\bar{z} > x_i^R$ (recall the derivation of equation (1)). Now, suppose that $\mu_i > \bar{z}$. Then, $x_{iB} > \mu_i > \bar{z}$, and it follows from (2) that:

$$-k_i + \int \max\{x_i, \bar{z}\} dF_i - \mu_i = -k_i + \int_{-\infty}^{\bar{z}} (\bar{z} - x_i) dF_i(x_i) < 0$$

P(U,n)=1: Assume now that $(\forall U' \leq U)(\forall n' \leq n)$, not both with equality, $P(U', n') = 1$. We show that $P(U, n) = 1$. Let \mathcal{U} be the set of boxes, $|\mathcal{U}| = U$, and let z denote the vector of already sampled prizes. Let $\mathcal{U}^{R<B} \subset \mathcal{U}$, $|\mathcal{U}^{R<B}| = \max\{U, n\}$. We use i to denote a box in $\mathcal{U}^{R<B}$, and j to denote a box in $\mathcal{U} \setminus \mathcal{U}^{R<B}$, whenever the latter is not empty.

We make two remarks. First, notice that if a box $j \in \mathcal{U} \setminus \mathcal{U}^{R<B}$ is inspected, then we move to continuation history $(\mathcal{U}', z \cup \{x_j\})$, where $\mathcal{U}' = \mathcal{U} \setminus \{j\}$, $\mathcal{U}'^{R<B} = \mathcal{U}^{R<B}$, and $|\mathcal{U}'| = U - 1$, and $\mathcal{U}'^{R<B} = n$ (note that if there was $j \in \mathcal{U} \setminus \mathcal{U}^{R<B}$, then it can't be the case that $|\mathcal{U}^{R<B}| = U$). Since, by the inductive step, we know that $P(U - 1, n) = 1$, then boxes in $\mathcal{U}^{R<B}$ are not inspected in any continuation history. Second, if a box $i \in \mathcal{U}^{R<B}$ were to be inspected, then we move to continuation history $(\mathcal{U}', z \cup \{x_i\})$, where $\mathcal{U}' = \mathcal{U} \setminus \{i\}$, $\mathcal{U}'^{R<B} = \mathcal{U}^{R<B} \setminus \{i\}$, and $|\mathcal{U}'| = U - 1$, $|\mathcal{U}'^{R<B}| = \max\{U - 1, n - 1\}$. Since, by the inductive step, we know that $P(U - 1, n - 1) = 1$, then boxes in $\mathcal{U}'^{R<B}$ are not inspected in any continuation history. The first remark implies that to prove $P(U, n) = 1$ it remains to show that no box in $\mathcal{U}^{R<B}$ is inspected in history (\mathcal{U}, z) . The second remark will be used when computing the payoff of inspecting a box in $i \in \mathcal{U}^{R<B}$.

Given the above, we want to show that:

$$\begin{aligned} & \max \left\{ \bar{z}, \max_{i \in \mathcal{U}^{R<B}} \mu_i, \max_{j \in \mathcal{U}^{B<R}} \mu_j, \max_{j \in \mathcal{U}^{B<R}} \left\{ -k_j + \int V^*(\mathcal{U} \setminus \{j\}, z \cup \{x_j\}) dF_j \right\} \right\} \\ & \geq \max_{i \in \mathcal{U}^{R<B}} \left\{ -k_i + \int V^*(\mathcal{U} \setminus \{i\}, z \cup \{x_i\}) dF_i \right\} \end{aligned} \quad (7)$$

where the LHS of the above expression denotes the payoff the agent can get by either stopping, and getting $\max\{\bar{z}, \max_{i \in \mathcal{U}^{R<B}} \mu_i, \max_{j \in \mathcal{U}^{B<R}} \mu_j\}$, or continuing search by inspecting a box in $\mathcal{U}^{B<R}$; the RHS denotes the payoff of inspecting a box in $\mathcal{U}^{R<B}$. The stars in V denote that the agent follows

the optimal policy in the continuation histories, and the two remarks above apply, by the inductive step to those histories.

Note that we can write, for any box $i \in \mathcal{U}^{R < B}$:

$$\begin{aligned}
& -k_i + \int V^*(\mathcal{U} \setminus \{i\}, z \cup \{x_i\}) dF_i \\
&= -k_i + \int \max \left\{ \begin{array}{c} x_i, \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \\ \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i, x_j\}) dF_j\} \end{array} \right\} dF_i \\
&= -k_i + \int \max \left\{ x_i, \max \left\{ \begin{array}{c} \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \\ \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i, x_j\}) dF_j\} \end{array} \right\} \right\} dF_i \\
&= \int_{x_i^R}^{+\infty} x_i^R + \max \left\{ 0, \max \left\{ \begin{array}{c} \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \\ \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i, x_j\}) dF_j\} \end{array} \right\} - x_i \right\} \\
&+ \int_{-\infty}^{x_i^R} \max \left\{ x_i, \max \left\{ \begin{array}{c} \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \\ \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i, x_j\}) dF_j\} \end{array} \right\} \right\} dF_i
\end{aligned}$$

where the first equality is by definition of the set of actions available to the agent, and we use the second remark above; the second equality is just a rearrangement of terms, and the third equality follows from using (1) for box i .

Notice that the second term in the first integrand:

$$\max \left\{ 0, \max \left\{ \begin{array}{c} \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \\ \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i, x_j\}) dF_j\} \end{array} \right\} - x_i \right\}$$

is decreasing in x_i : the slope of $-x_i$ is -1 , and the slope of the term in the $\max\{\cdot\}$ as a function of x_i is at most one (it would be 1 only if x_i is better than any of the terms in the $\max\{\cdot\}$ for all $x_i \in [x_i^R, +\infty]$). Therefore, we have that:

$$\begin{aligned}
& \int_{x_i^R}^{+\infty} x_i^R + \max \left\{ 0, \max \left\{ \begin{array}{c} \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \\ \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i, x_j\}) dF_j\} \end{array} \right\} - x_i \right\} \\
& \leq \int_{x_i^R}^{+\infty} \max \left\{ x_i^R, \max \left\{ \begin{array}{c} \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \\ \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i^R, x_j\}) dF_j\} \end{array} \right\} \right\} dF_i
\end{aligned}$$

Also, we have that:

$$\begin{aligned}
& \int_{-\infty}^{x_i^R} \max \left\{ x_i, \max \left\{ \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \right. \right. \\
& \quad \left. \left. \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i, x_j\}) dF_j\} \right\} \right\} dF_i \\
& \leq \int_{-\infty}^{x_i^R} \max \left\{ x_i^R, \max \left\{ \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \right. \right. \\
& \quad \left. \left. \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i^R, x_j\}) dF_j\} \right\} \right\} dF_i
\end{aligned}$$

since the integrand is increasing in x_i . Putting all of this together, we conclude that for all $i \in \mathcal{U}^{R < B}$, the following holds:

$$\begin{aligned}
& -k_i + \int V^*(\mathcal{U} \setminus \{i\}, z \cup \{x_i\}) dF_i \\
& = -k_i + \int \max \left\{ x_i, \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \right. \\
& \quad \left. \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i, x_j\}) dF_j\} \right\} dF_i \\
& \leq \max \left\{ x_i^R, \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \right. \\
& \quad \left. \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i^R, x_j\}) dF_j\} \right\}
\end{aligned}$$

But, then we conclude that, for all $i \in \mathcal{U}^{R < B}$:

$$\begin{aligned}
& \max \left\{ \bar{z}, \max_{i \in \mathcal{U}^{R < B}} \mu_i, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{j\}, z \cup \{x_j\}) dF_j\} \right\} \\
& \geq \max \left\{ x_i^R, \bar{z}, \max_{i' \in \mathcal{U}^{R < B} \setminus \{i\}} \mu_{i'}, \max_{j \in \mathcal{U}^{B < R}} \mu_j, \right. \\
& \quad \left. \max_{j \in \mathcal{U}^{B < R}} \{-k_j + \int V^*(\mathcal{U} \setminus \{i, j\}, z \cup \{x_i^R, x_j\}) dF_j\} \right\} \\
& \geq -k_i + \int V^*(\mathcal{U} \setminus \{i\}, z \cup \{x_i\}) dF_i
\end{aligned}$$

where the first inequality follows from $x_i^R < \mu_i$ for $i \in \mathcal{U}^{R < B}$, and the fact that taking box i without inspection and getting μ_i is always an option in the optimal policy in the first line, while not in the second.

Since the above holds for each $i \in \mathcal{U}^{R < B}$, we conclude that (7) holds, and, thus, $P(U, n) = 1$

□

References

- [1] M. Armstrong and J. Zhou. Search deterrence. *Oxford University Economics Series Working Papers*, (661), 2013.
- [2] E. Ben-Porath, E. Dekel, and B. L. Lipman. Optimal allocation with costly verification. Mar. 2013.
- [3] H. Chade and N. Kovrijnykh. Information acquisition, moral hazard, and rewarding for bad news. *Working Paper*, 2011.
- [4] H. Chade and L. Smith. Simultaneous search. *Econometrica*, 74(5):1293–1307, 2006.
- [5] G. De Clippel, K. Eliaz, and K. Rozen. Competing for consumer inattention. *Working Paper*, 2013.
- [6] J.-J. Ganuza and J. S. Penalva. Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica*, 78(3):1007–1030, 2010.
- [7] D. Klabjan, W. Olszewski, and A. Wolinsky. Attributes. *Mimeo*, 2011.
- [8] D. Kräbmer and R. Strausz. Optimal procurement contracts with pre-project planning. *The Review of Economic Studies*, 78(3):1015–1041, 2011.
- [9] H. Li and X. Shi. Discriminatory information disclosure. *Working paper*, 2013.
- [10] M. Mueller-Frank and M. Pai. Social search. *Available at SSRN 2166925*, 2012.
- [11] W. Olszewski. A more general pandora rule. *Mimeo*, 2012.
- [12] C. Papadimitriou and J. Tsitsiklis. The complexity of optimal queuing network control. *Mathematics of Operations Research*, pages 293–305, 1999.
- [13] P. Postl. Delegated search: Procedure matters. *Working Paper*, 2004.
- [14] T. Vishwanath. Parallel search for the best alternative. *Economic Theory*, 2(4):495–507, 1992.

- [15] M. Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.