Optimal Search with Costly Returns*

Sebastian Ertner†

VERY PRELIMINARY AND INCOMPLETE

February 21, 2026

Abstract

I study an extension of the sequential search problem of Weitzman (1979), where the agent has access to a second, alternative search technology to learn a box's value. Compared to the 'classic' technology of Weitzman (1979), this alternative has lower inspection costs but an additional return cost if the box is not ultimately chosen. An example of this setting is online shopping, where a product may be inspected more easily by having it delivered, but then a return cost is incurred if it is not kept. I characterize general features of the optimal search policy and fully determine the optimal policy under sufficient conditions. I find that the alternative technology becomes more valuable relative to the classic technology as search progresses.

JEL-Classification: D83

Keywords: Sequential search, Information acquisition, Product returns

1 Introduction

Consider a consumer that wants to buy a new phone. The consumer may go to a brick and mortar store and test a certain model. Alternatively, the consumer may simply order that phone online and try it out at home once it arrives. In the second case, the consumer has an easier time inspecting the phone to find out if they like it, but if they do not want to keep it, they face return costs.

Now, consider a firm that wants to hire a worker for a specific job. The firm can do multiple rounds of interviews and assessments to judge if an applicant fits the position. Alternatively,

*I thank Maarten Janssen and Daniel Garcia for their advice and support and Jacopo Gambato and audiences in Vienna for helpful discussions and suggestions. I acknowledge financial support from the Austrian Science Foundation under project number FG 6.

†Vienna Graduate School of Economics (VGSE), University of Vienna. E-mail: sebastian.ertner@univie.ac.at.

the firm could hire that applicant on the spot and see how they perform on the job. In the second case, the firm saves on upfront costs, but if the applicant turns out not to be a good fit for the position, the firm faces costs similar to the consumer who returns an ordered phone.

There are many[1] similar scenarios in which an agent searches for the best alternative out of finitely many and faces this kind of trade-off: For each alternative, they can either make an upfront investment to be fully informed before selecting that alternative, or they could select the alternative, learn its value and face the risk of paying return costs if they then want to back out again and choose a different alternative. To keep the language general I will from now on refer to an alternative as a box that contains a value.

In this paper, I identify properties of the optimal search strategy that an agent should employ in this kind of setting. I extend the classic model of Weitzman (1979) by allowing the agent to not only inspect a box before taking it, but alternatively to take the box right away, inspect it at a lower inspection cost, at which point they however have to pay a return cost should they ultimately decide not to keep the box. I will refer to the first type of action as 'pre-inspection' and to the second as 'post-inspection', in reference to the point in time where the box is taken.

My main result is that in the optimal search strategy, for a given search order the agent becomes weakly more likely to post-inspect boxes the later they are in the search process. This is the case because from an ex-post perspective, it is better to post-inspect a box if it is also ultimately chosen. This becomes more likely later in the process, when there are less boxes remaining that can still be inspected and which could potentially reveal a larger value. While over the course of the search process the agent becomes more likely to post-inspect boxes, the best observed value of all opened boxes also weakly increases. This can result in the agent going 'back and forth' between pre- and post-inspecting boxes, depending on the realizations these inspections yield.

The result is formulated for a given search order, covering a variety of settings. In many search models, boxes are assumed to be homogeneous, such that the search order does not matter and my main result applies. More generally, in many applications that allow for heterogeneous boxes, the focus of the analysis is on the agent's decision when to stop, while the order of inspection is pre-determined, for example by equilibrium considerations. Development of further concrete sufficient conditions is work in progress.

In terms of the initial example of the consumer shopping for a phone, my main result implies that the consumer may start out investing effort upfront to learn the value of a certain

[1] I discuss more examples in section 2.4.

phone. At this point, there are many different models available such that the consumer expects to still inspect more of them and is not willing to risk paying the return cost. However, over the course of the search less and less models remain uninspected and the consumer becomes more willing to take the risk of ordering a phone, making inspection easier while the risk of returning the phone has become smaller due to the limited options that remain available.

The optimal search policy cannot be characterized easily in the general case, but it is clear in specific cases. In the case with an infinite number of ex-ante homogeneous boxes[2] the optimal search strategy is such that the agent should inspect all boxes in the same way, either pre-inspect all or post-inspect all. While at first this seems to contradict the main result, it is consistent with it. Note that the main result states that pre-inspection becomes *weakly* better the earlier a box is inspected. In the special case where the agent strictly prefers to post-inspect the last box and boxes are homogeneous, the value of pre-inspection in fact stays constant for all earlier boxes, such that all boxes should be post-inspected. In all other cases, the value of pre-inspection strictly increases for earlier boxes. Thus if infinitely many boxes may be inspected afterwards, then the current box should be pre-inspected. This is however true for all boxes, such that all should be pre-inspected in this case.

Alternatively, I consider cases where it is not feasible to pre-inspect boxes, for example when it is prohibitively difficult to obtain information about a box before having access to it.[3] In such cases the optimal search strategy can be identified as that of Weitzman (1979) applied to a transformed version of the model. Further results can be derived, such as an eventual purchase theorem, as in Armstrong (2017) and Choi, Dai, and Kim (2018). A possible extension in this case is to account for limited information the agent may obtain already before taking a box, while full information may still only be revealed by post-inspection. An example of this is that an online shopper may find some information about a product by browsing online, while they fully learn whether they like the product only once they have ordered and inspected it. This can be achieved by adding a 'browsing'[4] stage before the post-inspection stage using the branching-bandit framework of Keller and Oldale (2003), with the derivation of an eventual purchase theorem being possible also in this configuration.

To derive these results I utilize a transformed version of the post-inspection action. It is insightful to discuss this transformation briefly as it gives intuition for when the agent wants to post-inspect a box. The idea is that the agent assumes when they take a box for post-

[2] An example for this case are consumer search models with a unit mass of firms that offer a homogeneous product.

[3] Similar results can be obtained when there are some boxes that may only be pre-inspected and others that may only be post-inspected.

[4] The terminology is borrowed from Heidhues, Johnen, and Kőszegi (2021).

inspection that they will later return it. They thus account for the payment of the return cost already at the time of inspection. Should the agent later however decide to keep the box, then they get back that return cost they considered lost. If the agent returns the box, then no payment has to be made as they already accounted for this event. In this transformed view, the return cost therefore does not affect the decision to return the box, but it instead affects the decisions whether to inspect and whether to keep the box. In the first instance, the return cost acts as a kind of additional inspection cost, making it more costly to inspect the box. In the second instance, the return cost increases the value found in the box as it gets paid back again if the box is ultimately chosen. This has the effect that the agent may pick this box over another one that nominally has a larger value. This provides us now with the intuition for when post-inspection is better: It results in a larger 'effective' inspection cost than pre-inspection, in exchange for an equally increased value of the box. The agent should thus choose this costly way of increasing a box's value only when it is likely that they will ultimately select that box.

This paper relates to two strands of the search literature. The first are papers that identify the optimal search policy in generalized settings. Here, the closest paper conceptually is Doval (2018), who extends Weitzman (1979) by alternatively allowing the agent to pick a box without inspecting it, thereby ending search. The main difference to this paper arises from the fact that in this paper, taking the alternative action of post-inspecting the box does not end search. The implication is that the agent may choose this alternative inspection for multiple boxes and thus more than once within a search. The model of Doval (2018) is a special case of the model in this paper, obtained in the limit where the cost of post-inspecting goes to zero while the return cost goes to infinity.

The second strand of literature considers search with product returns, where consumers search optimally in a setting where products can be returned once ordered and inspected. Matthews and Persico (2007) have consumers who can pay a cost to become informed about their value of a monopolist's product, or can stay uninformed, in which case they learn their value after buying the product and can after that return it at a cost. Consumers in Janssen and Williams (2024) may similarly inspect a product before or after purchasing it, with the option of returning the product in the latter case. They additionally allow their consumers to buy a product without inspecting it (as Doval (2018) does) and consumers have to pay a search cost to learn the price of a product before they can inspect or buy it. Ertner and Janssen (2025) study how a multi-(two-)product monopolist offering ex-ante homogeneous goods to a consumer may incentivise consumers to inspect the products after ordering. My approach of transforming post-inspection in the present paper is equivalent to the one employed in that paper. In these three papers the return cost is indirectly chosen by the firm(s) through the refund it provides to consumers who return a product. Petrikaitė (2018) studies a duopoly of horizontally differentiated single-product firms, where the consumer must purchase a product to learn its value, at which point they may purchase the product of the other firm as well and return the one they like less at a fixed return cost that is the same for either product. In that paper the consumer cannot pre-inspect products. What all these papers have in common and what differentiates this paper from them is that they do not derive the optimal policy for a general setting with heterogeneous options.[5] Further and importantly, in these papers the number of products/boxes among which the consumer searches is either one, two or infinitely many. The main result of this paper only emerges for a finite number of at least two boxes. While Ertner and Janssen (2025) with two products that may be inspected before or after ordering is in this regard the closest to this paper, that result is not observed there as the threshold where the consumer is indifferent between inspecting the first product before or after ordering is not derived.

There is further a related literature on search with correlated products and branching bandits including Keller and Oldale (2003), Greminger (2021) and Gambato (2024).[6] The difference these papers is that most of them consider correlation in early stages of the search process, but not in terminal nodes. This allows them to still derive an index policy. In this paper, there is correlation in the terminal nodes such that an index policy cannot be derived.

The remainder of this paper is organized as follows. Section 2 introduces the model and preliminary results. Section 3 presents the main results. Section 4 concludes with a discussion.

2 Model and preliminary results

There is an initial set $\mathcal{I} = \{1, \dots, n\}$ of uninspected boxes. Each box $i \in \mathcal{I}$ has an initially unknown value x_i that is distributed according to cdf F_i on the support $[\underline{x}_i, \overline{x}_i]$. To learn x_i, the agent has to inspect box i, which they can do in two alternative ways. The agent can pay an inspection cost c_i^{pre} to look inside the box and observe the realization of x_i. This I will call "inspecting the box before taking it", or in short 'pre-inspection'. This way of inspecting a box is the same as in Weitzman (1979). Alternatively, the agent may take the box and inspect it only after that, in which case the cost c_i^{post} to learn x_i is smaller, $c_i^{post} < c_i^{pre}$, but the agent has to pay an additional return cost $\sigma_i > 0$ if they don't ultimately want to keep

[5] While some of these papers may allow heterogeneity in terms of the price a firm charges, none has products that ex-ante differ in search/inspection cost or value distribution.

[6] See also related work in progress by Anderson, Engers, and Savelle (2021) and Zhang (2026).

the box. I will call this way of inspecting a box "inspecting the box after taking it" or in short 'post-inspection'. The pre and post prefixes thus are in reference to the point in time where a box is 'taken', but may still be returned. Note however, that when a box has been pre-inspected, then no new information can be learned about the box thereafter and thus a pre-inspected box will never be returned once it has been taken. To clearly differentiate between the action of taking a box, where it can still be returned, and the action of ultimately selecting a box and ending search, I will only refer to the former as 'taking the box'.[7]

In addition to the boxes in $\mathcal{I}$ there is an initial outside option with value x_0.[8] The agent can at any time instead of inspecting a new box pick one of the already inspected boxes or the initial outside option, received its value and thereby end search. After the agent has inspected an arbitrary number of boxes, denote by $N \subset \mathcal{I}$ the set of remaining uninspected boxes and by $\overline{N} := \mathcal{I} \setminus N$ the set of already inspected boxes. Assume without loss of generality that all boxes that were post-inspected and are not kept are returned only when the agent ends search. To simplify notation I will sometimes denote a box i that has been pre-inspected as having a return cost of $\sigma_i = 0$.

2.1 Transforming post-inspection

There is an alternative interpretation of the 'post-inspection' mode, which facilitates the derivation and interpretation of results and which therefore I will employ for the remainder of the paper. In this alternative interpretation, when the agent decides to post-inspect box i, they already assume that they will later return the box and pay the return cost σ_i. Thus the new cost of 'inspecting' box i in this way is the sum of its real inspection cost c_i^{post} and its return cost σ_i. If the agent later ends search and does not keep the box, then no additional payment has to be made as this event has already been accounted for. Should the agent however later decide to keep box i, then they 'get back' σ_i together with the value x_i in the box. We thus take the perspective that the agent treats returning the box as the default option. To summarize: When post-inspection is transformed in this way, it has a total inspection cost of $c_i^{post} + \sigma_i$, no return cost, and the agent gets $x_i + \sigma_i$ if they ultimately pick box i.

[7] The phrase 'taking the box' is chosen deliberately, as it associates the actions of 'taking' and 'returning' the box with physical acts, such that it is natural that returning the box can entail explicit costs. This underlines the fact that the agent does not just choose the box in their mind, but that by taking the box they commit to paying the return cost in case they do not keep the box.

[8] The initial outside option may also have a return cost $\sigma_0 > 0$. The only difference in that case is that the outside option has to be transformed like post-inspected boxes, such that its value becomes $x_0 + \sigma_0$ and we have to subtract σ_0 from the value function of the transformed problem – see the next section and the proof of Lemma 1. This does not affect the optimal policy or the results in any way.

This transformation has two advantages. First, we do not have to keep track of return costs of post-inspected boxes. This makes solving the model easier, as both modes of inspection are now similar in structure by having only an inspection cost and a value distribution. Second, this transformation adjusts the value of ultimately selecting the box in such a way that it directly reveals which already inspected box is currently the 'best', i.e. which one should be picked if the agent wanted to end search right then. Consider again the original representation of the model and the case where some boxes have already been inspected, i.e. $\overline{N}$ is non-empty. Then each $i \in \overline{N}$ has a certain realized value x_i and a return cost $\sigma_i \geq 0$.[9] Which is currently the best out of those already inspected boxes? When box i is selected and search ends, the agent receives x_i and pays the return costs σ_j of all other boxes $j \in \overline{N} \backslash \{i\}$. We can rewrite this utility of selecting i as follows: $x_i - \sum_{j \in \overline{N} \backslash \{i\}} \sigma_j = x_i + \sigma_i - \sum_{j \in \overline{N}} \sigma_j$. On the right, the sum of the return costs of all inspected boxes is identical for all i, so that the currently best box is the one with largest $x_i + \sigma_i$. This is clear, as by deciding to inspect them, the agent has already committed to paying the return cost for all of the post-inspected boxes except one. Thus to determine the best box, we only have to consider the sum of each box's individual value x_i and its return cost σ_i, as the agent can only prevent paying the latter by ultimately selecting i. Returning to the transformed view, we see that there the value of box i is identified as $x_i + \sigma_i$, correctly ranking the boxes by the value they provide when selected. Note also that in the transformed view, σ_i has already been paid at inspection for any box $i \in \overline{N}$, such that in this view $x_i + \sigma_i$ is in fact the full value of selecting box i to end search.

The following lemma proves that we can indeed make such a transformation, i.e. that the search problem with boxes that have a return cost can be transformed into one where boxes have no return costs, but adjusted inspection costs and value distributions.

Lemma 1. *The agent's optimal search policy can be found by solving the transformed problem.*

To prove the lemma I again make use of the fact that not only can post-inspection be transformed into pre-inspection, but the reverse is possible as well. This follows simply from the fact that a box that is pre-inspected can be viewed as having a return cost of $\sigma_i = 0$.

This transformation is mathematically identical to the transformation employed in Ertner and Janssen (2025), but there is a minor difference in terms of interpretation. In that paper, after the agent has ordered a product to post-inspect it, they may either buy the product, paying its retail price to the firm, or return the product and pay a return cost (e.g. a restocking fee) to the firm, which is smaller than the retail price. In that sense, by ordering and inspecting the product the agent commits to paying at least the return cost to the firm. Here in this

[9] If the box was pre-inspected, then $\sigma_i = 0$.

paper, the agent instead commits to paying the return cost only when they return the box.

We see that in the transformed view the return cost acts like an additional inspection cost, a fact that also Petrikaitė (2018) notes in her model without actual inspection costs. The return cost also acts as a positive shift of the value distribution of the box, which has the effect that the agent may choose this box instead of another one (or the outside option) that has a larger actual value. This fact does not emerge in her model, as there the outside option is so small that the consumer always buys one of the two products and both products also have the same return cost.

2.2 The agent's problem

Consider the case where a certain number of boxes has already been inspected, such that $\overline{N}$ is non-empty. We have concluded in the previous section that the best element in $\overline{N}$ is the one with largest $x_i + \sigma_i$. It is clear then that out of all the boxes in $\overline{N}$, the agent only considers choosing that box and disregards all other boxes.[10] Thus we define the so far best realized outcome z as the maximum of the initial outside option and the best opened box, $z := \max(\{x_0\} \cup \{x_i + \sigma_i\}_{i \in \overline{N}})$. Then at any point in the search process the state is fully defined by the best realized outcome z and the set of remaining uninspected boxes N. The value function of the agent can then be written down as follows:

$$\begin{aligned} V(N,z) = \max \Bigg\{ & z, \max_{i \in N} \left[-c_i^{pre} + F_i(z) V(N \backslash \{i\}, z) + \int_z^{\overline{x}_i} V(N \backslash \{i\}, x_i) f_i(x_i) dx_i \right], \\ & \max_{j \in N} \left[-c_j^{post} - \sigma_i + F_j(z - \sigma_j) V(N \backslash \{j\}, z) + \int_{z-\sigma_j}^{\overline{x}_j} V(N \backslash \{j\}, x_j + \sigma_j) f_j(x_j) dx_j \right] \Bigg\} \end{aligned}$$

The available actions of the agent in any state are thus to either take z and end search, or to pre- or post-inspect another uninspected box and transition to another state. When the agent pre-inspects box i, then they pay the inspection cost c_i^{pre} and transition to the state with the remaining boxes and where the best realized outcome is either still z, or becomes x_i, depending on which is bigger. The same happens when the agent post-inspects box j, except in addition to the inspection cost c_j^{post} the agent pays σ_j upfront and then compares z to $x_j + \sigma_j$ for the best realized outcome in the new state. Note that the element in N that maximizes the value function when the next element is pre-inspected does not have to be the same as when it is post-inspected, i.e. in the above value function i can be different from j. For later we define $V_i^{pre}(N,z)$ and $V_i^{post}(N,z)$ as the respective value functions of pre-/post-

[10] There may be multiple 'best' boxes with the same $x_i + \sigma_i$, in which case it does not matter which of those the agent chooses. Also, as before, if i was pre-inspected, then $\sigma_i = 0$.

inspecting box $i \in N$ in state (N, z), such that we can write the agent's value function as follows:[11]

$$V(N, z) = \max\left\{z, \max_{i\in N} V_i^{pre}(N, z), \max_{j\in N} V_j^{post}(N, z)\right\} \tag{1}$$

Throughout the paper I will make comparisons between $V_i^{pre}(N, z)$ and $V_i^{post}(N, z)$. It is important to note that in such comparisons, we ask the question 'if box i would be inspected next, would it be better to pre- or post-inspect it (or possibly not at all)?'. However, as just established, the decision which box to optimally inspect next can differ depending on whether it will be pre- or post-inspected.

For certain derivations and illustrations it is helpful to consider instead of the total value in state (N, z) the additional value that can be gained over the best realized outcome z. I denote this additional value over z by $\Delta V(N, z)$, defined as $\Delta V(N, z) := V(N, z) - z$. Then equivalently define $\Delta V_i^{pre}(N, z) := V_i^{pre}(N, z) - z$ and $\Delta V_i^{post}(N, z) := V_i^{post}(N, z) - z$.

2.3 The one box problem

As typical in search models, certain characteristic values can be defined for each box $i \in \mathcal{I}$, that determine which action the agent should take in state (N, z). The purpose of this section is to introduce these characteristic values in the simple case where there is only one box available, i.e. when $(\{i\}, z)$. I introduce the general versions for (N, z) in section 3.

For the inspection actions, we can define the well-known reservation values – the largest value of z for which the agent is willing to take the respective inspection action. We denote by r_i^{pre} the reservation value of pre-inspecting box i – it is the value of z at which the agent is indifferent between pre-inspecting i and not inspecting it and taking z immediately, i.e. $\Delta V_i^{pre}(\{i\}, z) = 0$. Equivalently, r_i^{post} is the reservation value of post-inspecting box i – it is the value of z at which the agent is indifferent between post-inspecting i and not inspecting it and taking z immediately, i.e. $\Delta V_i^{post}(\{i\}, z) = 0$. The reservation values are thus implicitly defined as the solutions to these equations, which simplify to:

$$\int_{r_i^{pre}} [1 - F_i(x_i)]dx_i = c_i^{pre} \qquad \text{and} \qquad \int_{r_i^{post}-\sigma_i} [1 - F_i(x_i)]dx_i = c_i^{post} + \sigma_i$$

We now consider how the relation between the reservation values affects how the agent should inspect box i. First, note that by taking the derivatives of V_i^{pre} and V_i^{post} with regard to z we find that $V_i^{pre}(\{i\}, z)' \geq V_i^{post}(\{i\}, z)'$ for all z, i.e. the slope of the value function of

[11] Thus $V_i^{pre}(N, z) := -c_i^{pre} + F_i(z)V(N\backslash\{i\}, z) + \int_z^{\overline{x}_i} V(N\backslash\{i\}, x_i)f_i(x_i)dx_i$ and $V_i^{post}(N, z) := -c_i^{post} - \sigma_i + F_i(z - \sigma_i)V(N\backslash\{i\}, z) + \int_{z-\sigma_i}^{\overline{x}_i} V(N\backslash\{i\}, x_i + \sigma_i)f_i(x_i)dx_i$.

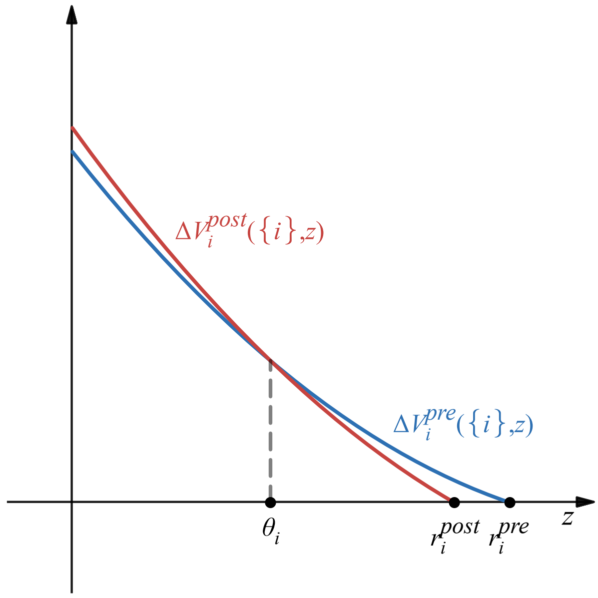

Figure 1: The characteristic values of box i as points of indifference between pre-/post-/no inspection when $r_i^{pre} > r_i^{post}$.

pre-inspection is always more positive than that of post-inspection. From this it follows that if $r_i^{pre} < r_i^{post}$, then $V_i^{pre}(\{i\}, z) \leq V_i^{post}(\{i\}, z)$ for all z. In other words, if the reservation value of post-inspection is larger, then the agent will never pre-inspect box i, regardless of the value of z.

If, on the other hand, the reservation value of pre-inspection is larger, V_i^{pre} and V_i^{post} may have a unique intersection point. In that case it is optimal to post-inspect i up to that threshold and to pre-inspect i above that and up to r_i^{pre}. Thus define the threshold θ_i as the value of z at which the agent is indifferent between pre- and post-inspecting, i.e. $V_i^{pre}(\{i\}, z) = V_i^{post}(\{i\}, z)$. The threshold θ_i is therefore implicitly defined by:

$$\int_{\theta_i - \sigma_i}^{\theta_i} F_i(x_i) dx_i = c_i^{pre} - c_i^{post} \tag{2}$$

The agent prefers to post-inspect i for $z \leq \theta_i$ and otherwise prefers to pre-inspect it. The threshold θ_i is well-defined in this way as long as neither inspection option is strictly better than the other. For the purposes of this section it is not necessary to go into these cases; I provide a more precise definition of θ_i in Proposition 1.

We can immediately make the following observation about the relation between the reservation values and the threshold. It is clear that the threshold value cannot lie between the reservation values. In that case the agent would be willing to undertake only the inspection action with the larger reservation value and not the other, but then it is impossible for the agent to be indifferent between the inspection actions. As we have argued that there is no point of indifference between the inspection modes when $r_i^{pre} < r_i^{post}$, the only possibility for θ_i to exist is if $\theta_i \leq r_i^{post} \leq r_i^{pre}$. Figure 1 shows the relation of the characteristic values in that case. It also demonstrates that $V_i^{pre}(\{i\}, z)' \geq V_i^{post}(\{i\}, z)'$.

2.4 Interpretation of the model

What the return cost σ_i represents differs by setting. It can correspond to costs connected to the return process, like hassle costs, costs of physically returning the object to the seller (postage), restocking fees or costs for terminating a contract. It can also correspond to a partial refund, where only a part of the initially paid price of the box is refunded (see e.g. Ertner and Janssen (2025)).

Another important consideration is when discarded boxes are returned. As noted above, I assume for simplicity and without loss of generality, that boxes are returned only at the end of search. However, this is not necessary. At any point during the search, the agent only considers the best realized outcome z out of the already inspected boxes. This implies that when the agent inspects a new box, they immediately observe if it is better than z and can return the worse of the two right away.

I now discuss some further examples that can be represented by this kind of model. When searching for a new flat, the agent may invest effort to learn everything about a certain flat before moving in, or they may save this effort and risk finding out later that they do not like the flat, at which point they would face a return cost. Here this cost may include the cost of moving again or the cost of terminating the contract. A related case is when the agent is searching to buy a house, where they can hire experts to assess the house. If they do not do that, the return cost here might represent the loss in value that may be incurred by having to resell the house and certainly psychological hassle costs. Similar to the example of the firm that wants to fill a position, we can consider the perspective of the applicant, who can exert effort to learn about the firm and position they are considering, for example by looking up experiences of former employees on websites that review companies. The motivating example in Doval (2018) also applies here: A student wants to searchers for a school to visit. The student can go to the visit day of a school to learn the fit. In Doval's model, if there is only one school left to consider, the student can save effort by not going to the visit day and simply choose that school. In my model, the student may similarly choose a school without going to the visit day, and if they find out they do not like it, they may transfer to a different school, again involving a cost that can include hassle costs or the opportunity costs of missed time in school.

3 Main results

With a finite set of boxes and the value function (1), it is a known result that the dynamic problem of the agent is solvable under regular assumptions and thus an optimal search policy must exist. The main question is whether there exists a simple characterization of that optimal policy, which in the case of search problems would usually take the form of an index policy.[12] However, we can immediately answer this question in the negative, as the correlation of pre- and post-inspection modes for the same box makes it impossible to derive indices for those actions.[13] Note that in particular the challenge is determining the correct inspection order and inspection mode; it is however clear that for determining when the agent should stop searching it is enough to compare z to the reservation values. While the optimal search policy is neither an index policy nor does it seem to have another simplifying form, it is still possible to identify characteristic patterns in the search behavior of an agent who follows the optimal strategy, which is what I do in this section. It is possible to derive the optimal search policy in specific cases, which I present in later subsections.

In section 2.3 I have derived characteristic values for an individual box i. Similar values can be derived if $i \in N$ is not the only remaining box. That is, for any box $i \in N$ we can ask the question when that box should be pre-, post- or not inspected at all, if the remaining boxes in $N \backslash \{i\}$ can still be inspected afterwards. Note, importantly, that it may not be optimal to inspect i next, but for the purpose of defining these characteristic values we ask what would be optimal if i was inspected next.

A note on notation: Throughout the section I will use the set $N \subset \mathcal{I}$ when referring to a specific state in the search process, as I have done so far. I will instead use the sets K, L with $L \subset K \subset \mathcal{I}$ when referring to changes over the course of the search process. The set K represents a state earlier in the search process than L. In other words, the set L arises from the set K after a number of additional boxes have been inspected. I will use these sets to express how the characteristic values of box i change over the course of the search process.

[12] Commonly in search models the index is a reservation value as in Weitzman (1979), a special case of the 'Gittins index' developed by Gittins and co-authors in the 1970s, see e.g. Gittins (1979).

[13] This aspect of the model is similar to Doval (2018), see her online appendix S.1 for a proof of the impossibility of an index policy in her model, which can be equivalently applied to this model.

First, we can take the derivatives of the value functions of pre- and post-inspecting box i with regard to z, such that we find, equivalently to the one-box case:

Lemma 2 (Single-crossing property)**.** *For any $i \in N \subset \mathcal{I}$ it holds that:*

(i) $0 \leq V_i^{post}(N,z)' < V_i^{pre}(N,z)' \leq 1$

(ii) $V_i^{pre}(N,z)$ *and* $V_i^{post}(N,z)$ *intersect at most once*

This makes it clear that we can again define a threshold value that separates for which values of z the agent wants to pre- or post-inspect box i. To properly define this threshold value θ_i^N, the equivalent to θ_i for general N, we first need to consider for which values of z the agent is willing to continue search at (N,z), when box i is inspected next.

Akin to the reservation values we can define two characteristic values for when the agent is indifferent between inspecting box i in a particular way (pre or post) and continuing with $N \setminus \{i\}$, and stopping search by taking z. Thus denote by $\phi_{i,N}^X$ the implicit solution to $\Delta V_i^X(N,z) = 0$, solved for z, with $X \in \{pre, post\}$. Note that while $\phi_{i,\{i\}}^X = r_i^X$, these values are conceptually different from reservation values, as becomes clear from the next result:[14]

Lemma 3. *$\phi_{i,K}^X \geq \phi_{i,L}^X$ for $X \in \{pre, post\}$ for all $L \subset K$.*

These characteristic values therefore may not stay constant over the search process. The proof is immediate: It is clear that $\Delta V_i^X(K,z) \geq \Delta V_i^X(L,z)$, as it can only be beneficial to have more boxes to choose from and so the value at which the agent is indifferent must be weakly larger for K. Intuitively it is clear that if $\phi_{i,K}^X > \phi_{i,L}^X$, there must be a box in K that is not in L, which has a reservation value larger than the largest reservation value of the boxes in L.

We can now define the threshold θ_i^N such that:

Proposition 1. *For every $i \in N$ there exists a threshold θ_i^N such that: If box i is inspected next in state (N,z), it is optimal to post-inspect i if $z < \theta_i^N$, pre-inspect i if $\theta_i^N \leq z < \max(\phi_{i,N}^{pre}, \phi_{i,N}^{post})$ and not inspect i if $\max(\phi_{i,N}^{pre}, \phi_{i,N}^{post}) \leq z$.*

The proposition shows that similar to the threshold θ_i in the one box case, we can define in cases where there are many boxes a similar threshold θ_i^N. In particular, $\theta_i = \theta_i^{\{i\}}$. When neither pre- nor post-inspection is strictly better, then θ_i^N is the solution of (2), as argued in section 2.3. When pre-inspection is strictly better, then θ_i^N is set to a value so low that the agent will always pre-inspect box i, as long as $z < \phi_{i,N}^{pre}$. When instead post-inspection is strictly better, then we set $\theta_i^N = \phi_{i,N}^{post}$, such that the agent post-inspects box i for all

[14] This is not specific to this model, it would also follow in Weitzman (1979) for a similar definition.

$z < \phi_{i,N}^{post}$. In these two extreme cases, θ_i^N therefore does not behave like a threshold between pre- and post-inspection, but this ensures that θ_i^N is consistently defined.

From these findings follows now the main result that establishes how the threshold evolves over the course of the search process:

Proposition 2. *In the optimal search strategy, for a given inspection order and $L \subset K \subset \mathcal{I}$:*

(i) $\forall i \in L : r_i^{pre} \geq r_i^{post} \Rightarrow \theta_i^K \leq \theta_i^L$

(ii) $\forall i \in L : r_i^{pre} < r_i^{post} \Rightarrow \theta_i^K \leq \theta_i^L$ *and* $\theta_i^K < \phi_{i,K}^{post}$ *if* $\theta_i^L < \phi_{i,L}^{post}$

In short, the proposition states that the threshold is weakly smaller the earlier in the search process box i is inspected (and under a condition in case *(ii)*). The reason for this is that pre-inspection becomes weakly better relative to post-inspection the more boxes can be inspected after box i. In that sense, the agent is more inclined to pre-inspect box i the earlier it is inspected. The intuition behind this result is clear: Post-inspection is ex-post better if the box is chosen at the end of search. This is more likely toward the end of search, as there are less unopened boxes remaining that could contain a larger value.

A distinction is necessary between boxes with a one-box reservation value (section 2.3) that is larger for pre-inspection, and the converse. In case *(i)*, where pre-inspection has the larger reservation value, we have shown in section 2.3 that there exists a threshold $\theta_i < r_i^{pre}$. The result above now states that in this case, the threshold is weakly smaller the earlier in the search box i is inspected. Case *(ii)*, where r_i^{post} is larger, states the same, but only under another condition. In the one-box case, we have seen that the agent strictly prefers to post-inspect box i when $r_i^{post} > r_i^{pre}$, and in Proposition 1 we have in that case defined $\theta_i^N = \phi_{i,N}^{post}$ for consistency only, as θ_i^N does not act like a threshold between pre- and post-inspection in that case. As shown in the proof of Proposition 2, pre-inspection becomes weakly better relative to post-inspection with every additional box that may be inspected after i. Then in the general case there exists an N with enough boxes besides i, such that $\theta_i^N < \phi_{i,N}^{post}$, i.e. post-inspection is not strictly better anymore, but it is now optimal to pre-inspect box i above this threshold and below $\phi_{i,N}^{pre}$. Proposition 2 *(ii)* then states that if the threshold has fallen below $\phi_{i,N}^{post}$ in this way, then the threshold is weakly smaller the earlier in the search box i is inspected, same as in case *(i)*. There are important special cases in which such an N large enough that $\theta_i^N < \phi_{i,N}^{post}$ does not exist when $r_i^{pre} < r_i^{post}$. One example is when $c_i^{pre} > c_i^{post} + \sigma_i$, such that pre-inspecting box i is strictly worse no matter whether it is ultimately chosen or not. Another notable example is the case of ex-ante homogeneous boxes. In that case, $r_i^{post} = r^{post}$, and so $\phi_{i,N}^{post} = r^{post}$ for all i. But then it is clear that $V_i^{post}(N, \phi_{i,N}^{post}) = 0$ and $V_i^{pre}(N, \phi_{i,N}^{post}) \leq 0$ for any N. Thus post-inspection is weakly better for any N; i.e. $\theta_i^N = r^{post}$ for all N with $i \in N$.

Proposition 2 has further implications. First, regardless of whether we are in case *(i)* or *(ii)*, it is clear that once θ_i^L reaches the lower bound[15] $\underline{x}$ such that pre-inspecting box i is strictly better, θ_i^K will also be equal to $\underline{x}$ for any K with $L \subset K$. Thus box i will only be pre-inspected if it is inspected earlier in the search process than in state $(L, .)$.

Further note that while the threshold increases with the progress of search, this does not imply that at some point the agent switches from only post-inspecting to only pre-inspecting boxes, not even when the boxes are ex-ante homogeneous. This is clear as in expectation z also weakly increases with every inspection, and so it may happen that the agent switches 'back and forth' between pre- and post-inspecting boxes.

Proposition 2 is formulated for a given inspection order. There are different settings in which this is the case. Many search models focus for their equilibrium analysis on either symmetric equilibria, in which the inspection order does not matter, as the boxes are essentially ex-ante homogeneous, or equilibria where the search order of a consumer is determined by the firms' prior decisions. In the case of online search, the inspection order may be suggested by a platform. A more general analysis is currently work in progress – it is conceivable that for example a stochastic order of the boxes' value distributions, e.g. such that earlier boxes stochastically dominate later boxes, may imply a corresponding inspection order, but this remains to be shown.

With this result established, I now discuss the similarities and differences to Doval (2018). In that paper, the alternative to pre-inspection is to choose a box without inspection and thereby end search. It is clear then that the agent may take this option at most once, necessarily at the end of search. In this paper, the agent may take the alternative inspection option of post-inspection for any box at any point of search. Thus we can identify how the value of post-inspection in relation to the classic pre-inspection varies over the length of the search process. This also allows us to observe the following fact: When an earlier box is post-inspected instead of pre-inspected, the observed value is artificially inflated by σ_i to $x_i + \sigma_i$. This means that the best observed outcome z weakly increases, making it weakly more likely that it surpasses the threshold θ_i^N of later boxes, such that they will be pre-inspected instead of post-inspected.

Doval's model emerges as the limit case of the model in this paper when we let $c_i^{post} \to 0$ and $\sigma_i \to \infty$ for all $i \in \mathcal{I}$. When it becomes infinitely costly to return a taken box, then the

[15] $\underline{x} := \min(\{x_0\} \cup \{\underline{x}_i\}_{i \in \overline{N}})$, the smallest value attainable during search.

decision to post-inspect a box implies that the box will be kept and search ends. Thus we can also see the relation between the threshold θ_i and the 'backup value' of her model. The backup value of box i is that value of z at which the agent is indifferent between pre-inspecting box i and selecting it without inspection. When we take the limit as described, it is clear that $\theta_i = \theta_i^{\{i\}}$ become equal to the backup value of box i.

3.1 Infinitely many homogeneous boxes

Many papers consider markets in which a unit mass of firms selling a homogeneous product operate, which represents a search problem with an infinite number of homogeneous boxes.

Proposition 3. *For an infinite number of homogeneous boxes the optimal search strategy is a reservation value policy where all boxes are post-inspected if $r^{post} > r^{pre}$ and otherwise all boxes are pre-inspected.*

This finding is consistent with Proposition 2 for a finite number of boxes. In the finite case with homogeneous boxes, if $r_i^{pre} < r_i^{post}$, then Proposition 2 implies that all boxes will be post-inspected, as discussed above. In the other case where $r_i^{pre} > r_i^{post}$, Proposition 2 implies that the threshold above which box i should be pre-inspected weakly decreases with the number of boxes that can still be inspected afterwards. In the case where the threshold is an internal solution of (2), it is strictly decreasing. If we let this number of boxes approach infinity, then it follows that all boxes should only be pre-inspected.

3.2 When pre-inspection is not feasible

In certain markets or for certain types of products pre-inspection may not be feasible, or post-inspection is simply much more beneficial. In these cases we can identify the optimal search policy.

A sufficient (but not necessary) condition for when pre-inspection is not feasible is when $c_i^{post} + \sigma_i < c_i^{pre}$, as in that case the transformed post-inspection has a smaller total inspection cost and provides a higher value $x_i + \sigma_i$, making it strictly better than pre-inspection.

When boxes are only post-inspected, then the optimal policy is immediately clear:

Proposition 4. *If boxes are only post-inspected, then the optimal policy is "Pandora's rule" of Weitzman (1979) applied to the transformed problem.*

Thus the optimal policy is clear in settings where boxes may only be post-inspected.[16,17] This

[16] This is a version of the setting of Petrikaitė (2018) with a finite number of products instead of two.

[17] An equivalent result holds if there are some boxes that may only be pre-inspected and others that

also immediately implies that all results for a setting as in Weitzman (1979) apply as well when all boxes are only post-inspected.

In particular, we can formulate an eventual purchase theorem - a discrete choice reformulation of the search model. Following the notation of Armstrong (2017) we find:

Proposition 5. *If boxes are only post-inspected, then the agent will ultimately choose the box that maximizes the following index:*

$$w_i^{post} := \min\{r_i^{post}, x_i + \sigma_i\}$$

The proof is analogous to the classic case in Armstrong (2017) and Choi, Dai, and Kim (2018), applied to the transformed problem. As a result we can further determine the expected demand for each box. If boxes represent products of firms, we can thus determine a firm's demand in settings with product returns, a helpful tool for empirical applications.

A possible extension in this case is to split the inspection of a box into an initial 'browsing' stage, where partial information about the box's value is revealed, and a following proper 'inspection' stage, where the full value is revealed. The agent's possible actions in a given state are then to browse a new box, inspect an already browsed box or choose an already inspected box or initial outside option to end search.

As the chosen terminology suggests, this extension may give additional insights in settings like online shopping, where some information about a product may readily be accessed before ordering by browsing, but the full value of the product is typically learned only after the product has been ordered and tried.

This extended model can be solved by applying the branching bandits framework, see e.g. Keller and Oldale (2003). In addition to the optimal search policy, an eventual purchase theorem can also be derived in this case[18], such that firm demand can be identified e.g. in the online shopping context.

4 Discussion

In this paper I have studied settings where an agent, who searches for the best alternative out of many, is given the option to 'post-inspect' an alternative after they have gained access to it, thereby reducing inspection costs but entailing a return cost if the alternative is returned.

may only be post-inspected, but none that may be inspected in both ways.

[18] See Greminger (2021) for the derivation of an EPT in a similar branching setting.

I have found that the agent makes use of this inspection option more often, relative to the classic 'pre-inspection' before gaining access, the less alternatives still remain to be inspected.

The inspection option of post-inspecting an alternative can be interpreted as treating the alternative as an experience good. In that sense, the inspection cost c_i^{post} may even be negative (representing a utility gain) if it is interpreted as the utility derived from using the good while learning its value. Then the transformation I employ for post-inspection can be interpreted as transforming an experience good into a search good.

References

Anderson, Simon P., Maxim Engers, and Daniel Savelle (2021). "Consumer Journeys with Nested Weitzman Scores and Informational Complementarities".

Armstrong, Mark (2017). "Ordered Consumer Search". In: *Journal of the European Economic Association* 15.5, pp. 989–1024.

Choi, Michael, Anovia Yifan Dai, and Kyungmin Kim (2018). "Consumer Search and Price Competition". In: *Econometrica* 86.4, pp. 1257–1281.

Doval, Laura (2018). "Whether or Not to Open Pandora's Box". In: *Journal of Economic Theory* 175, pp. 127–158.

Ertner, Sebastian and Maarten C.W. Janssen (2025). "Multiproduct Firms, Refunds and Product Returns".

Gambato, Jacopo (2024). "Consumer Search and Firm Strategy with Multi-Attribute Products".

Gittins, J. C. (1979). "Bandit Processes and Dynamic Allocation Indices". In: *Journal of the Royal Statistical Society: Series B (Methodological)* 41.2, pp. 148–164.

Greminger, Rafael P. (2021). "Optimal Search and Discovery". In: *Management Science* 68.5, pp. 3904–3924.

Heidhues, Paul, Johannes Johnen, and Botond Kőszegi (2021). "Browsing versus Studying: A Pro-market Case for Regulation". In: *The Review of Economic Studies* 88.2, pp. 708–729.

Janssen, Maarten and Cole Williams (2024). "Consumer Search and Product Returns in E-Commerce". In: *American Economic Journal: Microeconomics* 16.2, pp. 387–419.

Keller, Godfrey and Alison Oldale (2003). "Branching Bandits: A Sequential Search Process with Correlated Pay-Offs". In: *Journal of Economic Theory* 113.2, pp. 302–315.

Matthews, Steven A. and Nicola Persico (2007). "Information Acquisition and Refunds for Returns". SSRN Scholarly Paper. Rochester, NY.

Petrikaitė, Vaiva (2018). “A Search Model of Costly Product Returns”. In: *International Journal of Industrial Organization* 58, pp. 236–251.

Weitzman, Martin L. (1979). “Optimal Search for the Best Alternative”. In: *Econometrica* 47.3, pp. 641–654.

Zhang, Yutong (2026). “Nested Search”.

A Appendix

A.1 Proof of Lemma 1

The proof is in two parts. I first show that this transformation is valid if all boxes are only post-inspected. Using that result, I then show that we can also make this transformation when boxes may be pre- or post-inspected.

Part 1) I show that a search problem where all boxes have return costs can be transformed into a search problem without any return costs and which is of the same form as in Weitzman (1979). For this proof I deviate slightly from the notation and model definitions above. Let N denote a set of uninspected boxes, where each box i has a value distribution with cdf F_i on the support $[\underline{x}_i, \overline{x}_i]$, an inspection cost c_i and a return cost σ_i to be paid if the box is returned. Denote by $V(N, x_z, \sigma_z)$ the value function for a given set of uninspected boxes N as well as the value x_z and return cost σ_z of the so far best realized outcome z, which can be either an already inspected box or the initial outside option.[19] Note that while it is possible for box j that the realized value x_j is so negative that $x_j \leq -\sigma_j$, the best realized outcome z will always have $x_z > -\sigma_z$ (otherwise it is always better to return that box). This implies that $V(\emptyset, x_z, \sigma_z) = x_z$, that is, if there are no uninspected boxes left, then it is optimal to choose the best realized outcome with value x_z and the exact value of the return cost σ_z is irrelevant. The dynamic problem of search with costly returns is then:

$$\begin{aligned} V(N, x_z, \sigma_z) = \max \Bigg\{ x_z, \max_{i \in N} \Bigg\{ -c_i + \int_{-\infty}^{x_z - \sigma_i + \sigma_z} \left(V(N \backslash \{i\}, x_z, \sigma_z) - \sigma_i \right) dF_i(x_i) \\ + \int_{x_z - \sigma_i + \sigma_z}^{\infty} \left(V(N \backslash \{i\}, x_i, \sigma_i) - \sigma_z \right) dF_i(x_i) \Bigg\} \Bigg\} \end{aligned} \tag{3}$$

This value function can be understood as follows: When the agent is in state (N, x_z, σ_z), they may either take x_z and end search, or inspect another box i. In that case they pay the inspection cost c_i and compare keeping i and returning z to keeping z and returning i and

[19] As argued in section 2, if the best realized outcome is one of the inspected boxes, it is the box j with largest $x_j + \sigma_j$.

do whichever provides the larger utility. Thereby they transition to the next state, keeping either i or z as described.

Consider now the problem in Weitzman (1979), which has no return costs for all boxes and can be written as:

$$V(N^W, x_z^W, 0) = \max\left\{x_z^W, \max_{i \in N^W}\left\{-c_i^W + \int_{-\infty}^{x_z^W} V(N^W \backslash \{i\}, x_z^W, 0) dF_i^W(x_i^W) \right.\right. \\ \left.\left. \int_{x_z^W}^{\infty} V(N^W \backslash \{i\}, x_i^W, 0) dF_i^W(x_i^W)\right\}\right\} \tag{4}$$

where all elements in N^W have $\sigma_i^W = 0$. As we know the solution to (4), the remainder of this proof shows how we can transform any problem of the form (3) to be equivalent to (4).

$$\begin{aligned} V(N, x_z, \sigma_z) = & \max\left\{x_z, \max_{i \in N}\left\{-c_i + \int_{-\infty}^{x_z - \sigma_i + \sigma_z} \left(V(N \backslash \{i\}, x_z, \sigma_z) - \sigma_i\right) dF_i(x_i)\right.\right. \\ & \left.\left. + \int_{x_z - \sigma_i + \sigma_z}^{\infty} \left(V(N \backslash \{i\}, x_i, \sigma_i) - \sigma_z\right) dF_i(x_i)\right\}\right\} \\ = & \max\left\{x_z, \max_{i \in N}\left\{-c_i - \sigma_i + \int_{-\infty}^{x_z - \sigma_i + \sigma_z} \left(V(N \backslash \{i\}, x_z, \sigma_z)\right) dF_i(x_i)\right.\right. \\ & \left.\left. + \int_{x_z - \sigma_i + \sigma_z}^{\infty} \left(V(N \backslash \{i\}, x_i, \sigma_i) - \sigma_z + \sigma_i\right) dF_i(x_i)\right\}\right\} \\ = & - \sigma_z + \max\left\{x_z + \sigma_z, \max_{i \in N}\left\{-c_i - \sigma_i + \int_{-\infty}^{x_z - \sigma_i + \sigma_z} \left(V(N \backslash \{i\}, x_z, \sigma_z) + \sigma_z\right) dF_i(x_i)\right.\right. \\ & \left.\left. + \int_{x_z - \sigma_i + \sigma_z}^{\infty} \left(V(N \backslash \{i\}, x_i, \sigma_i) + \sigma_i\right) dF_i(x_i)\right\}\right\} \end{aligned}$$

We can now transform our search problem with return costs and show that the value function above is indeed equivalent to the value function of a search problem without return costs minus σ_z, such that overall we have $V(N, x_z, \sigma_z) = V(N', x_z + \sigma_z, 0) - \sigma_z$. To that end we define a new set N' with the same number of elements in the following way: for each element i in N, there is a corresponding element i' in N' such that $F_i'(x+\sigma_i) = F_i(x)$, $c_i' = c_i + \sigma_i$ and $\sigma_i' = 0$. Intuitively we have done the following: every element in N' has zero return cost, but instead the return cost is pre-emptively paid together with the inspection cost and reimbursed in those cases where the box is kept, i.e. it is added to the box value.

We now prove by induction that $V(N, x_z, \sigma_z) = V(N', x_z + \sigma_z, 0) - \sigma_z$ for any set N.[20] For the base case assume $N = \{i\}$ and $N' = \{i'\}$, i.e. there is only one uninspected box left.

[20] The subtraction of σ_z on the RHS is necessary for consistency regarding the best realized outcome z. In the original view (LHS), the potential payment of σ_z is a future decision, while in the transformed view (RHS), σ_z had to be paid when z was inspected and is reimbursed only when z is chosen in the end. Note that at the start of search, the initial outside option will typically have $\sigma_0 = 0$ in most applications. In that case, the value functions of original and transformed problem are exactly equal.

Then:

$$
\begin{aligned}
V(N,x_z,\sigma_z) =& \max\left\{x_z+\sigma_z, -c_i-\sigma_i+\int_{-\infty}^{x_z-\sigma_i+\sigma_z}\left(V(\emptyset,x_z,\sigma_z)+\sigma_z\right)dF_i(x_i)\right.\\
&\left.+\int_{x_z-\sigma_i+\sigma_z}^{\infty}\left(V(\emptyset,x_i,\sigma_i)+\sigma_i\right)dF_i(x_i)\right\}-\sigma_z=\\
=&\max\left\{x_z+\sigma_z, -c_i-\sigma_i+\int_{-\infty}^{x_z-\sigma_i+\sigma_z}\left(x_z+\sigma_z\right)dF_i(x_i)\right.\\
&\left.+\int_{x_z-\sigma_i+\sigma_z}^{\infty}\left(x_i+\sigma_i\right)dF_i(x_i)\right\}-\sigma_z=\\
=&\max\left\{x_z+\sigma_z, -c_i-\sigma_i+\int_{-\infty}^{x_z+\sigma_z}V(\emptyset,x_z+\sigma_z,0)dF_i'(x_i+\sigma_i)\right.\\
&\left.+\int_{x_z+\sigma_z}^{\infty}V(\emptyset,x_i+\sigma_i,0)dF_i'(x_i+\sigma_i)\right\}-\sigma_z=\\
=&V(N',x_z+\sigma_z,0)-\sigma_z
\end{aligned}
$$

Note that in the first line the second integral goes over all $x_i \geq x_z+\sigma_z-\sigma_i$, which implies $x_i \geq -\sigma_i$ and therefore $V(\emptyset,x_i,\sigma_i)=x_i$.

For the induction step assume that $V(N\backslash\{i\},x_z,\sigma_z)=V(N'\backslash\{i'\},x_z+\sigma_z,0)-\sigma_z$. Then:

$$
\begin{aligned}
V(N,x_z,\sigma_z) =& -\sigma_z+\max\left\{x_z+\sigma_z, \max_{i\in N}\left\{-c_i-\sigma_i+\int_{-\infty}^{x_z-\sigma_i+\sigma_z}\left(V(N\backslash\{i\},x_z,\sigma_z)+\sigma_z\right)dF_i(x_i)\right.\right.\\
&\left.\left.+\int_{x_z-\sigma_i+\sigma_z}^{\infty}\left(V(N\backslash\{i\},x_i,\sigma_i)+\sigma_i\right)dF_i(x_i)\right\}\right\}=\\
=& -\sigma_z+\max\left\{x_z+\sigma_z, \max_{i'\in N'}\left\{-c_i-\sigma_i+\int_{-\infty}^{x_z+\sigma_z}V(N'\backslash\{i'\},x_z+\sigma_z,0)dF_i'(x_i+\sigma_i)\right.\right.\\
&\left.\left.+\int_{x_z+\sigma_z}^{\infty}V(N'\backslash\{i'\},x_i+\sigma_i,0)dF_i'(x_i+\sigma_i)\right\}\right\}=V(N',x_z+\sigma_z,0)-\sigma_z
\end{aligned}
$$

Thus indeed $V(N,x_z,\sigma_z)=V(N',x_z+\sigma_z,0)-\sigma_z$ holds for any set N and so we have proven that the original problem with return costs can be transformed into an equivalent problem without return costs.

Part 2) For the proof to hold also when any box may be pre- or post-inspected, the following adjustments have to be made to Part 1: the set of uninspected boxes is twice as large, i.e. it contains two copies for each box i where one can only be pre-inspected and the other only post-inspected. Denote that set by N_2 with cardinality $|N_2|=2|N|$ as described. Define the function $\gamma(i)$ that returns the copy of box i. Inspecting one copy also removes the other copy from the set of remaining boxes, i.e. after inspecting copy i the set of uninspected boxes is $N_2\backslash\{i,\gamma(i)\}$. We then treat any copy in N_2 initially as a box with return cost σ_i where the copies that are pre-inspected simply have $\sigma_i = 0$. With small changes in the proof (we need to account for the fact that with each inspection two elements get removed from N_2) we can then apply the same argument as in Part 1, which transforms each copy into one without return costs. Thus the post-inspected copies get transformed as in Part 1, while the pre-inspected copies also get transformed, but as they have $\sigma_i = 0$ to begin with, their transformed version is identical to the initial version. $\square$

A.2 Proof of Lemma 2

(i) We first show that $V'(N,z) \in [0,1]$. With the definition of the value function in (1):

$$V'(N,z) = \begin{cases} 1 & \text{if } V(N,z) = z \\ F_j(z)V'(N\backslash\{j\},z) & \text{if } V(N,z) = \max_{j\in N} V_j^{pre}(N,z) \\ F_i(z-\sigma_j)V'(N\backslash\{i\},z) & \text{if } V(N,z) = \max_{i\in N} V_i^{post}(N,z) \end{cases}$$

We know that $V(\emptyset, z) = z$ and therefore $V'(\emptyset, z) = 1$. From that follows by induction that we indeed have $V'(N,z) \in [0,1]$ for any N.

Consider now the derivatives:

$$V_i^{pre}(N,z)' = F_i(z)V(N\backslash\{i\},z)' \qquad \text{and} \qquad V_i^{post}(N,z)' = F_i(z-\sigma_i)V(N\backslash\{i\},z)'$$

With $V(N,z)' \in [0,1]$ it follows that $0 \leq V_i^{post}(N,z)' < V_i^{pre}(N,z)' \leq 1$ for any N.
(ii) It follows immediately that there is at most one intersection point z of V_i^{pre} and V_i^{post}. $\square$

A.3 Proof of Proposition 1

When such an indifference point exists, we define θ_i^N as that z for which the searcher is indifferent between pre- and post-inspecting box i, with the possibility to inspect other boxes in set N afterwards. If the intersection point does not exist, then it is enough to set θ_i^N at either the lowest or highest value for which there is a positive gain from search to ensure the prescribed inspection behavior is indeed optimal.

We define $D_i^N(z) := V_i^{pre}(N,z) - V_i^{post}(N,z)$, the difference in the value function between pre- and post-inspecting box i. Then θ_i^N is defined as the z for which $D_i^N(z) = 0$, if it exists in $[\underline{x}_i, r_{i,N}^{pre}]$, otherwise it is set to the "closest border", i.e.:

$$\theta_i^N := \begin{cases} \underline{x} & \text{if } D_i^N(z) > 0 \ \forall z \in [\underline{x}_i, \phi_{i,N}^{pre}] \\ z : D_i^N(z) = 0 & \text{if } \exists z \in [\underline{x}_i, \phi_{i,N}^{pre}] : D_i^N(z) = 0 \\ \phi_{i,N}^{post} & \text{if } D_i^N(z) < 0 \ \forall z \in [\underline{x}_i, \phi_{i,N}^{pre}] \end{cases}$$

with $\underline{x} := \min(\{x_0\} \cup \{\underline{x}_i\}_{i \in \overline{N}})$. As shown in Lemma 2, if an intersection point with $D_i^N(z) = 0$ exists for some z, then it must be unique, as $D_i^N(z)' \geq 0$.[21] □

A.4 Proof of Proposition 2

First, for a given inspection order it holds that $V(K,z)' \leq V(L,z)'$ for any $L \subset K \subset \mathcal{I}$, as the following shows:

$$V'(K,z) = \begin{cases} 1 & \text{if } V(K,z) = z \\ F_i(z) V'(K \backslash \{i\}, z) & \text{if } V(K,z) = \max_{i \in K} V_i^{pre}(K,z) \\ F_i(z - \sigma_i) V'(K \backslash \{i\}, z) & \text{if } V(K,z) = \max_{i \in K} V_i^{post}(K,z) \end{cases}$$

If $V(K,z) = z$, then also $V(K \backslash \{i\}, z) = z$. With that it is clear that $V'(K,z) \leq V'(K \backslash \{i\}, z)$ and by induction $V(K,z)' \leq V(L,z)'$.

Now, recall the definitions of V_i^{pre} and V_i^{post} from section 2.2. We can rewrite V_i^{post}:

$$V_i^{post}(N,z) := -c_i^{post} - \sigma_i + F_i(z - \sigma_i) V(N \backslash \{i\}, z) + \int_z^{\overline{x}_i + \sigma_i} V(N \backslash \{i\}, x_i) f_i(x_i - \sigma_i) dx_i$$

Define the difference $D_i^N(z) := V_i^{pre}(N,z) - V_i^{post}(N,z)$ such that:

$$\begin{aligned} D_i^N(z) &= \left[-c_i^{pre} + F_i(z) V(N \backslash \{i\}, z) + \int_z^{\overline{x}_i} V(N \backslash \{i\}, x_i) f_i(x_i) dx_i \right] \\ &- \left[-c_i^{post} - \sigma_i + F_i(z - \sigma_i) V(N \backslash \{i\}, z) + \int_z^{\overline{x}_i + \sigma_i} V(N \backslash \{i\}, x_i) f_i(x_i - \sigma_i) dx_i \right] = \\ &= -(c_i^{pre} - c_i^{post} - \sigma_i) + [F_i(z) - F_i(z - \sigma_i)] V(N \backslash \{i\}, z) \\ &+ \int_z^{\overline{x}_i} V(N \backslash \{i\}, x_i) [f_i(x_i) - f_i(x_i - \sigma_i)] dx_i - \int_{\overline{x}_i}^{\overline{x}_i + \sigma_i} V(N \backslash \{i\}, x_i) f_i(x_i - \sigma_i) dx_i \end{aligned}$$

Then consider:

[21] Note that θ_i^N is set to $\phi_{i,N}^{post}$ and not $\phi_{i,N}^{pre}$ in the third case, as in that case $\phi_{i,N}^{post} > \phi_{i,N}^{pre}$.

$$
\begin{aligned}
D_i^K(\theta_i^L) - D_i^L(\theta_i^L) = &[F_i(\theta_i^L) - F_i(\theta_i^L - \sigma_i)][V(K\backslash\{i\},\theta_i^L) - V(L\backslash\{i\},\theta_i^L)] \\
&+ \int_{\theta_i^L}^{\overline{x}_i} [V(K\backslash\{i\},x_i) - V(L\backslash\{i\},x_i)][f_i(x_i) - f_i(x_i - \sigma_i)]dx_i \\
&- \int_{\overline{x}_i}^{\overline{x}_i+\sigma_i} [V(K\backslash\{i\},x_i) - V(L\backslash\{i\},x_i)]f_i(x_i - \sigma_i)dx_i
\end{aligned}
$$

This expression is the most negative if f_i is decreasing. Supposing that, and as we know that $\partial/\partial z[V(K\backslash\{i\},z) - V(L\backslash\{i\},z)] \leq 0$ for a given inspection order:

$$
\begin{aligned}
&\int_{\theta_i^L}^{\overline{x}_i} [V(K\backslash\{i\},x_i) - V(L\backslash\{i\},x_i)][f_i(x_i) - f_i(x_i - \sigma_i)]dx_i \geq \\
\geq &[V(K\backslash\{i\},\theta_i^L) - V(L\backslash\{i\},\theta_i^L)] \int_{\theta_i^L}^{\overline{x}_i} [f_i(x_i) - f_i(x_i - \sigma_i)]dx_i = \\
= &[V(K\backslash\{i\},\theta_i^L) - V(L\backslash\{i\},\theta_i^L)] \left\{[1 - F_i(\overline{x}_i - \sigma_i)] - [F_i(\theta_i^L) - F_i(\theta_i^L - \sigma_i)]\right\}
\end{aligned}
$$

And:

$$
\begin{aligned}
&- \int_{\overline{x}_i}^{\overline{x}_i+\sigma_i} [V(K\backslash\{i\},x_i) - V(L\backslash\{i\},x_i)]f_i(x_i - \sigma_i)dx_i \geq \\
\geq &- [V(K\backslash\{i\},\overline{x}_i) - V(L\backslash\{i\},\overline{x}_i)] \int_{\overline{x}_i}^{\overline{x}_i+\sigma_i} f_i(x_i - \sigma_i)dx_i = \\
= &- [V(K\backslash\{i\},\overline{x}_i) - V(L\backslash\{i\},\overline{x}_i)][1 - F_i(\overline{x}_i - \sigma_i)]
\end{aligned}
$$

Then together we have:

$$
\begin{aligned}
&D_i^K(\theta_i^L) - D_i^L(\theta_i^L) \geq \\
&\geq [F_i(\theta_i^L) - F_i(\theta_i^L - \sigma_i)][V(K\backslash\{i\},\theta_i^L) - V(L\backslash\{i\},\theta_i^L)] \\
&+ [V(K\backslash\{i\},\theta_i^L) - V(L\backslash\{i\},\theta_i^L)] \left\{[1 - F_i(\overline{x}_i - \sigma_i)] - [F_i(\theta_i^L) - F_i(\theta_i^L - \sigma_i)]\right\} \\
&- [V(K\backslash\{i\},\overline{x}_i) - V(L\backslash\{i\},\overline{x}_i)][1 - F_i(\overline{x}_i - \sigma_i)] = \\
&= \left\{[V(K\backslash\{i\},\theta_i^L) - V(L\backslash\{i\},\theta_i^L)] - [V(K\backslash\{i\},\overline{x}_i) - V(L\backslash\{i\},\overline{x}_i)]\right\} [1 - F_i(\overline{x}_i - \sigma_i)] \geq 0
\end{aligned}
$$

Where we again use that $V(K\backslash\{i\},z)' - V(L\backslash\{i\},z)' \leq 0$ for a given inspection order.

We have now shown that $D_i^K(\theta_i^L) \geq D_i^L(\theta_i^L)$, i.e. pre-inspection becomes weakly better relative to post-inspection the earlier we are in the search process. How this translates into the evolution of the threshold value θ_i^N depends on whether θ_i^L is an internal solution of (2) or a corner case.

- If θ_i^L is an internal intersection point $\theta_i^L \in [\underline{x}_i, \phi_{i,L}^{pre}]$, then we know that $D_i^L(\theta_i^L) = 0$ and therefore $D_i^K(\theta_i^L) \geq 0$. To achieve $D_i^K(\theta_i^K) = 0$ it must be that $\theta_i^K \leq \theta_i^L$ as $D_i^N(z)' \geq 0$.

- If pre-inspection is strictly better, then $\theta_i^L = \underline{x}$ and we know that $D_i^L(\theta_i^L) > 0$ and thus $D_i^K(\theta_i^L) > 0$. Decreasing the threshold to achieve $D_i^K(\theta_i^K) = 0$ is impossible, and so pre-inspection is still strictly better for K and we set $\theta_i^K = \theta_i^L = \underline{x}$. Note that this immediately implies that any J with $L \subset K \subset J \subset \mathcal{I}$ will also have $\theta_i^J = \underline{x}$.

- If post-inspection is strictly better, then $\theta_i^L = \phi_{i,L}^{post}$ and we know that $D_i^L(\theta_i^L) < 0$ so that it is possible that $D_i^K(\theta_i^L)$ is either positive or negative. If it is positive, then due to $D_i^K(z)' \geq 0$ we get $\theta_i^K < \theta_i^L = \phi_{i,L}^{post}$. If it is negative, then we would like to raise the threshold. This is the only case that can result in $\theta_i^K > \theta_i^L$, however the maximum is $\theta_i^K = \phi_{i,K}^{post}$ (note that post-inspection being strictly better requires $\phi_{i,K}^{pre} < \phi_{i,K}^{post}$). As Lemma 3 shows that $\phi_K^{post} \geq \phi_L^{post}$, an increase of $\theta_i^K > \theta_i^L$ is only possible when $\phi_K^{post} > \phi_L^{post}$. If the latter increases more strongly, then we are in the situation that the agent strictly prefers post-inspection in L, but will pre-inspect in K for some values of z, even though $\theta_i^K > \theta_i^L$! Note also that since $\phi_{i,N}^{post}$ increases with the cardinality of N, we find that once $\theta_i^K < \phi_{i,K}^{post}$, we are at an internal solution and therefore can never go back to $\theta_i^J = \phi_{i,J}^{post}$ for any J with $L \subset K \subset J \subset \mathcal{I}$.

How does this relate to the one-box reservation values?

Case (i) When $r_i^{pre} \geq r_i^{post}$, then there is either an internal point of indifference or pre-inspection is strictly better, as we have seen in section 2.3. We have just argued that in both these cases we will have $\theta_i^K \leq \theta_i^L$.

Case (ii) When $r_i^{pre} < r_i^{post}$, then post-inspection is strictly better, as we have seen in section 2.3.[22] We have just seen that in that case, when $\theta_i^L = \phi_{i,L}^{post}$, either $\theta_i^K = \phi_{i,K}^{post}$, or $\theta_i^K < \phi_{i,K}^{post}$. In the latter case we transition to an internal threshold value, and thus for any earlier search stage the threshold value will also be either an internal value, or pre-inspection will be strictly better, thus $\theta_i^L < \phi_{i,L}^{post} \Rightarrow \theta_i^K \leq \theta_i^L$. □

A.5 Proof of Proposition 3

Assume there is an infinite number of homogeneous boxes, that is $c_j^{pre} = c^{pre}, c_j^{post} = c^{post}, \sigma_j = \sigma, F_j(x_j) = F(x)$ for all j. Suppose that $i-1$ boxes have already been opened and the agent decides whether to open box i as well, with $i \geq 1$. Denote then by $z_{i-1} = \max\{z_0, \xi_1, \ldots, \xi_{i-1}\}$ the so far largest observed value, with z_0 being the initial outside option and the ξ_j being the observed values of the opened boxes for $j \in \{1, \ldots, i-1\}$.[23] Then the agent can choose one of the following three actions: (i) take z_{i-1} and stop search, (ii) pre- or (iii) post-inspect

[22] In particular, $\theta_i^{\{i\}} = \phi_{i,\{i\}}^{post} = r_i^{post}$.

[23] $\xi_j := x_j$ if box j was pre-inspected and $\xi_j := x_j + \sigma$ if it was post-inspected.

another box ("box i").

Now suppose that it is optimal for the agent to inspect box i with either action (ii) or (iii) and observe value ξ_i inside. If $\xi_i > z_{i-1}$, then $z_i = \xi_i$ and z_{i-1} will not be relevant for any later decision, so the agent can forget about it. If instead $\xi_i \leq z_{i-1}$, then $z_i = z_{i-1}$ and the agent will instead forget about ξ_i. But note that in this latter case, the agent is in the exact same position for box $i+1$ as they were for box i. This implies that it must be optimal to take the same action for box $i+1$ as for box i. In particular, the agent will inspect another box. But that means whatever the outcome of the ith search, the agent will never come back and end search by selecting z_{i-1} – if they start opening another box, they will open more boxes at least until they find a value larger than z_{i-1}. Therefore following an optimal search policy implies that the agent will never come back to any already opened boxes if they decide to open another one. Thus for an infinite number of homogeneous boxes the agent optimally searches as if they had no recall.

Then the dynamic problem of the agent can be written down as follows:

$$V_i(z_{i-1}) = \max\{z_{i-1}, \int V_{i+1}(x_i)dF(x_i) - c^{pre}, \int V_{i+1}(x_i + \sigma)dF(x_i) - c^{post} - \sigma\}$$

The three expressions in the maximum function correspond to actions (i), (ii) and (iii) above. The second and third expressions give the respective expected values of discarding z_{i-1} and starting new by pre-/post-inspecting a box. It is immediately clear that one of those two terms will be larger than the other for any value of z_{i-1} (as they are independent of it), and therefore it can only be optimal to either pre-inspect all boxes or post-inspect all boxes.[24]

We solve for the value function when all boxes are pre-inspected compared to all being post-inspected to determine when it is optimal to do either of those. For that we first simplify the dynamic problem further. As we know that there is an infinite number of boxes it follows that $V_i(z) = V_{i+1}(z) =: V(z)$. If the agent then pre-inspects all boxes, we find:

$$V(z) = \max\{z, \int V(x)dF(x) - c^{pre}\}$$

It is clear that the agent will only stop by choosing z if they expect to get less from a new search. Therefore there will be a threshold r^{pre} for z above which the agent stops and below which they make a new search. At $z = r^{pre}$ they are indifferent between stopping and

[24] Excluding the edge case of indifference when both expressions are equal.

continuing:

$$V(r^{pre}) = r^{pre} = \int V(x)dF(x) - c^{pre}$$

Then $V(x) = x$ for $x > r^{pre}$ and $V(x) = V(r^{pre}) = r^{pre}$ for $x \leq r^{pre}$. So r^{pre} is well defined as the solution to $\int \max\{x - r^{pre}, 0\}dF(x) - c^{pre} = 0$.

If the agent instead post-inspects all boxes, we find:

$$V(z) = \max\{z, \int V(x)dF(x) - c^{post} - \sigma\}$$

It is clear that the agent will only stop by choosing z if they expect to get less from a new search. Therefore there will be a threshold r^{post} for z above which the agent stops and below which they make a new search. At $z = r^{post}$ they are indifferent between stopping and continuing:

$$V(r^{post}) = r^{post} = \int V(x)dF(x) - c^{post} - \sigma$$

Then $V(x) = x$ for $x > r^{post}$ and $V(x) = V(r^{post}) = r^{post}$ for $x \leq r^{post}$. So r^{post} is well defined as the solution to $\int \max\{x + \sigma - r^{post}, 0\}dF(x) - c^{post} - \sigma = 0$.

Then it follows for our dynamic problem with given outside option z:

$$V(z) = \max\{z, r^{post}, r^{pre}\}$$

□

A.6 Proof of Proposition 4

The proof follows immediately from Part 1 in the proof of Lemma 1. Notice how $V(N', x_z + \sigma_z, 0)$ is equivalent to (4), the problem in Weitzman (1979), with the following correspondences: $x_i^W \leftrightarrow x_i + \sigma_i, F_i^W(x_i^W) \leftrightarrow F_i'(x_i + \sigma_i), c_i^W \leftrightarrow c_i + \sigma_i, \sigma_i^W = \sigma_i' = 0$ and therefore $N^W \leftrightarrow N'$. As the subtraction of the parameter σ_z does not affect the solution of the dynamic problem, it follows that the optimal search policy for the search problem with return costs is Pandora's Rule applied to its transformed problem. □