Whether or not to open Pandora's box

Laura Doval*†

First draft: November 2012. This draft: October 2014

Abstract

Weitzman's [12] 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 identify conditions under which 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 René Leal Vizcaíno, 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 [12] model has been used to study situations 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 (see Olszewski [8], and the references therein). 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 [12] 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. Our paper addresses Weitzman's search problem without this assumption. Within this framework, we find sufficient conditions under which the optimal *order* coincides with Weitzman's. However, the optimal stopping rule is different and we characterize it (see Section 1.1, and Theorems 1, 2 and 3 in Section 4).

Before discussing our results in detail, consider the following application where Weitzman's assumption is counterintuitive (see Section 6 for more applications). Say the agent is a student who has to make a choice between schools A and B, to which he has been admitted, or not going to school. He derives 0 utility from the latter option. 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.

We now use the example to show how the optimal policy changes without Weitzman's assumption. Assume that each school's distribution over prizes is given by the following table, based on an example by Postl [10]:

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

Table 1: Prize distribution for each school

Under Weitzman's assumption, the following is the optimal policy: school A is 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 (see Appendix B.1 for a proof). Intuitively, this is because (i) going to school dominates not going, (ii) the agent can only go to a school if he attends its visit day first, and (iii) attending visit days is costly. Hence, the student visits at least one school, and, by visiting school A first, if $x_A = 5$, he saves school B's inspection cost.

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 policy, because the value of visiting school B first improves now that the student, after finding out $x_B = 0$, can save on the inspection costs of school A (since $x_A > 0$).[1] Indeed, in the optimal policy, school B is 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. The optimal policy differs from Weitzman's rule both in that school B is visited first and in the stopping rule: when $x_A = 5$, search stops in the first round under Weitzman's rule, while it may continue (if $x_B = 3$) when his assumption does not hold.

1.1 Discussion of results

In our model, each box is characterized by *two* cutoff values, defined formally in Section 2.2 (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

[1]The example is stark for expositional purposes. For an example in which the same effect obtains, and school A is not ex-ante better than school B see Appendix B.2.

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 optimal to inspect the box if the maximum sampled prize is below 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, $\overline{z}$, when there is only one box left to inspect with mean μ:

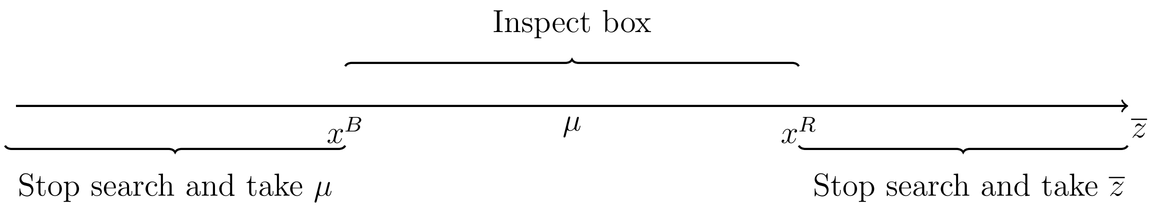

Figure 1: Optimal policy for one box

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 2.2 implies that $x^B < x^R$. This assumption rules out uninteresting cases (see Appendix A.5 for details).

Two properties of the optimal policy in the one-box case carry over to the case with $N > 1$ boxes left to be inspected. First, when the maximum sampled prize is higher than the highest backup value amongst uninspected boxes, then the optimal order and stopping policy coincide with Weitzman's from that moment on (this corresponds when $N = 1$ to $x^B \leq \overline{z}$ in Figure 1); see Proposition 1. Second, if the agent finds it optimal to stop, and take a box without inspection, then this box is the box with the highest backup *and* reservation value amongst uninspected boxes (this holds trivially when there is one box left to be inspected); this is recorded in Proposition 2. It follows from this property that the agent would like to leave the boxes with high backup values to be taken without inspection, while Weitzman shows that boxes should be inspected in decreasing order of their reservation values. Since boxes with high backup values may be the boxes with high reservation values, this introduces a trade-off when deciding which box to inspect next; Proposition 2 shows that search can only end with taking a box without inspection when this trade-off is maximal: the best box to inspect next (highest reservation value) is the best to take without inspection (highest backup value).

However, one property of the one-box case no longer holds when $N > 1$ boxes remain. While in the one-box case a box's cutoff values and the maximum sampled prize determine the optimal policy, this need not be the case when there are $N > 1$ boxes left to be inspected. As illustrated in Example B.2, two sets of boxes can share the same cutoffs, and yet have different optimal policies.[2]

The reason why more than the cutoff values matter to determine the optimal policy is that they don't necessarily determine the full "value" of a box. To see this, consider the example in Section 1. If only school A is available, it is optimal to accept school A without inspection. Now add school B, and note that it is worse than school A both to inspect *and* to take without inspection (Appendix B.1 shows school A has a higher reservation and backup value). We would then expect that the optimal policy remains the same when adding school B. However, this is not the case, because what dominates taking school A without inspection is inspecting school B *and* then choosing, given x_B, whatever is best between inspecting or taking school A without inspection. Thus, the comparison of the boxes' cutoffs alone is not enough to determine the optimal policy.

Our main contribution is to identify conditions under which the optimal *order* policy coincides with Weitzman's; the optimal *stopping* rule, however, is different and we characterize it as well (Section 4, Theorems 1-3). Identifying conditions under which Weitzman's order is optimal is useful because we retain the simplicity

[2]Formally, this is because 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 is shown to be P-SPACE hard in [9], and is usually not indexable. Our paper may be seen as a contribution to the restless bandit literature by providing instances of a restless bandit problem where the optimal policy can be described as an index rule.

of the reservation value rule, which is valuable for applications. Moreover, the conditions we identify have been used in the search literature to enable characterizing optimal policies in environments where, without these assumptions, said characterization has proved elusive (see the references below). Our results, then, show that the usefulness of these conditions extends to our environment as well.

Theorem 1 in Section 4 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 [12], 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. Theorem 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 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 [11] to obtain the reservation value rule in her parallel search model, while Chade and Smith [3] apply it in their simultaneous search model.

The conditions for Theorems 2 and 3 imply that, 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 in the probability of obtaining the highest prize (Theorem 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 (Theorem 3). The conditions in (i) allow us to extend Chade and Smith's [3] 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 to which the agent has been admitted.

Unlike Theorem 1, where the optimal stopping rule coincides with Weitzman's for all but the last box, under the conditions of Theorems 2 and 3 the agent may choose to stop and take a box without inspection, even before reaching the last box.

1.2 Related Literature

Postl [10] postulates this search problem explicitly within the context of a principal-agent model. He focuses on the two-box-equal-inspection-costs version of our search problem, and discusses an analogue of Theorem 1 in this simplified setting. Theorem 1 in our paper generalizes this result, showing that it is not necessary to assume two boxes nor that the boxes have equal costs.

Klabjan, Olszewski and Wolinsky [5] study a search for attributes model in which, contrary to our setting, the agent's utility function is given by the sum of the prizes (attributes). Like in our setting, the agent does not have to inspect all attributes in order to keep the object: he can accept the object, taking the rest of the attributes without inspection. Under sequential search, and with two boxes, the authors characterize the optimal solution when attribute distributions are symmetric around 0. The rule coincides with inspecting attributes in decreasing order of their reservation value (see Theorem 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 [2], Krähmer and Strausz [6]), and in Klabjan, Olszewski and Wolinsky [5]. Moreover, the backup value plays a crucial role in the optimal mechanism of Ben-Porath, Dekel and Lipman [1]. However, none of these papers provide a solution for the search problem analyzed here.

The rest of the paper is organized as follows. Section 2 describes the model, provides a formal definition of the cutoffs and intuition of their role in the search problem. Section 3 provides a series of properties the optimal sampling policy must satisfy regardless of the environment. Section 4 focuses on the optimal order and stopping policies, and contains our main results. The statements of Theorems 1-3 are presented informally to streamline notation; the Appendix contains the formal statements. Section 5 describes the optimal policy when the agent possesses only two boxes. Section 6 concludes. Proofs are relegated to the Appendix.

2 Model

An agent possesses a set $\mathcal{N} = \{1, ..., N\}$ 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 independently distributed, and, for all $i \in \mathcal{N}$, $\int |x_i| dF_i(x_i) < +\infty$. The agent has an initial outside option, x_0, normalized to 0. Given a vector z, we denote by $\overline{z}$, its highest coordinate. We assume the agent is risk neutral, and given a vector of realized prizes $z = (z_1, ..., z_n)$, his utility function is given by $u(z) = \overline{z}$.

2.1 Sampling Policy

The agent inspects boxes sequentially, and search is with recall. Given a set of uninspected boxes $\mathcal{U}$, and a vector of realized sampled prizes z, the agent decides whether to stop, or continue search; if he decides to continue search, he decides which box to inspect next. Let $\varphi(\mathcal{U}, z) \in \{0, 1\}$ denote the decision to stop search ($\varphi = 0$), or to continue search ($\varphi = 1$) at decision node $(\mathcal{U}, z)$; if $\varphi(\mathcal{U}, z) = 1$, let $\sigma(\mathcal{U}, z) \in \mathcal{U}$ denote the box which he inspects next. If $\varphi = 0$, the agent chooses between any prize in z, and any uninspected box in $\mathcal{U}$. If $\varphi = 1$, he inspects box σ, pays k_σ, and observes its prize x_σ. Having observed x_σ, the agent is now at decision node $(\mathcal{U} \backslash \{\sigma\}, z \cup \{x_\sigma\})$, and selects $\varphi(\mathcal{U} \backslash \{\sigma\}, z \cup \{x_\sigma\})$, and $\sigma(\mathcal{U} \backslash \{\sigma\}, z \cup \{x_\sigma\})$. Given a decision node $(\mathcal{U}, z)$, the strategy φ, σ, together with the distributions $\{F_i\}_{i \in \mathcal{U}}$ determine a probability distribution over continuation paths in the natural way, and the agent's expected payoff at that decision node, which we denote $V(\mathcal{U}, z)$. We use stars to denote the optimal strategies, and the payoff V when it results from using the optimal policy in $(\mathcal{U}, z)$.

At decision node $(\mathcal{U}, z)$, the agent's optimal strategy solves the following problem:

$$V^*(\mathcal{U}, z) = \max\{\overline{z}, \max_{i\in\mathcal{U}} \mu_i, \max_{i\in\mathcal{U}} -k_i + \int V^*(\mathcal{U}\backslash\{i\}, z \cup \{x_i\}) dF_i(x_i)\}$$

2.2 Cutoffs

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

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

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

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

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 $\overline{z} \leq x_i^R$. Using equation (1), we can write the payoff from inspecting box i when $\overline{z} \leq x_i^R$ as follows:

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

This shows that the reservation value represents the highest prize the agent expects to get from inspecting box i, after internalizing inspection costs, since it is *as if* the agent's payoff from inspecting box i is bounded above by x_i^R.

Consider now the case $\overline{z} \leq \mu_i$. If the agent stops, he takes the box without inspection. Therefore, the agent inspects box i if and only if the following holds:[3]

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

The above expression shows that, conditional on inspecting box i, 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 left hand side. Define the *backup value* to be the value x_i^B such that:

$$k_i = \int_{-\infty}^{x_i^B} (x_i^B - x_i) dF_i(x_i) \tag{2}$$

i.e., x_i^B 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 $x_i^B \leq \bar{z}$; otherwise, he takes it without inspection. Equation (2) can be written as:

$$\mu_i = -k_i + \int_{-\infty}^{x_i^B} x_i^B dF_i(x_i) + \int_{x_i^B}^{+\infty} x_i dF_i(x_i) \tag{3}$$

Using equation (1), equation (3) can be written as:

$$\mu_i = \int_{-\infty}^{x_i^B} x_i^B dF_i(x_i) + \int_{x_i^B}^{x_i^R} x_i dF_i(x_i) + \int_{x_i^R}^{+\infty} x_i^R dF_i(x_i) \tag{4}$$

Equation (4) illustrates that x_i^B is the lowest prize the agent expects to get from box i when he takes it without inspection, after internalizing that he did not pay box i's inspection cost. We refer to x_i^B as box i's backup value because, when agent i takes box i without inspection, it is *as if* his payoff is bounded below by x_i^B.

Throughout, we make the following assumption to ensure that $x_i^B \leq \mu_i \leq x_i^R$ always holds:

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

If the set of boxes $\mathcal{N}$ contains at least one box that violates Assumption 1, then

[3] We use that $\mu_i = \int_{-\infty}^{+\infty} x_i dF_i(x_i)$ to go from the LHS to the RHS.

said boxes are never inspected (see Appendix A.5 for a proof, and for a discussion of why Assumption 1 rules out such uninteresting cases).

When there is one box left, the optimal policy is determined by comparing the maximum sampled prize, $\overline{z}$, with the cutoffs, x^B, x^R. This is summarized in Figure 1 in Section 1.1, and is recorded in Proposition 0 below.

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

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

Results similar to Proposition 0 have appeared in the one-box-settings of Chade and Kovrijnykh [2], and Krähmer and Strausz [6], in the two-box setting of Postl [10], as well as in the attributes model of Klabjan, Olszewski and Wolinsky [5].

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

2.3 A different interpretation of x^R and x^B

We conclude Section 2 by providing a different interpretation for the cutoffs. Consider again the case $N = 1$. When $\overline{z} > \mu$, if the agent stops, he chooses to take $\overline{z}$. However, if he inspects the box, he may discover that it contains a prize better than $\overline{z}$, increasing his *ex-post* payoff by $x - \overline{z}$. 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: $V_I(\overline{z}) = \int_{\overline{z}}^{+\infty}(x - \overline{z})dF(x)$. When $\overline{z} \leq \mu$, 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 $\overline{z}$, which yields an *ex-post* loss of $\overline{z} - x$. By inspecting the box, the agent avoids taking boxes that are worse than what he has (type II error). *Ex-ante*, this is worth $V_{II}(\overline{z}) = \int_{-\infty}^{\overline{z}}(\overline{z} - x)dF(x)$ to the agent.

Thus, given the outside option, $\overline{z}$, the value of the new information for the agent is given by:

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

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

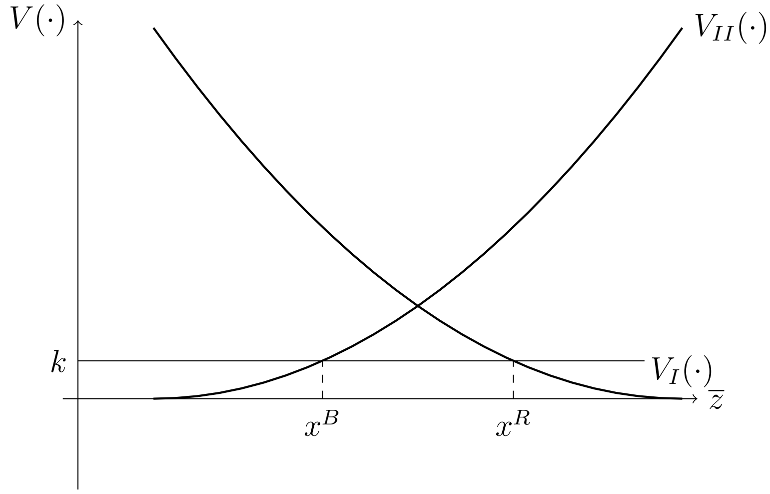

Figure 2: The value of information

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

$V_I(\overline{z})$ 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 prize. $V_{II}(\overline{z})$ 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 Proposition 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.

3 Preliminary results

Section 3 presents three building blocks in determining the optimal policy. Propositions 1 and 2 below formalize the claim that the backup value of box i represents the value of the option of taking box i without inspection. Proposition 3 illustrates the reasons for deviating from Weitzman's ordering. Propositions 1-3 also help simplify the taxonomy of the problem: when the conditions in Section 4 don't hold, and hence, the optimal policy must be computed by backward induction, the results in this section help narrow down the cases to be considered. This is illustrated in Section 5, where we characterize the optimal policy when $N = 2$. To state the results, 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, $\overline{z}$, is greater than μ_i, then, from then onwards, the optimal sampling policy is given by applying Weitzman's rule to the boxes in $\mathcal{U}$. Proposition 1 shows that, while sufficient, 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.

Proposition 1. *Let $(\mathcal{U}, z)$ denote the set of boxes, and the vector of realized prizes, respectively. If for all $i \in \mathcal{U}$, $x_i^B \leq \overline{z}$, then Weitzman's optimal sampling policy is optimal in all continuation histories.*

Proposition 2 shows that if the agent finds it optimal to stop, and take a box without inspection at decision node $(\mathcal{U}, z)$, then the chosen box, which is the box with the highest μ_i, satisfies three properties: (i) it has the highest backup value, (ii) it has the highest reservation value, and (iii) its mean is higher than the second highest reservation value. The first property shows that the agent takes the box with the highest backup value amongst remaining boxes: if there is a box with a higher mean but lower backup value, then it cannot be optimal to stop. The second property illustrates that when the agent stops, and takes a box without inspection, the tension between searching and stopping is "maximal": the candidate box to take without inspection is also the best box with which to continue search. The third property follows from (1): if (iii) does not hold, inspecting the box with the second highest reservation value, and choosing whatever is best between the sampled prize, and taking the chosen box without inspection, dominates taking the chosen box without inspection right away.

Proposition 2. *Let $(\mathcal{U}, z)$ denote the set of boxes, and the vector of realized prizes. Assume $\overline{z} < \max_{i\in\mathcal{U}} x_i^B$. If $\varphi^*(\mathcal{U}, z) = 0$, i.e., if search stops, then:*

1. $\arg\max_{i\in\mathcal{U}} x_i^B = \arg\max_{i\in\mathcal{U}} x_i^R \subseteq \arg\max_{i\in\mathcal{U}} \mu_i$,
2. *The sets* $\arg\max_{i\in\mathcal{U}} x_i^B$ *and* $\arg\max_{i\in\mathcal{U}} x_i^R$ *are singletons,*
3. $\max_{i\in\mathcal{U}} \mu_i > \max_{j\in\mathcal{U}\setminus\{\arg\max_{i\in\mathcal{U}} x_i^R\}} x_j^R$

Our next result, Proposition 3, shows that there are two reasons why, given a set of boxes $\mathcal{U}$ and a maximum sampled prize $\overline{z}$, 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, for some x_j such that $\overline{z}$ and x_j are below x_l^R, he might either (i) take box l without inspection in $\mathcal{U}\setminus\{j\}$, or (ii) continue search in $\mathcal{U}\setminus\{j\}$, but deviate yet again from Weitzman's order.

To understand (i), consider the school choice example in Section 1. There, the agent inspects school B (box j) first, which is the one with the lowest reservation value. 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 $\mathcal{U} = \{1, 2, 3\}$, let $X_i = \{0, 10\}$ be the set of prizes, and $p_i = P(X_i = 10)$. Assume $p_1 = 0.8 > p_2 =$

$0.65 > p_3 = 0.6$, $k_1 = 1 > k_2 = k_3 = 0.85$, and $\overline{z} = 0$. It can be checked that $x_1^R > x_2^R > x_3^R, x_1^B > x_2^B > x_3^B$. The optimal policy inspects box 2 first; if $x_2 = 10$ search stops, while if $x_2 = 0$ box 3 is inspected. If $x_3 = 10$ the agent takes $x_3 = 10$, and if $x_3 = 0$, box 1 is taken without inspection. In this example, $l = 1, j = 2$, and when $x_j = 0$, the agent continues search (inspects box 3), but deviates once more from Weitzman's order ($x_1^R > x_3^R$), since in the last stage box 1 is taken without inspection.[4]

Proposition 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 for all x_j such that $\max\{x_j, \overline{z}\} \leq x_l^R$, then $\varphi^*(\mathcal{U} \backslash \{j\}, z \cup \{x_j\}) = 1$ and $\sigma^*(\mathcal{U} \backslash \{j\}, z \cup \{x_j\}) = l$.*

4 Optimal Policy: Order and Stopping

This section introduces three sufficient conditions under which the optimal order coincides with Weitzman's. However, the optimal stopping rule is different, so we also characterize it.

Our first result, Theorem 1, requires that, given any two boxes i, j, $x_j^R \leq x_i^R$ if and only if $x_i^B \leq x_j^B$. Proposition 2 implies that then, since the box with the highest reservation value is always the box with the lowest backup value, when there is more than one box to be inspected, stopping and taking a box without inspection is never optimal. Hence, Weitzman's stopping rule applies to all but the last box. 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. Theorem 1 states the result formally, and Corollary 1 provides conditions under which the conditions in Theorem 1 are satisfied.[5]

[4] The example also shows that Theorem 2 does not extend when boxes have different inspection costs; in particular, when boxes with higher p_i have higher inspection costs. This is also true of other search models where conditions have to be imposed to derive the optimal search policy (see Chade and Smith [3], Vishwanath [11] and Klabjan, Olszewski, and Wolinsky [5])

[5] An analogue of Theorem 1 is discussed in Postl [10] 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 conditions on the primitives of the model under which Theorem 1 holds.

Theorem 1. *Fix a set $\mathcal{N} = \{1, ..., N\}$ of boxes. Assume that boxes can be labelled so that $[x_i^B, x_i^R]$ 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.*

Corollary 1 shows conditions on the primitives such that the ordering of the cutoffs is the one in Theorem 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 Theorem 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 boxes are not good backups since they can also contain worse draws. As discussed in Section 1.1, the same assumptions as in Corollary 1 are used in Vishwanath's [11] to obtain the reservation value rule in her parallel search model, and in the working paper version of Chade and Smith [3] 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 enough for Theorem 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 $\phi : \mathbb{R} \mapsto \mathbb{R}_+$, $\int \phi(x)dF_i(x) \leq \int \phi(x)dF_{i'}(x)$.[6]

[6]Mean preserving spreads, or the convex-order as it is defined in Ganuza and Penalva [4], and Li and Shi [7], requires the condition to hold for all convex functions.

Theorem 1 is not enough to characterize the optimal sampling policy in every environment: we need to consider the case in which both $x_j^R < x_i^R$ and $x_j^B < x_i^B$. When $x_j^B < x_i^B$, inspecting box i 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, and equal inspection costs (Theorem 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 (Theorem 3). It is interesting to note that similar conditions have been used before in search models: Chade and Smith [3] use binary prizes in their simultaneous search model, while Klabjan, Olszewski and Wolinsky [5] 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 $x_j^B \leq x_i^B$. The theorems 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.

Theorem 2. *Fix a set $\mathcal{N} = \{1, ..., N\}$ of boxes. Assume that boxes have binary prizes, i.e. $x_i \in \{y, x\}$, $x > \max\{x_0, y\}$, $p_i = P(x_i = x)$, and all boxes have the same inspection cost. Label boxes so that $p_N < ... < p_1$, and define recursively from N to 1:*

$$\begin{aligned} v_N &= \max\{x_0, x_N^B\} \\ v_n &= \max\{x_n^B, p_{n+1} \max\{x_{n+1}^R, x_0\} + (1 - p_{n+1}) v_{n+1}\} \end{aligned}$$

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 *Assume boxes $\{1, ...n - 1\}$ have been inspected. Search stops when: (i) the maximum sampled prize is above x_n^R, or (ii) the maximum sampled prize is y, and $x_n^B = v_n$. If (i), take the maximum sampled prize; if (ii) take box n without inspection.*

Before explaining the different elements of the policy in Theorem 2, it is useful to understand why Weitzman's order still holds under these conditions. Since all boxes contain either x or y, the agent wants to maximize the chance of obtaining a prize of x subject to inspecting the least number of boxes. Moreover, search stops as soon as the agent obtains x. Conditional on inspecting at least one box, by inspecting the box with the highest probability of obtaining x first, the agent satisfies these two objectives since search ends with a higher probability by starting with this box. However, this box is also the one with the highest backup value, so the agent may consider changing the inspection order to save this box to take later without inspection. However, changing the inspection order is not optimal: any other box he inspects has a higher chance of yielding a prize of y than the box with the highest backup value, which implies a higher probability of either stopping and taking the highest backup value box without inspection in the next round, or of continuing search. Therefore, the agent prefers to either stop right away, taking the box with the highest backup value without inspection, or to continue search with the high backup value box; what the optimal policy is depends on the continuation values. In particular, if the agent expects that, by continuing search, he does not earn more than the current highest backup value, then search stops.

Theorem 2 states that to determine the optimal stopping rule, we need to calculate new cutoffs, $\{v_n\}_{n=1}^N$. Note that these cutoffs are a function of the backup values, reservation values, and the initial outside option, and hence can be calculated together with these to determine the optimal policy. When the agent has inspected boxes $\{1, ...n-1\}$, he can either take box n without inspection, or inspect it. The first term of v_n represents the value of the first policy in the event that box n contains a prize of y.[7] The second term of v_n represents the value of inspecting box n in the event that it contains a prize of y: the agent earns the right to either inspect box $n+1$, or take it without inspection. When the agent decides whether to stop or not, he compares both terms in v_n, and stops whenever the continuation value after inspecting box n, v_n, is not bigger than x_n^B.

That $\{v_n\}_{n=1}^N$ can be computed using only the backup values, reservation val-

[7] Recall that the backup value is the lowest prize the agent expects to get from a box after internalizing he did not pay its inspections costs.

ues, and the initial outside option, is a consequence of the simple structure of the environment; in general, the continuation values could depend on more statistics of the distributions, and the details of the optimal policy.

Theorem 3. *Let $\mathcal{N} = \{1, 2\}$, and assume for all $i \in \mathcal{N}$ x_i is distributed according to F_i where F_i has pdf $\hat{f}_i(x) = f(x_i - \mu_i)$, and f 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 $x_2^B \geq 0$, search stops when x_1 is either (i) less than x_2^B, or (ii) higher than x_2^R. If (i), take box 2 without inspection; if (ii) take x_1. If $x_2^B < 0$, search stops only if (ii). Search starts if, and only if, $\max\{\mu_1, x_0\}$ is less than the payoff of inspecting box 1 first. In particular, if $\mu_1 \leq x_2^R$, search starts.*[8]

Theorem 3 follows because when boxes have symmetric distributions and equal inspection costs, the (unconditional) expected value of the prizes of each box above the reservation value coincides with the negative of the (unconditional) expected value of the prizes of each box below the backup value. When the agent compares the benefits and the costs of starting with box 1 (recall we are assuming $x_1^R > x_2^R$, and hence $x_1^B > x_2^B$), he compares the upper tails of boxes 1 and 2 with their lower tails: box 1 has a fatter upper tail, and hence is better for search; box 2 has a fatter lower tail, and hence box 1 is better to take without inspection. Given the above property, the costs and benefits exactly offset each other when $x_0 < x_2^B$, and hence box 2 is taken without inspection after starting with box 1. When $x_0 > x_2^B$, the benefit outweighs the cost, since in that case x_0 is a better buffer than the lower tail of box 2, as captured by x_2^B, when prizes in both boxes are too low.

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

Example 1. 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

[8] See Section 5 for a detailed discussion of why, in general, it is difficult to give a more precise statement for when it is optimal to start search.

$F_1 >_{FOSD} F_2$. Assume that $k_1 = k_2 = 1$. It can be shown that $x_1^B = \frac{14}{3} > x_2^B = 2.8$, 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 a 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 Theorems 2 and 3. It also shows why Corollary 1 cannot be relaxed to second-order stochastic dominance.

5 Two boxes

To further the understanding of the difficulties involved when characterizing the optimal policy when the conditions of Section 4 do not hold, this section characterizes the optimal policy when there are two boxes. Hence, for the rest of the section, $\mathcal{N} = \{1, 2\}$, and the outside option is given by $\overline{z}$.

Given that Proposition 0 characterizes the optimal continuation when there is one box left for inspection, we only need to determine which of the following three alternatives yields the highest payoff to characterize the optimal policy for two-boxes: (i) stop, taking $\max\{\overline{z}, \mu_1, \mu_2\}$, (ii) inspect box 1 first, and apply the optimal policy in Proposition 0 to box 2, and (iii) inspect box 2 first, and apply the optimal policy in Proposition 0 to box 1. Let Π_1 denote the payoff of (ii), and

Π_2 denote the payoff of (iii).[9]

Proposition 4 below describes the optimal policy when $\mathcal{N} = \{1, 2\}$:

Proposition 4. *Fix a set of boxes $\mathcal{N} = \{1, 2\}$, and let $\bar{z}$ denote the outside option. Assume without loss of generality that $x_2^R < x_1^R$. The following is the optimal policy:*

1. *If $\bar{z} > x_1^B$ and $\bar{z} > x_2^B$, then the optimal policy is given by Weitzman's rule.*
2. *If $x_1^B < x_2^B$, then it is optimal to inspect box 1 first. The optimal continuation policy is given by Proposition 0.*
3. *If $x_2^B < x_1^B$, $\bar{z} < x_1^B$, and $\mu_1 \leq x_2^R$, it is optimal to inspect at least one box. If $\Pi_1 > \Pi_2$, box 1 is inspected first; otherwise, box 2 is inspected first. In both cases, the optimal continuation is as in Proposition 0.*
4. *Otherwise, if $x_2^B < x_1^B$, $\bar{z} < x_1^B$, and $x_2^R < \mu_1$, it is optimal to inspect box 1 first if $\Pi_1 > \max\{\Pi_2, \mu_1\}$, to inspect box 2 first if $\Pi_2 > \max\{\Pi_1, \mu_1\}$; otherwise, box 1 is taken without inspection. If search does not stop, the optimal continuation policy is as in Proposition 0.*

Item 1 follows from Proposition 1, and item 2 follows from Theorem 1. When $x_2^R < x_1^R$ and $x_2^B < x_1^B$, Proposition 2 allows us to simplify the taxonomy by considering two cases: $\mu_1 \leq x_2^R$ and $x_2^R < \mu_1$. In the first case (item 3), the agent only has to decide which box to inspect next, i.e. the optimal policy is determined by $\max\{\Pi_1, \Pi_2\}$. In the second case (item 4), the agent has to choose either to stop, taking box 1 without inspection, or which box to inspect next.

To gain intuition about what may determine which option the agent chooses when $x_2^R < x_1^R$ and $x_2^B < x_1^B$, we analyze the differences $\Pi_1 - \Pi_2$, $\Pi_2 - \mu_1$, and $\Pi_1 - \mu_1$. The first determines the optimal policy in item 3, and all three determine the optimal policy in item 4.[10]

[9] Π_i is the payoff from inspecting box i first, and: (i) if $\max\{x_i, \bar{z}\} > x_j^R$ stop, and take $\max\{x_i, \bar{z}\}$, (ii) if $\max\{x_i, \bar{z}\} \in [x_j^B, x_j^R]$ inspect box j, and take $\max\{x_i, x_j, \bar{z}\}$, (iii) if $\max\{x_i, \bar{z}\} < x_j^B$ stop, and take μ_j

[10] The supplementary material contains a more complete analysis of these differences. In particular, we define three cutoffs (x_O, x_1^S, x_2^S) that must be compared with x_1^B and x_2^R to determine the optimal policy. These cutoffs are not necessarily easier to compute than the actual payoffs of the different strategies; they only provide a formal definition of conditions on parameters of

We start with $\Pi_1 - \Pi_2$. It is immediate, if somewhat tedious, to show that it is given by:[11]

$$\begin{aligned} \Pi_1 - \Pi_2 &= \int_{x_2^R}^{+\infty} \int_{x_2^R}^{+\infty} (\min\{x_1^R, x_2, x_1\} - x_2^R) dF_2 dF_1 \\ &+ \int_{-\infty}^{x_1^B} \int_{-\infty}^{x_1^B} (\max\{x_1, x_2, \max\{x_2^B, \overline{z}\}\} - x_1^B) dF_2 dF_1 \end{aligned} \tag{5}$$

Recall that we are assuming that $x_1^R > x_2^R$, and $x_1^B > x_2^B$, so that the first term in (5) is positive, and the second is negative. Equation (5) shows that inspecting first box 1 has a benefit, which is given by the possibility of obtaining higher prizes, net of inspection costs, and a cost, which is given by the possibility of obtaining really low prizes, in which case keeping box 1 to take without inspection would act as a buffer. A somewhat loose intuition is that the higher the backup value of box 1, or the higher the reservation value of box 2, the higher the cost of inspecting box 1 first, and hence the optimal policy would start with box 2.[12]

Equation (5) alone determines the optimal policy when $x_2^R < x_1^R, x_2^B < x_1^B, \mu_1 \leq x_2^R$. When $\mu_1 > x_2^R$, by Proposition 2, the agent may find it optimal to stop, and take box 1 without inspection. Hence, we also need to compare Π_1 to μ_1, and Π_2 to μ_1.

Consider first the choice of whether to inspect box 2 first, or take box 1 without inspection. It is immediate that if $x_2^R > \mu_1 (> x_1^B > \overline{z})$, then stopping cannot be optimal: inspecting box 2 and then taking box 1 without inspection whenever $x_2 < \mu_1$ certainly dominates stopping and taking box 1 without inspection. It is also immediate that if $x_2^R < x_1^B$, then stopping dominates inspecting box 2 first: x_2^R is the maximum prize the agent expects to get from box 2 after inspection, while x_1^B is the lowest prize the agent expects to get from box 1 when taking it without inspection. To sharpen this intuition, note that the difference $\Pi_2 - \mu_1$ is

the environment for which the different strategies are optimal.

[11] Equations (5)-(7) are derived in Appendix S.2 for completeness.

[12] The intuition is loose because some changes in x_1^B (x_2^R) may change also x_1^R (x_2^B).

given by:

$$\Pi_2 - \mu_1 = -k_2 + \int_{x_1^B}^{+\infty} \int_{-\infty}^{+\infty} \left(\begin{array}{c} \max\{x_2, \min\{x_1^R, \max\{x_1, x_2\}\}\} \\ -\min\{x_1^R, \max\{x_1, x_1^B\}\} \end{array} \right) dF_1 dF_2 \quad (6)$$

When $x_2 < x_1^B$, box 1 is taken without inspection, after inspecting box 2, and this determines the integration limits in the outer integral in (6). Recall from equation (4) that when taking box 1 without inspection, the agent expects to gain no more than x_1^R, and no less than x_1^B, and this determines the second term in the integrand. The first term is the gain from inspecting box 2 first, followed by inspecting box 1: by not taking box 1 without inspection, the agent gets the possibility of getting the prize inside box 2, though this comes at the cost of paying k_2.

Equation (6) resembles the equation that determines the reservation value for box 2, but where now the outside option is μ_1. As the previous intuition suggests, as long as it is worth inspecting box 2 (i.e., x_2^R is high compared to x_1^B), the above expression should favor inspecting at least one box.

Finally, we need to compare Π_1 and μ_1. The difference $\Pi_1 - \mu_1$ can be written as:

$$\Pi_1 - \mu_1 = \int_{-\infty}^{x_2^R} \int_{-\infty}^{+\infty} \min\{x_2^R, \max\{x_1, x_2, \max\{x_2^B, \overline{z}\}\}\} - \max\{x_1, x_1^B\} dF_2 dF_1 \quad (7)$$

The difference between Π_1 and μ_1 is that by inspecting box 1 first, the agent retains the option of inspecting box 2 (the first term in the integrand), while he loses the option to take box 1 without inspection (the second term in the integrand). The equation resembles the computation of the backup value of box 1, but with an inspection cost of 0. When the agent inspects box 1 first, he gives up the backup value of box 1; hence, if box 2 is sufficiently good for search, the possibility of searching with box 2 may compensate for this. This, in turn, favors inspecting at least one box over stopping, and taking box 1 without inspection.

Equations (5)-(7) and the discussion above show that, even in the case $N = 2$, it is not always simple to determine the optimal policy by just looking at the boxes' cutoff values. This, in turn, highlights the value of the conditions in Section 4, which allow us to characterize the optimal policy by only looking at these cutoffs, and thus retain tractability which is useful for applications.

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 identify sufficient conditions under which the optimal policy involves following Weitzman's inspection order, and characterize the optimal stopping rule in those cases. These conditions have been used elsewhere in the search literature to simplify other problems, and thereby obtain otherwise unavailable characterizations of the optimal policy. Moreover, by retaining the simplicity of the reservation value rule, the conditions help deliver results which are useful for applications. Section 1 already discussed one application of interest. Two other applications of particular interest where our results could be applied to are: (i) the choice amongst technologies with which to produce a good when the agent can invest in pre-project planning to find out the true production cost, but has the option to produce without making this investment (Krähmer and Strausz [6] consider a one-technology version of this problem), and (ii) the allocation of a good to one of several agents when the principal can find out which agent would generate the highest payoff from obtaining the good, but can allocate it without further investigation, as in Ben-Porath, Dekel and Lipman [1].

Finally, we also provide properties of the optimal policy that must hold across all environments (Propositions 1-3), and illustrated in Section 5 how they can be used to reduce the taxonomy when the sufficient conditions identified in Section 4 do not hold.

References

[1] E. Ben-Porath, E. Dekel, and B. L. Lipman. Optimal allocation with costly verification. Mar. 2013.

[2] H. Chade and N. Kovrijnykh. Information acquisition, moral hazard, and rewarding for bad news. *Working Paper*, 2011.

[3] H. Chade and L. Smith. Simultaneous search. *Econometrica*, 74(5):1293–1307, 2006.

[4] 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.

[5] D. Klabjan, W. Olszewski, and A. Wolinsky. Attributes. *Mimeo*, 2011.

[6] D. Krähmer and R. Strausz. Optimal procurement contracts with pre-project planning. *The Review of Economic Studies*, 78(3):1015–1041, 2011.

[7] H. Li and X. Shi. Discriminatory information disclosure. *Working paper*, 2013.

[8] W. Olszewski. A more general pandora rule. *Mimeo*, 2012.

[9] C. Papadimitriou and J. Tsitsiklis. The complexity of optimal queuing network control. *Mathematics of Operations Research*, pages 293–305, 1999.

[10] P. Postl. Delegated search: Procedure matters. *Working Paper*, 2004.

[11] T. Vishwanath. Parallel search for the best alternative. *Economic Theory*, 2(4):495–507, 1992.

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

A Appendix

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

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

Proposition 1. Let $\mathcal{U}$ be a set of boxes, and let z be the vector of realized prizes. If $(\forall i \in \mathcal{U}) : \overline{z} \geq x_i^B$, 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) $(\forall \mathcal{U}) : (|\mathcal{U}| = U), (\forall z) : (\overline{z} \geq \max_{i \in \mathcal{U}} x_i^B)$, the order and stopping policy indicated in Proposition 1 is optimal.

Step 1: $P(1) = 1$ This follows from Proposition 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 Proposition 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:

$\overline{z} \geq x_l^R$ By contradiction, suppose that it is optimal to continue search, and box $j \in \mathcal{U}$ is inspected. Since $|\mathcal{U} \backslash \{j\}| = U$, and $\max_{i \in \mathcal{U} \backslash \{j\}} x_i^B \leq x_l^R < \max\{x_j, \overline{z}\}$, then by the inductive hypothesis search stops. Thus, the payoff of continuing search with box $j \in \mathcal{U}$ is: $-k_j + \int \max\{x_j, \overline{z}\} dF_j(x_j) < \overline{z}$. The last inequality follows from equation (1), and $x_j^R < \overline{z}$ for all $j \in \mathcal{U}$.

$\overline{z} < x_l^R$ If $\max\{\overline{z}, \max_{i \in \mathcal{U}} \mu_i\} \neq \mu_l$, then, by equation (1), inspecting box l and stopping dominates stopping and obtaining payoff $\max\{\overline{z}, \max_{i \in \mathcal{U}} \mu_i\}$, since $\max\{\overline{z}, \max_{i \in \mathcal{U}} \mu_i\} < x_l^R$. If $\max\{\overline{z}, \max_{i \in \mathcal{U}} \mu_i\} = \mu_l$, since $\overline{z} \geq x_l^B$, we have that $\max\{\overline{z}, \max_{i \in \mathcal{U} \backslash \{l\}} \mu_i\} \geq x_l^B$, and hence, by equation (2), inspecting box l and stopping dominates stopping and taking box l without inspection.

Finally, when $\overline{z} < x_l^R$, we need to show that inspecting box l first is optimal. Let $j \in \mathcal{U} \backslash \{l\}$ be any other box. Note that $x_j^R < x_l^R$. Consider the following two policies:

P.J Inspect 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.L Inspect box l first. If $x_l^R \leq x_l$, stop. Otherwise, inspect box j, and stop, or continue search according to the rule described in Proposition 1.

Let $h = \arg\max_{i \in \mathcal{U} \setminus \{l\}} x_i^R$. The payoff from policies **P.J** and **P.L** can be written as:

$$\begin{aligned}
\textbf{P.J} &= -k_j + \int_{x_l^R}^{+\infty} x_j dF_j + \int_{x_h^R}^{x_l^R} (-k_l + \int \max\{x_l, x_j, \overline{z}\} dF_l) dF_j \\
&+ \int_{-\infty}^{x_h^R} (-k_l + \int_{x_h^R}^{+\infty} x_l dF_l + \int_{-\infty}^{x_h^R} V^*(\mathcal{U} \backslash \{l, j\}, \overline{z} \cup \{x_j, x_l\}) dF_l) dF_j
\end{aligned}$$

$$\begin{aligned}
\textbf{P.L} &= -k_l + \int_{x_l^R}^{+\infty} x_l dF_l + \int_{x_h^R}^{x_l^R} (-k_j + \int \max\{x_l, x_j, \overline{z}\} dF_j) dF_l \\
&+ \int_{-\infty}^{x_h^R} (-k_j + \int_{x_h^R}^{+\infty} x_j dF_j + \int_{-\infty}^{x_h^R} V^*(\mathcal{U} \backslash \{l, j\}, \overline{z} \cup \{x_j, x_l\}) dF_j) dF_l
\end{aligned}$$

The difference in payoffs between both policies is given by:

$$\begin{aligned}
\text{P.L} - \text{P.J} &= (1 - F_j(x_l^R))[\int_{x_l^R}^{+\infty} x_l dF_l - k_l] - (1 - F_l(x_l^R))[\int_{x_l^R}^{+\infty} x_j dF_j - k_j] \\
&= (1 - F_l(x_l^R))(1 - F_j(x_l^R))(x_l^R - x_j^R) + \int_{x_j^R}^{x_l^R} (x_j - x_j^R) dF_j \geq 0
\end{aligned}$$

where the second equality follows from equation (1) for boxes l, and j respectively. Thus, inspecting box l dominates inspecting any other box $j \in \mathcal{U} \backslash \{l\}$. This completes our proof. □

Proposition 2. Let $(\mathcal{U}, z)$ denote the set of boxes, and the vector of realized prizes. Assume $\overline{z} < \max_{i \in \mathcal{U}} x_i^B$. If $\varphi^*(\mathcal{U}, z) = 0$, i.e., if search stops, then:

1. $\arg\max_{i \in \mathcal{U}} x_i^B = \arg\max_{i \in \mathcal{U}} x_i^R \subseteq \arg\max_{i \in \mathcal{U}} \mu_i$,

2. The sets $\arg\max_{i \in \mathcal{U}} x_i^B$ and $\arg\max_{i \in \mathcal{U}} x_i^R$ are singletons,

3. $\max_{i \in \mathcal{U}} \mu_i > \max_{j \in \mathcal{U} \backslash \{\arg\max_{i \in \mathcal{U}} x_i^R\}} x_j^R$

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 $\arg\max_{i\in\mathcal{U}} x_i^B \subseteq \arg\max_{i\in\mathcal{U}} \mu_i$. We do so by contradiction. Assume that $(\exists j,k \in \mathcal{U})(j \neq k) : \mu_j < \mu_k = \max_{\in\mathcal{U}} \mu_i; x_k^B < x_j^B = \max_{i\in\mathcal{U}} x_i^B$. Note that $\overline{z} < x_j^B$ by assumption. We show that inspecting box k first, and then applying the policy in Proposition 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 Proposition 0 to inspect/take without inspection box j, his payoff is:

$$\Pi_k = -k_k + \int_{x_j^R}^{+\infty} x_k dF_k + \int_{x_j^B}^{x_j^R} (-k_j + \int \max\{x_j, x_k\} dF_j) dF_k + \int_{-\infty}^{x_j^B} \mu_j dF_k$$

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

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

Therefore, we can write:

$$\begin{aligned}
& \Pi_k - \mu_k = \\
- & \int_{x_k^B}^{x_j^R} x_k dF_k + \int_{x_j^B}^{x_j^R} (-k_j + \int \max\{x_j, x_k\} dF_j) dF_k + \int_{-\infty}^{x_k^B} (\mu_j - x_k^B) dF_k \\
+ & \int_{x_k^B}^{x_j^B} \mu_j dF_k \\
= & \int_{x_j^B}^{x_j^R} (\int_{x_j^R}^{+\infty} x_j^R dF_j + \int_{-\infty}^{x_j^R} \max\{x_j, x_k\} dF_j - x_k) dF_k + \int_{x_k^B}^{x_j^B} (\mu_j - x_k) dF_k \\
+ & \int_{-\infty}^{x_k^B} (\mu_j - x_k^B) dF_k
\end{aligned}$$

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 $\mu_j > x_j^B > x_k$ by assumption 1, and (iii) the third integrand is positive because $\mu_j > x_j^B > x_k^B$. This shows that $\arg\max_{i\in\mathcal{U}} x_i^B \subseteq \arg\max_{i\in\mathcal{U}} \mu_i$.

Now suppose that $\arg\max_{i\in\mathcal{U}} x_i^R \not\subseteq \arg\max_{i\in\mathcal{U}} x_i^B$. Then, there exists $(\exists j, l \in$

$\mathcal{U})(j \neq l) : x_j^R < x_l^R = \max_{\in \mathcal{U}} x_i^R; x_l^B < x_j^B = \max_{i \in \mathcal{U}} x_i^B$. Note that $\overline{z} < x_j^B < x_j^R < x_l^R$, and hence $\mu_j < x_j^R < x_l^R$. Consider the following policy: inspect box l first; if $x_l > \mu_j$, stop and take x_l, if $x_l \leq \mu_j$, stop and take box j without inspection. Applying equation (1) we know that:

$$-k_l + \int_{\mu_j}^{+\infty} x_l dF_l + \int_{-\infty}^{\mu_j} \mu_j dF_l > \mu_j$$

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

We now show that $\arg\max_{i \in \mathcal{U}} x_i^R, \arg\max_{i \in \mathcal{U}} x_i^B$ are singletons. Suppose not. Then $(\exists l, l' \in \mathcal{U}) : x_l^R = x_{l'}^R, x_l^B = x_{l'}^B$. 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 Proposition 0 for inspecting box l'. This improves upon stopping and taking box l', because:

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

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 $\mu_{l'} = \int_{x_{l'}^R}^{+\infty} x_l^R + \int_{x_{l'}^B}^{x_{l'}^R} x_{l'} dF_{l'} + \int_{-\infty}^{x_{l'}^B} x_{l'}^B dF_{l'}$ (note that the inequality is an equality only when $x_l^R = \mu_l = x_l^B$).

Therefore, we conclude that $\arg\max_{i \in \mathcal{U}} x_i^R, \arg\max_{i \in \mathcal{U}} x_i^B$ are both singletons.

Finally, letting box l denote the box with the highest reservation value, and hence one of the boxes with the highest mean, suppose that $\mu_l < \max_{i \in \mathcal{U} \setminus \{l\}} x_i^R$, and let box h denote the box with the second highest reservation value. Consider the following policy: inspect box h first; if $x_h > \mu_l$, stop and take x_h, and if $x_h \leq \mu_l$, take box l without inspection. Equation (1), and the definition of x_h^R imply that:

$$-k_h + \int_{\mu_l}^{+\infty} x_h dF_h + \int_{-\infty}^{\mu_l} \mu_l dF_h > \mu_l$$

This contradicts the optimality of stopping search. Thus, $x_h^R < \mu_l$. □

Proposition 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 $(\forall x_j) : \max\{x_j, \overline{z}\} \leq x_l^R$, $\sigma^*(\mathcal{U}\backslash\{j\}, z \cup \{x_j\}) = l$, and $\varphi^*(\mathcal{U}\backslash\{j\}, z \cup \{x_j\}) = 1$.

Proof. Suppose $\sigma^*(\cdot) = \{j\}$ and the optimal continuation policy dictates inspecting box l whenever $\max\{x_j, \overline{z}\} \leq x_l^R$. The following policy improves on this, as shown by the proof of 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 $\mathcal{U} = \mathcal{U}\backslash\{l, j\}$. □

A.2 Proof of Theorem 1 and Corollary 1

Theorem 1. *Fix a set $\mathcal{U}$ of boxes, and let z be the vector of previously realized prizes. Assume that boxes are labelled so that $[x_i^B, x_i^R]$ forms a monotone sequence in the set inclusion order, that is $i < i'$ implies $[x_i^B, x_i^R] \subset [x_{i'}^B, x_{i'}^R]$. The following is the optimal policy:*

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

Stopping *If $|\mathcal{U}| > 1$, $\varphi^*(\mathcal{U}, z) = 0$ if and only if $\overline{z} > \arg\max_{i\in\mathcal{U}} x_i^R$. If $|\mathcal{U}| = 1$, $\varphi^*(\mathcal{U}, z) = 0$ if and only if (i) $\overline{z} > \arg\max_{i\in\mathcal{U}} x_i^R$, or (ii) $\overline{z} < \arg\max_{i\in\mathcal{U}} x_i^B$*

Proof. 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 Theorem 1, the order and stopping rules in Theorem 1 are optimal.

Proposition 0 shows that $P(1) = 1$. We show that Theorem 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:

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

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

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

$\overline{z} < x_2^R$ If $\max\{\overline{z}, \max_i \mu_i\} \neq \mu_2$, then inspecting box 2 alone, and stopping dominates stopping and obtaining payoff $\max\{\overline{z}, \max_{i\in\mathcal{U}} \mu_i\}$ by (1). If $\max\{\overline{z}, \max_i \mu_i\} = \mu_2$, since $\max\{\overline{z}, \mu_1\} > x_1^B \geq x_2^B$, 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 $\overline{z} < x_2^R$. This follows from equation (5).

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 $y = \max\{\overline{z}, \max_{i\in\mathcal{U}} \mu_i\}$ 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.[13] Since $\mu_{U+1} < x_{U+1}^R$, and $x_{U+1}^R < y$ we have that $x_{U+1}^R < \max\{\overline{z}, \max_{i\in\mathcal{U}\backslash\{U+1\}} \mu_i\}$. Thus, $(\forall x_{U+1}) \max\{\overline{z}, x_{U+1}, \max_{i\in\mathcal{U}\backslash\{U+1\}} \mu_i\} > x_{U+1}^R$. Since $|\mathcal{U}\backslash\{U+1\}| = U$, it is optimal to stop by the inductive hypothesis. Therefore, the payoff from inspecting box $U+1$ is:

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

which follows from (1), since $\overline{z} > x_{U+1}^R$ (see footnote 13). Therefore, it is optimal to stop.

2. Assume that $y \leq x_{U+1}^R$. Since, by assumption 1, $x_k^B < \mu_k$, and $(\forall k \in \mathcal{U}\backslash\{U+$

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

$1\})x^B_{U+1} \le x^B_k$, it must be the case that $y > x^B_{U+1}$, and $\max\{\bar{z}, \max_{i\in\mathcal{U}\setminus\{U+1\}} \mu_i\} > x^B_{U+1}$. Then, inspecting box $U+1$ and stopping dominates the payoff the agent obtains by stopping. If $\mu_{U+1} = y$, this follows from $\max\{\bar{z}, \max_{i\in\mathcal{U}\setminus\{U+1\}} \mu_i\} > x^B_{U+1}$, and equation(2); if $\mu_{U+1} < y$, this follows from 1. above.

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

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

P.U+1 Inspect box $U+1$ first. If $x^R_{U+1} \le x_{U+1}$, stop. Otherwise, inspect box j, and stop, or continue search according to the rule described in Theorem 1.

Note that when $U+1 = 3$, after inspecting boxes $\{U+1, j\}$, the continuation policy may imply that box U is taken without inspection (this would be the case if $\bar{z} < x^B_U$). This, a priori, may indicate that a different proof is needed when $U+1 = 3$, and when $U+1 > 3$. This is actually not the case since policies **P.J**, and **P.U+1** coincide whenever $(x_{U+1}, x_j) \in (-\infty, x^R_U]^2$, and, therefore, the difference in payoffs between the two policies does not depend on whether box U is taken without inspection or not. However, for completeness sake, we distinguish between the two cases. In particular, we present the proof for when $U+1 = 3$, and $\bar{z} < x^B_U$, and $U+1 > 3$. (When $U+1 = 3$, and $x^B_U \le \bar{z}$, box U is not taken without inspection and, hence, the continuation policy looks like the case in which $U+1 > 3$).

Suppose $U+1 = 3$, and $\bar{z} < x^B_U$. The payoffs from policies **P.J** and **P.U+1** can be written as:

$$\mathbf{P.J} = -k_j + \int_{x^R_{U+1}}^{+\infty} x_j dF_j + \int_{x^R_U}^{x^R_{U+1}} \left(-k_{U+1} + \int \max\{x_{U+1}, x_j\} dF_{U+1}\right) dF_j$$
$$+ \int_{x^B_U}^{x^R_U} (-k_{U+1} + \int_{x^R_U}^{+\infty} x_{U+1} dF_{U+1} + \int_{-\infty}^{x^R_U} (-k_U + \int_{-\infty}^{+\infty} \max\{x_U, x_{U+1}, x_j\} dF_U) dF_{U+1}) dF_j$$
$$+ \int_{-\infty}^{x^B_U} \begin{pmatrix} -k_{U+1} + \int_{x^R_U}^{+\infty} x_{U+1} dF_{U+1} \\ + \int_{x^B_U}^{x^R_U} (-k_U + \int_{-\infty}^{+\infty} \max\{x_U, x_{U+1}, x_j\} dF_U) dF_{U+1} + \int_{-\infty}^{x^B_U} \mu_U dF_{U+1} \end{pmatrix} dF_j$$

$$\mathbf{P.U+1} = -k_{U+1} + \int_{x^R_{U+1}}^{+\infty} x_{U+1} dF_{U+1} + \int_{x^R_U}^{x^R_{U+1}} \left(-k_j + \int \max\{x_{U+1}, x_j\} dF_j\right) dF_{U+1}$$

$$+ \int_{x^B_U}^{x^R_U} (-k_j + \int_{x^R_U}^{+\infty} x_j dF_j + \int_{-\infty}^{x^R_U} (-k_U + \int_{-\infty}^{+\infty} \max\{x_U, x_{U+1}, x_j\} dF_U) dF_j) dF_{U+1}$$

$$+ \int_{-\infty}^{x^B_U} \left(\begin{array}{c} -k_j + \int_{x^R_U}^{+\infty} x_j dF_j \\ + \int_{x^B_U}^{x^R_U} (-k_U + \int_{-\infty}^{+\infty} \max\{x_U, x_{U+1}, x_j\} dF_U) dF_j + \int_{-\infty}^{x^B_U} \mu_U dF_j \end{array} \right) dF_{U+1}$$

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

$$\begin{aligned} \mathbf{P.U+1\text{-}P.J} &= (1 - F_j(x^R_{U+1})) \left[\int_{x^R_{U+1}}^{+\infty} x_{U+1} dF_{U+1} - k_{U+1} \right] \\ &- (1 - F_{U+1}(x^R_{U+1})) \left[\int_{x^R_{U+1}}^{+\infty} x_j dF_j - k_j \right] \end{aligned} \tag{A.1}$$

Using the definition of the reservation value for boxes $U+1$ and j, equation (A.1) can be written as:

$$\mathbf{P.U+1\text{-}P.J} = (1 - F_j(x^R_{U+1}))(1 - F_{U+1}(x^R_{U+1}))(x^R_{U+1} - x^R_j) + \int_{x^R_j}^{x^R_{U+1}} (x_j - x^R_j) dF_j$$

That **P(U+1)**=1 follows from $x^R_{U+1} > x^R_j$.

Consider now the case in which $U+1 > 3$ (or, $x^B_U \leq \overline{z}$). Let $\Phi(x_{U+1}, x_j) = \mathbb{E}[V^*(\mathcal{U} \backslash \{U+1, j\}, z \cup \{x_{U+1}\} \cup \{x_j\})]$ to be the expected (continuation) payoff the agent obtains by applying the rule in Theorem 1 when the set of boxes is $\mathcal{U} \backslash \{U+1, j\}$, and the vector of realized prizes is $z \cup \{x_{U+1}\} \cup \{x_j\}$. Since $|\mathcal{U} \backslash \{U+1, j\}| < U+1$, and the boxes in $\mathcal{U} \backslash \{U+1, j\}$ can be enumerated as in Theorem 1, the policy in Theorem 1 is optimal when applied to that set. Consider again the payoffs obtained from following policies **P.J** and **P.U+1**:

$$\mathbf{P.J} = -k_j + \int_{x^R_{U+1}}^{+\infty} x_j dF_j + \int_{x^R_U}^{x^R_{U+1}} \left(-k_{U+1} + \int \max\{x_{U+1}, x_j\} dF_{U+1}\right) dF_j$$

$$+ \int_{-\infty}^{x^R_U} (-k_{U+1} + \int_{x^R_U}^{+\infty} x_{U+1} dF_{U+1} + \int_{-\infty}^{x^R_U} \Phi(x_{U+1}, x_j) dF_{U+1}) dF_j$$

$$\mathbf{P.U{+}1} = -k_{U+1} + \int_{x^R_{U+1}}^{+\infty} x_{U+1} dF_{U+1} + \int_{x^R_U}^{x^R_{U+1}} \left(-k_j + \int \max\{x_{U+1}, x_j\} dF_j \right) dF_{U+1}$$

$$+ \int_{-\infty}^{x^R_U} (-k_j + \int_{x^R_U}^{+\infty} x_j dF_j + \int_{-\infty}^{x^R_U} \Phi(x_{U+1}, x_j) dF_j) dF_{U+1}$$

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

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 Theorem 1.

Proof. It suffices to show that if $i < i'$, then $[x^B_i, x^R_i] \subseteq [x^B_{i'}, x^R_{i'}]$. To see this, rewrite equation (1) for box i as:

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

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

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

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

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

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

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

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

Therefore, we conclude that $[x^B_i, x^R_i] \subseteq [x^B_{i'}, x^R_{i'}]$. □

A.3 Proof of Theorem 2

The proof is divided in two parts. First, we show the following modified version of Theorem 2:

Theorem A.1. *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* $k_i \equiv k$. *Assume that boxes are labeled so that* $p_{|\mathcal{U}|} < ... < p_1$. *The following is the optimal policy:*

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

Stopping $\varphi^*(\mathcal{U}, z) = 0$ *if and only if* $\overline{z} > x^R_{\sigma^*(\mathcal{U},z)}$, *or* $\overline{z} = \max\{x_0, y\} < x^B_{\sigma^*(\mathcal{U},z)}$ *and* $V^*(\mathcal{U} \backslash \{\sigma^*(\mathcal{U}, z)\}, z \cup \{y\}) < x^B_{\sigma^*(\mathcal{U},z)}$

Second, we show that $V^*(\mathcal{U} \backslash \{\sigma^*(\mathcal{U}, \overline{z})\}, \overline{z} \cup y)$ coincides with the cutoffs $\{v_n\}$ defined in Theorem 2. We start with the proof of Theorem A.1.

Proof. 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 Theorem A.1, the order and stopping rules in Theorem A.1 are optimal.

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

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

Let $\mathcal{U}$ be such that $|\mathcal{U}| = U + 1$, and let 1 be the box with the highest p_i. We first show that when $\overline{z} \geq x^R_1$ it is optimal to stop. Note that $\overline{z} > \mu_1 = \max_{i \in \mathcal{U}} \mu_i$, hence, if the agent stops he selects $\overline{z}$ as payoff. Moreover, if the agent inspects any box $k \in \mathcal{U}$, $|\mathcal{U} \backslash \{k\}| = U$, 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 $\overline{z} < x^R_1$. We first show that, if a box is to be opened first, it has to be box 1. Let j be any other box $j \neq l$. Let $V^*(\mathcal{U} \backslash \{1, j\}, z \cup \{x_1\} \cup \{x_j\})$ denote the value function in the continuation problem after inspecting boxes $1, j$.[14]

[14] That is, the payoff the agent obtains by following the policy stated in Theorem 2, which is optimal by the inductive step since $|\mathcal{U} \backslash \{1, j\}| < U + 1$

In a slight abuse of notation, define $V^*(\mathcal{U}, \max\{x_0, y\}) \equiv V^*(\mathcal{U}, z)$ whenever $\overline{z} < x$. If $V^*(\mathcal{U}\backslash\{1, j\}, \max\{x_0, y\}) \geq x_1^B$ consider the following two policies:

P.1 Open box 1 first. If $x_1 = x$, 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.1** is:

$$p_1 x + (1 - p_1)[p_j x + (1 - p_j)V^*(\mathcal{U}\backslash\{1, j\}, \max\{x_0, y\}) - k] - k$$

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

$$p_j x + (1 - p_j)[p_1 x + (1 - p_1)V^*(\mathcal{U}\backslash\{1, j\}, \max\{x_0, y\}) - k] - k$$

The comparison of the payoffs yields the result, by noticing that it reduces to the case in which there are only two boxes $1, j$. When $V^*(\mathcal{U}\backslash\{1, j\}, \max\{x_0, y\}) < x_1^B$, compare the following two policies:

P.1 Open box l first. If $x_1 = x$, 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.1] is $p_1 x + (1 - p_1)(p_j x + (1 - p_j)y) - k$, and that of policy [P.J] is $p_j x + (1 - p_j)(p_1 x + (1 - p_1)y) - k$. The difference is null. Then, opening box 1 first dominates (weakly) opening any other box $j \neq 1$. Now, to verify the rest of the stopping rule, note that if $\overline{z} = \max\{x_0, y\}$ and $\overline{z} < x_1^B$,

$$\begin{aligned} p_1 x + (1 - p_1)V^*(\mathcal{U}\backslash\{1\}, \max\{y, 0\}) - k &\geq p_1 x + (1 - p_1)y \\ \Leftrightarrow V^*(\mathcal{U}\backslash\{1\}, \max\{y, 0\}) &\geq y + \frac{k}{1 - p_1} = x_1^B \end{aligned}$$

□

Now, we prove the following:

Claim A.1 (Stopping rule for Theorem 2)**.** Let $\mathcal{U}$ be a set of boxes as in the assumptions of Theorem 2, and let x_0 be the initial outside option. Then, for

$n \leq |\mathcal{U}| - 1$:

$$V^*(\mathcal{U}\backslash\{1,...,n\}, \max\{x_0, y\}) = p_{n+1}\max\{x_n^R, x_0\} + (1-p_{n+1})v_{n+1}$$
$$v_{|\mathcal{U}|} = \max\{x_0, x_{\mathcal{U}}^B\}$$

Proof. We prove it for $n = |\mathcal{U}| - 1$, and then extend it inductively for $n < |\mathcal{U}| - 1$. Consider then $n = |\mathcal{U}| - 1$. Want to show that $V^*(\mathcal{U}\backslash\{1...|\mathcal{U}|-1\}, \max\{x_0, y\}) = p_{|\mathcal{U}|}\max\{x_0, x^R_{|\mathcal{U}}|\} + (1-p_{|\mathcal{U}|})\max\{x_0, x^B_{|\mathcal{U}|}\}$. We have the following cases:

1. If $x_0 > x^R_{|\mathcal{U}|}$, then box $|\mathcal{U}|$ is not inspected by Theorem A.1. Hence, $V^*(\mathcal{U}\backslash\{1...|\mathcal{U}|-1\}, \max\{x_0, y\}) = x_0$.

2. $x^R_{|\mathcal{U}|} \geq x_0 \geq x^B_{|\mathcal{U}|}$. Then, box $|\mathcal{U}|$ is inspected, and if the prize is y, the agent takes the outside option. Thus $V^*(\mathcal{U}\backslash\{1...|\mathcal{U}|-1\}, \max\{x_0, y\}) = -k + p_{|\mathcal{U}|}x + (1-p_{|\mathcal{U}|})x_0 = p_{|\mathcal{U}|}(x - \frac{k}{p_{|\mathcal{U}|}}) + (1-p_{|\mathcal{U}|})x_0 = p_{|\mathcal{U}|}x^R_{|\mathcal{U}|} + (1-p_{|\mathcal{U}|})x_0$.

3. If $x^B_{|\mathcal{U}|} > x_0$, box $|\mathcal{U}|$ is taken without inspection, and $V^*(\mathcal{U}\backslash\{1...|\mathcal{U}|-1\}, \max\{x_0, y\}) = p_{|\mathcal{U}|}x + (1-p_{|\mathcal{U}|})y = -k + p_{|\mathcal{U}|}x + (1-p_{|\mathcal{U}|})y + k = p_{|\mathcal{U}|}(x - \frac{k}{p_{|\mathcal{U}|}}) + (1-p_{|\mathcal{U}|})(y + \frac{k}{1-p_{|\mathcal{U}|}}) = p_{|\mathcal{U}|}x^R_{|\mathcal{U}|} + (1-p_{|\mathcal{U}|})x^B_{|\mathcal{U}|}$.

The three cases complete the proof that $V^*(\mathcal{U}\backslash\{1...|\mathcal{U}|-1\}, \max\{x_0, y\}) = p_{|\mathcal{U}|}\max\{x_0, x^R_{|\mathcal{U}|}\} + (1-p_{|\mathcal{U}|})\max\{x_0, x^B_{|\mathcal{U}|}\}$.

Suppose the claim is true for all $n' > n$, and we show that $V^*(\mathcal{U}\backslash\{1...n\}, \max\{x_0, y\}) = p_{n+1}\max\{x_0, x^R_{n+1}\} + (1-p_{n+1})v_{n+1}$. Consider the following cases:

1. $x_0 > x^R_{n+1}$. Then, box $n+1$ is not inspected, and $V^*(\mathcal{U}\backslash\{1...n\}, \max\{x_0, y\}) = x_0$ (note that no more boxes are inspected, and indeed $v_{n'} = x_0$ for all $n' > n$).

2. $x^R_{n+1} > x_0 > x^B_{n+1}$. Then, $v_{n+1} \geq x_0$, and hence box $n+1$ is not taken without inspection. Thus, $V^*(\mathcal{U}\backslash\{1...n\}, \max\{x_0, y\}) = -k + p_{n+1}x + (1-p_{n+1})v_{n+1} = p_{n+1}x^R_{n+1} + (1-p_n)v_{n+1}$.

3. $x^B_{n+1} > x_0$, and $v_{n+1} > x^B_{n+1}$. Then, box $n+1$ is inspected, and the continuation, after a prize of y, is v_{n+1}. The same as above obtains.

4. $x^B_{n+1} \geq \max\{x_0, v_{n+1}\}$. Then, the agent stops and takes box $n+1$ without inspection. Thus, $V^*(\mathcal{U}\backslash\{1...n\}, \max\{x_0, y\}) = p_{n+1}x + (1-p_{n+1})y = p_{n+1}x^R_{n+1} + (1-p_{n+1})x^B_{n+1} = p_{n+1}x^R_{n+1} + (1-p_{n+1})v_{n+1}$.

The four steps complete the proof. □

A.4 Proof of Theorem 3

We first establish a preliminary result on the cutoff values when the conditions in Theorem 3 hold:

Lemma A.1 (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, $(\exists \underline{b}, \overline{b}) : x^B = \mu - \underline{b}, x^R = \overline{b}$.

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 $x^R = \mu + \overline{b}$, for some $\overline{b} > 0$.

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

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

$$k = \int_{\overline{b}}^{+\infty} (u - \overline{b}) dF(u) \tag{A.2}$$

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

Corollary A.1. Consider the same assumptions as before. If F is symmetric around 0 then $\overline{b} = \underline{b} = b > 0$

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:

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

Now, symmetry of F implies that:

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

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

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

$$\begin{aligned} \Pi_i - \Pi_j &= \int_{-\infty}^{x_i^B} \int_{-\infty}^{x_i^B} (\max\{x_i, x_j, x_j^B\} - x_i^B) dF_i dF_j \\ &+ \int_{x_j^R}^{+\infty} \int_{x_j^R}^{+\infty} (\min\{x_i, x_j, x_i^R\} - x_j^R) dF_i dF_j \end{aligned}$$

$$
\begin{aligned}
&= (1-F_i(x_i^R))(1-F_j(x_i^R))(x_i^R - x_j^R) + \int_{x_j^R}^{x_i^R}\int_{x_i}^{+\infty}(x_i - x_j^R)dF_j dF_i \\
&+ \int_{x_j^R}^{x_i^R}\int_{x_j^R}^{x_i}(x_j - x_j^R)dF_j dF_i + (1-F_i(x_i^R))\int_{x_j^R}^{x_i^R}(x_j - x_j^R)dF_j \\
&+ F_i(x_j^B)F_j(x_j^B)(x_j^B - x_i^B) + F_i(x_j^B)\int_{x_j^B}^{x_i^B}(x_j - x_i^B)dF_j \\
&+ \int_{x_j^B}^{x_i^B}\int_{x_i}^{x_i^B}(x_j - x_i^B)dF_j dF_i + \int_{x_j^B}^{x_i^B}\int_{-\infty}^{x_i}(x_i - x_i^B)dF_j dF_i
\end{aligned}
$$

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 $G(a) \equiv 0$.

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

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

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

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 $x_j^B \leq 0 \leq x_i^B$, we have that:

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

Since the previous proof never used the fact that $x_j^B \geq 0$, and $x_j^B < 0$ in this case, this shows that the previous difference is positive. Finally, when $x_j^B \leq x_i^B \leq 0$, the problem is exactly as Weitzman's, hence we know that the difference is strictly positive. This completes the proof.

A.5 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 $x_i^R < x_i^B$, 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 $x_i^R < x_i^B$, our analysis carries through by focusing on the boxes for which $x_i^B < x_i^R$.

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

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

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

Proposition A.1. *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 A.1.

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:

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

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

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

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

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

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 $\mathcal{U}' =$

$\mathcal{U}\backslash\{j\}, \mathcal{U}'^{R<B} = \mathcal{U}^{R<B}$, and $|\mathcal{U}'| = U - 1$, and $\mathcal{U}'^{R<B} = n$ (note that if there was $j \in \mathcal{U}\backslash\mathcal{U}^{R<B}$, 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 $\mathcal{U}' = \mathcal{U}\backslash\{i\}$, $\mathcal{U}'^{R<B} = \mathcal{U}^{R<B}\backslash\{i\}$, 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:

$$\begin{aligned} &\max\left\{\overline{z}, \max_{i\in\mathcal{U}^{R<B}} \mu_i, \max_{j\in\mathcal{U}^{B<R}} \mu_j, \max_{j\in\mathcal{U}^{B<R}} \{-k_j + \int V^*(\mathcal{U}\backslash\{j\}, z \cup \{x_j\}) dF_j\}\right\} \\ \geq\ &\max_{i\in\mathcal{U}^{R<B}} \{-k_i + \int V^*(\mathcal{U}\backslash\{i\}, z \cup \{x_i\}) dF_i\} \end{aligned} \tag{A.3}$$

where the LHS of the above expression denotes the payoff the agent can get by either stopping, and getting $\max\{\overline{z}, \max_{i\in\mathcal{U}^{R<B}} \mu_i, \max_{j\in\mathcal{U}^{B<R}} \mu_j\}$, 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 $i \in \mathcal{U}^{R<B}$:

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

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:

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

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:

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

Also, we have that:

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

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:

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

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

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

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 (A.3) holds, and, thus, $P(U, n) = 1$

□

B Examples

B.1 Optimal policy in the example of Section 1

We now prove that the policies described for the example in Section 1 (see Table 1) are indeed optimal.

Claim B.1. Suppose Weitzman's assumption holds. Then, the following is the optimal policy: school A is 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.

Proof. The parameters of the model imply that $x_A^R = 4$, and $x_B^R = \frac{5}{2}$ (this can be shown by applying equation (1)). The result follows immediately from applying Weitzman's [12] rule. □

Claim B.2. Suppose Weitzman's assumption does not hold. Then, the following is the optimal policy: school B is visited first; if the prize is $x_B = 0$, search stops and school A is accepted without inspection; if $x_B = 3$, the agent visits school A, and chooses the school with the highest realized prize.

Proof. 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 Proposition 0, conditional on $x_B = 3$, visiting school A and selecting the best school is optimal since $x_A^B = 1 < 3 < x_A^R = 4$. Third, if the agent visits school A first, the optimal continuation policy is the same as when the student cannot accept a school without attending the visit day: this follows from Proposition 0, and the fact that, for all x_A, $x_B^B = \frac{1}{2} < x_A$. 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, comparing the payoffs from visiting school A first, and from visiting school B first, we conclude that visiting school B first is optimal. □

B.2 Example footnote 1 in Section 1

Section 1 contains a two-box example that illustrates that Weitzman's policy may no longer be optimal when the agent can take a box without inspection conditional on stopping. The example is, in part, driven (and, made simple) by the fact that the worst realization in box A is better than the worst realization in box B. Thus, after bad news from box B, it does not pay to inspect box A. The following example illustrates that this is not in fact necessary to obtain such a result. In particular, the example features two boxes A and B such that box B's prize distribution has a higher mean, both boxes coincide in their worst realization, box B has a lower reservation and backup value, and the optimal policy consists in inspecting box B first.

Example B.1. Table B.1 describes the prize distribution, and inspection costs of boxes A and B It can be verified that $x_A^R = 4 > x_B^R = 3.9$, $x_A^B = 1 > x_B^B = \frac{1}{2}$,

A	Prize	0	1	5	Inspection cost
	Probability	0.10	0.80	0.10	0.10
B	Prize	0	0.5	4.3	Inspection cost
	Probability	0.2	0.55	0.25	0.10
Z	Prize		0		Inspection cost
	Probability		1		0

Table 2: Prize distribution for each box

and $\mu_2 = 1.35 > \mu_1 = 1.3$. Thus, in Weitzman's model, the agent inspects box A first; if $x_A = 5$, search stops, and, if $x_A < 5$, the agent inspects box B and keeps the highest realized prize.

If we relax Weitzman's assumption, by Proposition 0, after inspecting box A, the agent inspects box B only when $x_A = 1$; if $x_A = 5$, search stops and the agent keeps the prize in box A, and when $x_A = 0$ he takes box B without inspection. If, instead, he starts with box B, he never inspects box A: if $x_B = 4.3$, he stops search, and takes the prize, while if $x_B \in \{0, 0.5\}$, he takes box A without inspection. That is, he takes box A without inspection when $x_B \leq \frac{1}{2}$ even if box A may contain a prize worse than $\frac{1}{2}$. However, box A assigns a very high probability to $x_A = 1$, which in turns makes box A's backup value higher than box B's. Thus, the combined effect of saving on inspection costs when box B has a low enough prize and the "certainty" of a not so low prize from box A induce the agent to inspect box B first. In fact, the following is the optimal policy:

Claim B.3. The agent inspects box B first. If $x_B = 4.3$, search stops, and he takes $x_B = 4.3$. If $x_B < 4.3$, the agent takes box A without inspection.

Remark 2. One can modify the above example so that $x_B = 0$ with probability $\frac{25}{48}$, $x_B = 1.2$ with probability $\frac{11}{48}$, and $x_B = 4.3$ with probability 0.25, and still obtain the result that despite box B having a higher mean value than box A, the reversal in the search order occurs all the same.[15] The only difference is that now box A is only taken without inspection when $x_B = 0$, and in that case the worst prize in box A is the same as the agent's current outside option.

B.3 Cutoffs don't determine the optimal policy if $N \geq 2$

The following example demonstrates the claim made in Section 1:

Example B.2. Suppose $N = 2$. Box 1 is as in Example 1: $X_1 = \{0, 2, 10\}$, $P(X_1 = 2) = 0.2, P(X_1 = 10) = 0.7$, and $k_1 = 1$. Box 2 is such that $X_2 = \{0, 9\}$, $P(X_2 = 9) = 0.7368$, and $k_2 = \frac{14}{9}$. It is immediate to show that cutoffs are exactly the same as the ones of the boxes in Example 1. However, the optimal policy now inspects box 1 first; search stops if $X_1 = 10$, and the agent gets $X_1 = 10$, while box 2 is taken without inspection when $X_1 \leq 2$.

[15]One can verify that the change in the distribution of box B left the mean, reservation and backup values of box B unaltered.