Search with Interdependent Values*

Renjie Bao — Princeton[†]

Sanxi Li — RUC[‡]

Jun Yu — SUFE[§]

January 17, 2023

Abstract

We propose a sequential search theory where values of different items are interdependent via a common attribute. Compared to the model with independent values, search here not only reveals the value of the current item, but also provides information on uninspected ones. This interdependence leads to a "cockroach" effect, that is, a bad realization can signal poor values for other items. Therefore, the agent will stop searching when the current realization is either good enough, which aligns with the literature, or bad enough, which is new to the search literature. The cockroach effect also adds up to the trade-off in determining the optimal search order, making small-variance items valuable as they contain more information on the common attribute. Finally, our theory predicts recall behavior as well as history-dependent search order, which further complements the existing search theory.

Keywords: Search; Interdependent Values; Stopping Rule; Search Order; Recall; History Dependence

*We sincerely thank Mark Armstrong, Jan Eeckhout, Shuo Liu, Xi Weng, Yonghui Zhang, Jie Zheng and Jidong Zhou for many useful comments and insightful discussions, as well as seminar audiences at Society for the Advancement of Economic Theory, Peking University, and Universitat Pompeu Fabra. All errors are ours.

[†]Department of Economics, Princeton University; Julis Romo Rabinowitz Building, Washington Rd, Princeton, NJ 08544, renjie.bao@princeton.edu

[‡]School of Economics, Renmin University of China; Center for Digital Economy Research, Renmin University of China, Beijing, 100872, China, sanxi@ruc.edu.cn

[§]School of Economics, Shanghai University of Finance and Economics; Key Laboratory of Mathematical Economics (SUFE), Ministry of Education, Shanghai, 200433, China, yu.jun@mail.shufe.edu.cn

1 Introduction

Search is the key element in economic activities that captures how people acquire information on specific subjects. In the classic search theory, values of different items are usually assumed to be independent, which is hardly true in the real world. For example, computers built with the same CPU can have similar performances (product market); students from the same college may share some common ability (labor market); different bonds also suffer from the same aggregate shock (financial market). In these cases, search is not only a process that uncovers values of the sampled items, but also a way for people to learn about the remaining uninspected ones. This paper contributes in analyzing how this interdependence reshapes an agent's optimal search strategy.

Specifically, we build a parsimonious search model based on the framework of Pandora's Box by Weitzman (1979). In the original problem, Pandora is faced with boxes containing different and independent values. Although she knows the value distribution of each box, she has to open them sequentially to find the exact realization, which incurs a search cost. The optimal search strategy, also called the *Pandora's Rule*, asks her to open boxes in the descending order of the reservation price, which is an index computed from each box's characteristic, and to stop sampling if and only if the current best realization is greater than the reservation prices for all closed boxes. Therefore, she will stop opening new boxes only if the current realization is good enough, will prioritize a large-variance box because it induces a higher reservation price, and will not return to old offers before opening all boxes. Moreover, the optimal order here is independent of her search history.

Relaxing the assumption of independence, we allow the values of different boxes to be interdependent, which creates a new mechanism in the search theory. Let the value of each box consists of a common and an idiosyncratic attribute. Although Pandora can only see the total match value from opening one box, she can update her belief on the common attribute via Bayesian rule and hence has better information on the remaining closed boxes. This mechanism leads to a brand new effect, which we call "cockroach" effect, that a bad realization can also be valuable in the sense that it signals a poor common attribute, which can help her make better search decisions in the future.[1] This cockroach effect makes the Pandora's Rule no longer work.

In the baseline model with two boxes, we find two cutoff values in her optimal stopping rule after the agent samples the first box. Consistent with the conventional wisdom, she will stop sampling if the first realization is good enough. We interpret this result as the "luck effect", that is, a good match value indicates a high idiosyncratic attribute for the first box, which stops her from opening the second one. On the other hand, we complement the literature by bringing out

[1] The name of "cockroach" comes from the fact that, when we see one cockroach, there are likely many more that we have not seen. In our case, when the agent finds a low-value box, it is likely that remaining boxes are also of low value.

a lower threshold under which she will also stop searching. This lower cutoff value is due to the cockroach effect, and typically appears when a poor realization signals a bad common attribute that convinces her to take the outside option instead of taking one more inspection.

Furthermore, the cockroach effect will also show up in her expected utility, and therefore leads to a new trade-off in determining the optimal search order. When the two boxes differ in the variance of the idiosyncratic attribute, conventional wisdom suggests that the agent always open the large-variance box first, as it maximizes the luck effect by making her more likely to achieve a good realization in the first shot. However, the new cockroach effect will push her to search the small-variance box first because it contains better information on the common attribute. When the first realization is bad, she can be more informative regarding the value of the second box, which helps her make a better decision on whether to open that box or not. Therefore, the optimal search order depends on the trade-off between these two opposite effects: when the agent already has good information on the common attribute (the variance of the common attribute is small), starting with the large-variance box is optimal for her; otherwise, the cockroach effect becomes important and makes the small-variance box more valuable to be opened first. The former case complements the existing literature that shows a riskier box is more favorable for the consumer (e.g., Weitzman, 1979; Armstrong, 2017).

We then extend our discussion into the case with more than two boxes and find that introducing interdependence brings two interesting insights that align with the reality. First, we find that the agent may return to previously sampled boxes even if she has not sampled all of them, which cannot be captured by the classic search theory. Experimentalists have provided evidence for frequent recall in consumer search behavior, and they account for this observation mainly by irrationality (e.g., Schotter and Braunstein, 1981; Kogut, 1990). This paper contributes in providing a theory-based argument that recall can happen when the agent finds the common attribute is too bad for her to inspect the remaining closed boxes during search. Second, we find that the optimal search order can be history-dependent. This result again comes from the trade-off between the luck effect and the cockroach effect. Depending on different previous realizations, the agent may either want to acquire more information on the common attribute and therefore sample a small-variance box, or to maximize the luck effect and sample the large-variance box next.

As far as we know, this paper is the first one that completely captures the optimal search strategy when values of items are interdependent. A closely related paper is Ke and Lin (2020), who also use a correlated search model to construct complementarity among competing products. The biggest difference is that they assume the consumer is able to distinguish the exact value of the common attribute after sampling a product, while the agent in our model can only observe the total match values. While they have a brief discussion on this situation in their extensions, we formally solve the model and generate a rich set of insights regarding the agent's search behaviors.

Related literature. Our paper is directly built on the literature with non-random sequential searches. The first-generation studies focus on the search behavior under the exogenous search order (e.g., Perry and Wigderson, 1986; Arbatskaya, 2007; Armstrong, Vickers, and Zhou, 2009; Haan and Moraga-González, 2011; Zhou, 2011). Starting from Weitzman (1979), the second-generation models discuss the optimal search strategy with the endogenous search order. For example, Doval (2018) considers the possibility that an agent can take a Pandora's Box without opening it; other papers investigate the consumer search with multi-dimensional attributes, partial information, and gradual learning (e.g., Branco, Sun, and Villas-Boas, 2012; Klabjan, Olszewski, and Wolinsky, 2014; Ke, Shen, and Villas-Boas, 2016; Ke, Tang, Villas-Boas, and Zhang, 2022). The framework of consumer search has also been widely implemented into the discussion of firms optimality (e.g., Diamond, 1971; Wolinsky, 1986; Anderson and Renault, 1999; Chen and Zhang, 2018; Ding and Zhang, 2018). Most recently, Armstrong (2017) shows that the Pandora's rule can be formulated as a static discrete choice problem. We contribute to this strand of literature by fully investigating searchers' behaviors when values of different items are interdependent.

Our work is also related to the literature on learning. For example, under the framework of multi-armed bandits problem, Bolton and Harris (1999) discusses the case where an agent can learn from the current experimentation of other agents, and the information is therefore a public good; Eeckhout and Weng (2015) investigates the experimentation with common values where agents can continue to learn about the same underlying state variable. Gonzalez and Shi (2010) examine labor market search where workers can learn about their job-finding probability from search.

The remainder of this paper is organized as follows. Section 2 sets up our baseline model. We first discuss the optimal stopping rule in Section 3, then we show the trade-off in the optimal search order in Section 4. Section 5 reports our results on recall and history-dependence. Finally, Section 6 concludes.

2 Model Setup

Environment. There are two closed boxes at the beginning of our scenario. Each box is indexed by $i \in \{1, 2\}$. The value of box i is:

$$V_i = X + E_i,$$

where $X \sim F_X(x)$ and $E_i \sim F_{E_i}(e_i)$ are the common and idiosyncratic attributes, respectively. The two boxes' values are therefore correlated through the common part X. We assume that random variables X, E_1 and E_2 are continuously defined on the entire real space.[2] Furthermore, we assume that they are independently distributed and their priors are common knowledge.

Search. To learn the value of a box, an agent has to sample it at a search cost s. We assume that sampling the first box is costless.[3] Note that, after inspecting box i, she can only learn its total value v_i without knowing the exact values for the common and idiosyncratic attributes.[4] The agent can choose the order for her sampling and can choose whether to quit the market at any stage of the search. For simplicity, we assume that the agent has zero reservation utility.[5] Finally, search is with perfect recall.

The search problem. The agent is faced with a dynamic optimization problem, which can be described in the following three stages.

1. In the first stage, the agent chooses a box $a \in \{1,2\}$ and sample it. Her expected utility is:

$$U = \max_{a\in\{1,2\}} \left\{ \mathbb{E}_{v_a}\left[U_1\left(v_a; a\right)\right] \right\}, \tag{1}$$

   where $U_1\left(v_a; a\right)$ is her expected utility after observing that the value of box a is v_a. This utility function is defined in stage 2.

2. In stage two, she has known the realization of box a is v_a. She can choose whether to sample the other box b at a cost s, or stop sampling. The problem in this stage can be written as:

$$U_1(v_a; a) = \max\left\{ \underbrace{\max\{0, v_a\}}_{\text{Stop}}, \underbrace{-s + \mathbb{E}_{V_b}\left[U_2\left(v_a, V_b\right) | V_a = v_a\right]}_{\text{Continue}} \right\}, \tag{2}$$

   where $U_2(v_a, v_b)$ is her utility when the values of both boxes are revealed in stage 3. Note that the expected utility of continue sampling is conditional on the first box's realization, because v_a contains information for the box b through the common value X. If she stops

[2]It is essential to assume those distributions are continuous. A discrete prior may lead to different belief updates by construction. We are only interested in the most general situation.

[3]This assumption is harmless and will make sure the agent is always willing to start searching. We are not interested in the case where she does not search at all.

[4]Throughout the paper, we use uppercase letters to represent random variables and the corresponding lowercase letters to represent their realizations.

[5]The zero outside option assumption is not crucial to our results. Conceptually, the two-box case with an exogenous outside option is equivalent to a three-box case with no exogenous outside option. When the agent has sampled the first box, that realization becomes the outside option for her to open the remaining two boxes. For simplicity, we assume there is an exogenous outside option and conduct our discussion within two boxes.

searching, then the game ends at this stage; otherwise, she will sample box b and enter stage 3.

3. In the last stage, the agent has observed the realized values of both boxes, v_a and v_b. She will therefore choose the best option, i.e.,

$$U_2(v_a, v_b) = \max\Big\{0, v_a, v_b\Big\}. \tag{3}$$

   Note that the search order does not matter for this stage because both boxes have been sampled.

3 Optimal Stopping Rule

First suppose that both the search order $\{a, b\}$ and the first box's value v_a are given. The search problem is described by the optimization problem (2), where the agent chooses between: (1) stop sampling; and (2) sample box b.

The realization of box a will influence the agent's belief on V_b. Before analyzing the search decison, it is helpful to introduce some properties of this belief updating. Intuitively, a better realization of v_a indicates better common value X as well as better private value E_a. These insights are shown as the standard assumption of Monotone Likelihood Ratio Property (MLRP) in Assumption 1.[6]

Assumption 1 *The densities $\{f_X(\cdot|V_i = v)\}$ and $\{f_{E_i}(\cdot|V_i = v)\}$ have the strict monotone likelihood ratio property.*

Problem (2) indicates that the agent will keep searching if and only if the marginal benefit exceeds the additional cost of search, i.e.,

$$g_a(v_a) := \mathbb{E}_{V_b}\Big[\max\{0, v_a, V_b\}\Big| V_a = v_a\Big] - \max\{v_a, 0\} > s. \tag{4}$$

The first realization v_a determines not only the posterior for the value of box b but also the outside option for sampling it. While the former, information effect is always there, v_a can influence the outside option only when it exceeds the reservation utility 0. Therefore, we formalize our discussion based on whether $v_a \geq 0$.

[6]For more discussion about the MLRP, please see Milgrom (1981).

Case I: $v_a \leq 0$**.** In this case, the value of box a is lower than the reservation utility, so the agent will never take this box. However, v_a can still influence her decision by shaping her belief on the value of box b. The equivalent condition (16) reduces to:

$$g_a(v_a) = \mathbb{E}_{V_b}\left[\max\{0, \underbrace{X + E_b}_{V_b}\}\Big| V_a = v_a\right] > s \quad , \quad \forall\, v_a \leq 0. \tag{5}$$

Here, the private value E_b is independent, so v_a enters only through the common value X. Based on strict MLRP, a larger realization of V_a is more favorable than a small one, which suggests that $g_a(v_a)$ is strictly increasing over v_a when $v_a \leq 0$. Therefore, the incremental utility $g_a(v_a)$ is maximized when $v_a = 0$, which implies the highest value for X. On the other hand, as the prior F_X is defined on the entire real space, a sufficiently low realization v_a can lead to a sufficiently bad posterior for the common value. These properties are summarized in the following equations:

$$0 = g_a(-\infty) < g_a\left(v_a'\right) < g_a\left(v_a''\right) \leq g_a(0) \quad , \quad \forall\, v_a' < v_a'' \leq 0. \tag{6}$$

As a result, when the search cost $s < g_a(0)$, there will be a unique threshold $\underline{v}_a < 0$ such that the agent will decide to stop sampling and take the reservation utility if and only if $v_a < \underline{v}_a$. This result corresponds to the intuition that a bad realization signals a poor common value, which may stop the agent from further searching. When the search becomes more costly, that is, $s \geq g_a(0)$, any negative realization will be bad enough to keep her out of the search market.

Case II: $v_a \geq 0$**.** An increase in v_a here will not only lead to a better posterior for X, but also increase the agent's outside option — the next box is valuable only if its value is greater than v_a. Under this scenario, condition (16) can be written into:

$$g_a(v_a) = \mathbb{E}_{E_b}\left[\max\{0, E_b - E_a\}\Big| V_a = v_a\right] > s \quad , \quad \forall\, v_a \geq 0. \tag{7}$$

Interestingly, the common attribute X is cancelled out in this continue-sampling condition. The intuition is that, when comparing values of the two boxes, only the idiosyncratic part matters. Again, E_b is independent of the first box, so v_a enters this condition only through the posterior of E_a. Therefore, based on the strict MLRP from Assumption 1, we learn that $g_a(v_a)$ is decreasing over v_a when $v_a \geq 0$, which leads to following conclusion:

$$g_a(0) \geq g_a\left(v_a'\right) > g_a\left(v_a''\right) > g_a(+\infty) = 0 \quad , \quad \forall\, v_a'' > v_a' \geq 0. \tag{8}$$

The lower bound is zero because it is nearly impossible to find a better box when the first realization is sufficiently good.

Therefore, when the search cost $s < g_a(0)$, there exists a unique cutoff value $\overline{v}_a > 0$ that makes the agent unwilling to continue sampling if and only if $v_a > \overline{v}_a$. In this case, a sufficiently high realization for the box a signals an extremely high idiosyncratic value E_a, so the agent does not want to spend additional search cost sampling the next box, which will be very likely worse than v_a. On the other hand, when $s \geq g_a(0)$, the search cost will be high enough to stop she sampling box b for any $u_a \geq 0$.

Combining the above two cases, we summarize the agent's optimal stopping rule in stage two in the following proposition.

Proposition 1 (Optimal stopping rule) *Given the realization of box a, v_a, the agent will continue sampling box b if and only if:*

1. $s < g_a(0)$*; and*
2. $\underline{v}_a \leq v_a \leq \overline{v}_a$*, where* $\underline{v}_a < 0$ *and* $\overline{v}_a > 0$ *are defined such that* $g_a(\underline{v}_a) = g_a(\overline{v}_a) = s$*.*

Otherwise, she will stop searching and take the outside option $\max\{0, u_a\}$*.*

To see the intuition more clearly, Figure 1a illustrates the expected utility of the agent. If she decides to stop searching, she will get the outside option $\max\{0, v_a\}$, which is depicted by the blue, solid line. In an extreme economy without search cost, it will always be lucrative for her to sample box b (the red, solid line). However, as the search becomes moderately frictional (the red, dashed line), the two cutoff values $\underline{v}_a < 0 < \overline{v}_a$ appear. The agent will stop sampling if the realization of box a indicates a sufficiently low common value (that is, when $v_a < \underline{v}_a$) or if it indicates a sufficiently high private value (when $v_a > \overline{v}_a$). Finally, if the search cost gets too high (the red, dotted line), she will quit searching regardless of the realization v_a. Figure 1b plots the incremental utility of continuing sampling, which is defined by $g_a(v_a)$ in equation (16). Clearly, this increment is largest when $v_a = 0$, whose value becomes the threshold for the search cost s.

We are interested in the case with active search. Therefore, we from now on impose a small search cost assumption based on Proposition 1.

Assumption 2 *Search cost* $s < \min\{g_1(0), g_2(0)\}$ *such that both boxes are possible to be sampled.*

The information value. In the absence of belief updating, Weitzman (1979) shows that the agent will continue searching if and only if the current realization is lower than an upper threshold.

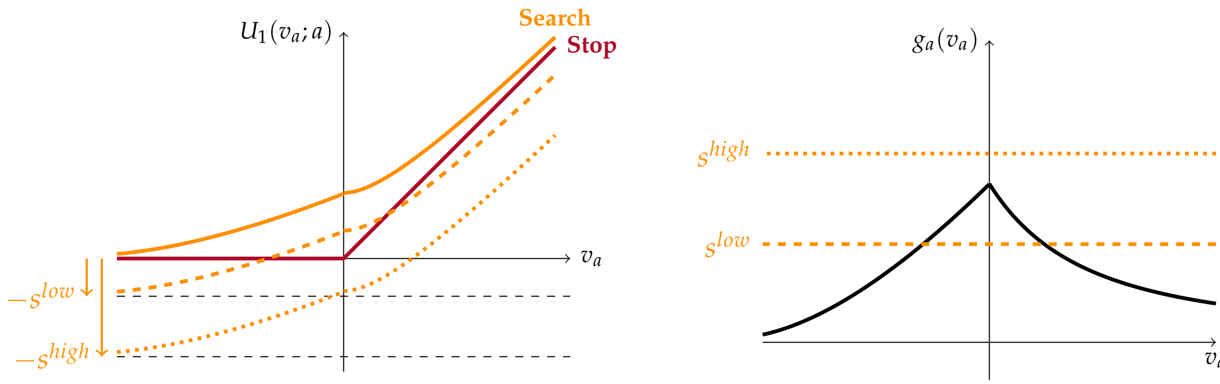

(a) Expected utility

(b) Increment utility from searching

Figure 1: Utility comparison after sampling the first box

When boxes are correlated, however, the existence of the lower threshold $\underline{v}_a$ in Proposition 1 demonstrates the importance of the information. In this section, we will briefly discuss this information value.

Imagine that there is a myopic agent (denoted by superscript $\mathcal{M}$) who fails to see the correlation between the two boxes. Given that the realization of box a is v_a, she will decide whether to sample the box b based on Weitzman's Pandora rule. That is, sample it if and only if $v_a < \overline{v}_a^{\mathcal{M}}$, where the "reservation price" $\overline{v}_a^{\mathcal{M}}$ is defined as:

$$\overline{v}_a^{\mathcal{M}} = \arg_z \left\{ \mathbb{E}_{V_b}\Big[\max\{z, V_b\} \Big] - z = s \right\}$$

$$\Leftrightarrow \quad \overline{v}_a^{\mathcal{M}} = \arg_z \left\{ \mathbb{E}_{E_b}\Big[\max\{0, E_b - (z - X)\} \Big] = s \right\},$$

which shares the same idea as the condition for $\overline{v}_a$ (i.e., Equation (7)) except for the fact that the myopic agent does not update her belief on X from the search history u_a.[7] Compared to the optimal stopping rule, we find that the myopic stopping rule deviates in both the upper and lower thresholds, which we summarize in the Proposition 2 below.

Proposition 2 (Information value) *Compared to the myopic stopping rule, the information value of the first realization u_a shows up in two ways:*

1. *a lower threshold $\underline{v}_a$ that informs the agent of a bad common value when $u_a < \underline{v}_a$; and*
2. *an upper threshold $\overline{v}_a$ that is adjusted by the belief updating process.*

[7] Assumption 2 ensures that the myopic agent will inspect the next box when $u_a = 0$. This property indicates that $\overline{v}_a^{\mathcal{M}} > 0$, which further means that $\overline{v}_a^{\mathcal{M}}$ is independent of the reservation utility 0.

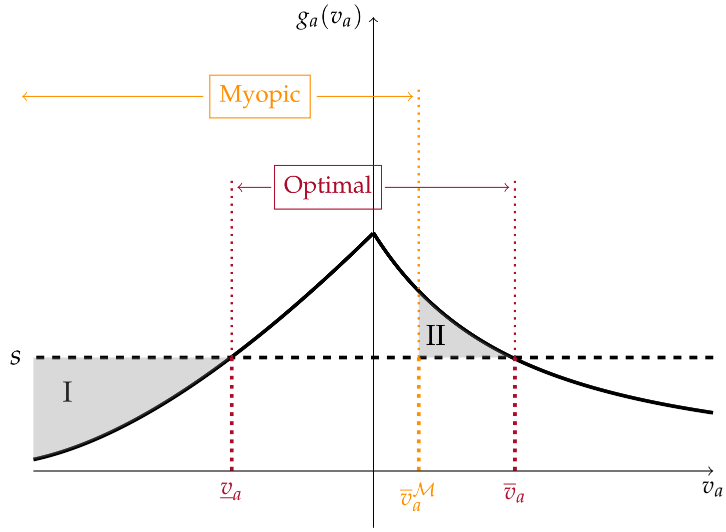

Figure 2: The information value in the optimal stopping rule

We interpret the difference in the expected utility between the optimal and myopic stopping rules as the information value of the search history, which is graphically shown by the two shaded areas in Figure 2. Part I corresponds to the first point in Proposition 2, that the agent can benefit from a bad realization which signals a low common value. On the other hand, part II comes from the agent's belief updating. Observing the value of box a, a rational agent can update her belief on the common attribute X, which further enables her to make a better decision in whether to continue searching box b.

4 Optimal Search Order

Given the optimal stopping rule in Proposition 1, we now discuss the search problem (1) in stage one, where the agent decides which box to sample first. The expected utility of sampling box a first can be written as:

$$\mathbb{E}_{v_a}[U_1(v_a;a)] = \underbrace{\mathbb{E}_{v_a}\Big[\max\{0, v_a\}\Big]}_{\text{Exp. value - box } a} + \underbrace{\int_{\underline{v}_a}^{\overline{v}_a} \Big[g_a(v_a) - s\Big] dF_{V_a}(v_a)}_{\text{Exp. incremental value - box } b}. \tag{9}$$

The total expected value consists of the expected utility from both boxes. The first part indicates the utility gain by sampling box a, while the second component, being consistent with Proposition 1, means the expected incremental utility due to the possibility of searching the second box b. Furthermore, defining the first-best benchmark as the case without search costs, we have following property:

Lemma 1 (Expected utility) *The expected utility* (9) *can be written into:*

$$\mathbb{E}_{v_a}[U_1(v_a;a)] = U^{fb} - \int_0^s \left[F_{V_a}(\overline{v}_a(t)) - F_{V_a}(\underline{v}_a(t))\right] dt, \tag{10}$$

where U^{fb} is the first-best expected utility where search is costless, and t represents for the search cost. Furthermore, the difference in expected utility between first sampling box 1 and 2, i.e., $\Delta U := \mathbb{E}_{v_2}[U_1(v_2;2)] - \mathbb{E}_{v_1}[U_1(v_1;1)]$, is:

$$\Delta U = \underbrace{\int_0^s \left[F_{V_1}(\overline{v}_1(t)) - F_{V_2}(\overline{v}_2(t))\right] dt}_{\text{Luck effect}} + \underbrace{\int_0^s \left[F_{V_2}(\underline{v}_2(t)) - F_{V_1}(\underline{v}_1(t))\right] dt}_{\text{Cockroach effect}}. \tag{11}$$

Proof. See Appendix A. ■

Lemma 1 expresses the agent's expected utility in terms of the difference to the first-best case. Note that this first-best utility is independent of search order, which therefore becomes a good reference point for comparing different orders. In addition, Proposition 1 shows that the two cutoff values of stopping searching depends on the search cost t. When t rises, the agent has to pay higher search cost for additional search in the case of $\underline{v}_a(t) < v_a < \overline{v}_a(t)$, which eventually becomes the disutility due to the search cost that is captured by the integral part in Equation (10).

Furthermore, Equation (11) decomposes the difference in expected utility incurred by different search orders into two components: the "luck" effect and the "cockroach" effect. The former luck effect is captured by the difference in the upper threshold $\overline{v}_a$, which means how good the first sampled box should be such that the agent can consider herself lucky enough to stop searching. On the contrary, the lower threshold captures the situation where a terrible realization suggests a bad common attribute that stops her from sampling the next one. We call this part the cockroach effect.[8] We will discuss the optimal search order based on this decomposition.

4.1 Parametric assumption

Lemma 1 suggests that the choice of search order is based on a (potential) trade-off between the luck effect and the cockroach effect. To make this comparison explicit, however, we have to impose stronger distributional assumptions in our model. Specifically, we assume that the values of both common and idiosyncratic attributes, X and E_i, are normally distributed. Denoting the mean as μ and variance as σ^2, we can write the distribution as $X \sim \mathcal{N}(\mu_X, \sigma_X^2)$ and $E_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$.

[8] This name comes from the fact that, when you see one cockroach, there are likely many more that you have not seen.

Lemma 2 (Belief updating with normality) *Given the value of box a as u_a, the agent will update her belief on the common value X into a normal posterior with*

$$X|(V_a = v_a) \sim \mathcal{N}\left(\frac{\sigma_X^2(v_a - \mu_a) + \sigma_a^2 \mu_c}{\sigma_X^2 + \sigma_a^2}, \frac{\sigma_X^2 \sigma_a^2}{\sigma_X^2 + \sigma_a^2}\right).$$

The belief updating under the normal distribution is presented in Lemma 2. A higher value of u_a will increase the mean of the posterior distribution of X, which dominates the one with lower u_a in the sense of FSD. Moreover, the variance of the updated distribution is irrelevant to the first realization. Sampling a small-variance box can bring the agent a more accurate information on X. As we will see, this insight will become a part of the trade-off for her in deciding the search order that is new to the literature.

Furthermore, the utility difference ΔU in Equation (11) can be rewritten as:

$$\Delta U = \int_0^s \left[\Phi\left(\frac{\overline{v}_1(t) - \mu_X - \mu_1}{\sqrt{\sigma_X^2 + \sigma_1^2}}\right) - \Phi\left(\frac{\overline{v}_2(t) - \mu_X - \mu_2}{\sqrt{\sigma_X^2 + \sigma_2^2}}\right)\right] dt + \int_0^s \left[\Phi\left(\frac{\underline{v}_2(t) - \mu_X - \mu_2}{\sqrt{\sigma_X^2 + \sigma_2^2}}\right) - \Phi\left(\frac{\underline{v}_1(t) - \mu_X - \mu_1}{\sqrt{\sigma_X^2 + \sigma_1^2}}\right)\right] dt, \quad (12)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. To make the comparison clear and complete, we will consider the case where the two boxes differ in idiosyncratic mean μ_i and variance σ_i in Section 4.2 and 4.3, respectively.

4.2 Boxes with different idiosyncratic mean

We start with a simple case where the two boxes only differ in the idiosyncratic mean. Without loss of generosity, we assume that $\mu_1 < \mu_2$ and $\sigma_1 = \sigma_2 = \sigma$. As the two boxes have the same variance, they contain the same amount of ex-ante information on the common attribute according to Lemma 2. Intuitively, sampling the high-expectation box 2 should be optimal. We confirm this intuition and show that doing so benefits the agent through both the luck and the cockroach channels.

Proposition 3 (Optimal search order with different means) *When the two boxes differ in private mean, first sampling the larger-mean box dominates the other search order in both luck and cockroach channels.*

Proof. See Appendix B. ■

First of all, the box with larger mean is more likely to induce a favorable realization, which therefore dominates in the luck effect. Moreover, although the two boxes contain the same level of information for X, the box with a more favorable idiosyncratic attribute can better compensate for the potential low common value. As a result, first sampling this box also dominates the other search order in terms of the cockroach effect.

4.3 Boxes with different idiosyncratic variance

We now investigate a more interesting case where the two boxes have different private variance. Without loss of generosity, assume that $\mu_1 = \mu_2 = 0$ and $\sigma_1 < \sigma_2$.[9] Conventional wisdom suggests the agent to search box 2 first because it is more likely to realize a good result. Our model covers this insight through the luck effect. However, a lower variance box can better inform the agent about the common value, which adds up to the trade-off in deciding the order of search.

Channel I. Luck Effect. We first discuss how the choice of the first-sampled box will influence the agent's expected utility through the luck channel. Intuitively, inspecting a large-variance box has a higher probability to realize a good result, which maximizes the luck effect by minimizing the cost to do an additional search. In mathematics, that intuition is equivalent to:

$$F_{V_1}(\overline{v}_1(t)) > F_{V_2}(\overline{v}_2(t)) \quad \Leftrightarrow \quad \frac{\overline{v}_1(t) - \mu_X}{\sqrt{\sigma_X^2 + \sigma_1^2}} > \frac{\overline{v}_2(t) - \mu_X}{\sqrt{\sigma_X^2 + \sigma_2^2}} \quad , \forall\, 0 < t < s. \tag{13}$$

In Proposition 4, we will confirm that first sampling box 2 strictly dominates the other order in the luck effect by showing Equation (13).

Channel II. Cockroach Effect. Nevertheless, in addition to the well-recognized benefit in the luck effect, we also discover a cost for prioritizing the riskier box, that is, the cockroach effect. When observing a low match value in the first search, a agent who has sampled the large-variance box cannot tell is it caused by a low common value, or it is it just because she is unlucky in the idiosyncratic attribute. On the contrary, if she has sampled a low-variance box, she can feel more certain in inferring a poor common attribute. Therefore, a less informed agent will keep sampling the next box unless the first realization is extremely bad, which induces additional search cost due

[9] It is harmless to assume $\mu_i = 0$ because the private mean can be absorbed by the common attribute.

to the lack in information. Mathematically, this claim corresponds to:

$$F_{V_1}(\underline{v}_1(t)) > F_{V_2}(\underline{v}_2(t)) \quad \Leftrightarrow \quad \frac{\underline{v}_1(t) - \mu_X}{\sqrt{\sigma_X^2 + \sigma_1^2}} > \frac{\underline{v}_2(t) - \mu_X}{\sqrt{\sigma_X^2 + \sigma_2^2}} \quad , \forall\, 0 < t < s, \tag{14}$$

which will be shown true when the mean of the common value μ_X is large enough.

Summary. In summary, the optimal search order here is determined by the trade-off between these two opposing effects. Among all the cases, the comparative static of the common attribute's variance, σ_X, is of our interest because it captures the relative importance of the information in search with common value. Proposition 4 summarizes the corresponding insights.

Proposition 4 (Optimal search order with different variances) *Suppose* $\min\{g_1(\mu_X), g_2(\mu_X)\} > s$ *and* $\mu_X > \sqrt{\frac{\sigma_X^2(\sigma_1^2+\sigma_2^2)+\sigma_1^2\sigma_2^2}{\sigma_X^2+\sigma_1^2}}$. *Starting searching with the large-variance box 2 is dominating in the luck effect, but dominated in the cockroach effect. The trade-off depends on the common variance* σ_X^2. *There exist two cutoff values* $\overline{\sigma}_X^2 \geq \underline{\sigma}_X^2$, *such that:*

1. *if* $\sigma_X^2 \geq \overline{\sigma}_X^2$, *the agent will sample box 2 first; and*
2. *if* $\sigma_X^2 \leq \underline{\sigma}_X^2$, *the agent will sample box 1 first.*

Proof. See Appendix C. ■

Specifically, in an extreme case with $\sigma_X^2 = 0$, the agent has complete information on the common attribute X. In this example, the equilibrium regresses to Weitzman (1979) where only the luck effect matters. Therefore, first searching the large-variance box 2 is her strictly dominating strategy. On the other hand, when σ_X^2 is large enough, the difference in luck effect will be trivial because the large variance in common value dominates the small difference that lies in the idiosyncratic attributes. In this case, the information (cockroach) effect will become dominating and will make it lucrative for the agent to sample box 1 first.

5 Extensions with more boxes

This section further extends our discussion into a more general case with more-than-two boxes. We find two new insights that are complementary to the literature: (1) recall without sampling all boxes; and (2) history-dependent search order.

5.1 Recall without sampling all goods

Recall, a common strategy for consumers in the real world, is usually absent in the framework of traditional consumer search literature. It only appears as an extreme case when the consumer has sampled all products and then chooses the best one among her searching history. By implementing the interdependent value, we complement the literature with the finding that the agent can recall for a previously sampled one even if she has *not* inspected all the boxes.

Consider a special case with N boxes whose idiosyncratic attributes are independently drawn from the same distribution F_E.[10] As boxes are ex-ante identical, the agent only needs to decide when to stop sampling new boxes. Without loss of generality, we use the subscript $i \in \{1, 2, ..., N\}$ to index the ith box to be opened. The search problem can be summarized as follows: in each stage $k + 1 \leq N$, she has observed the sampling history $\boldsymbol{\vartheta}_k := \{v_1, ..., v_k\}$ and is deciding whether to open the box $k + 1$. Denote $U(\boldsymbol{\vartheta})$ as the agent's expected utility given previous realizations $\boldsymbol{\vartheta}$, the Bellman equation can be written as:

$$U(\boldsymbol{\vartheta}_k) = \max \left\{ \underbrace{\max \left\{0, \boldsymbol{\vartheta}_k\right\}}_{\text{Outside option}}, \underbrace{\mathbb{E}_{V_{k+1}}\left[U\left(\{\boldsymbol{\vartheta}_k, V_{k+1}\}\right)\middle|\boldsymbol{\vartheta}_k\right] - s}_{\text{Expected utility from continuing}} \right\}. \tag{15}$$

Again, the Bellman equation (15) indicates that she will keep sampling if and only if the marginal benefit exceeds the additional cost of search, i.e.,

$$g_k(\boldsymbol{\vartheta}_k) := \mathbb{E}_{V_{k+1}}\left[U\left(\{\boldsymbol{\vartheta}_k, V_{k+1}\}\right)\middle|\boldsymbol{\vartheta}_k\right] - \max\left\{0, \boldsymbol{\vartheta}_k\right\} > s. \tag{16}$$

The same mechanisms leading to Proposition 1 in the duopoly case apply here as well. When the latest sampled box k is not her outside option (i.e., $v_k < \max\{0, \boldsymbol{\vartheta}_k\}$), an increase in v_k will raise the agent's expectation for closed boxes due to the strict MRPL, and the incremental utility $g_k(\boldsymbol{\vartheta}_k)$ is therefore increasing in v_k. However, when $v_k > \max\{0, \boldsymbol{\vartheta}_k\}$, box k becomes her outside option and the expected incremental utility is determined by the idiosyncratic attributes alone. A higher v_k now indicates a better idiosyncratic attribute for the outside option, which will lead to a lower $g_k(\boldsymbol{\vartheta}_k)$. Given this analysis, the following proposition summarizes her optimal search strategy in this N-box case.

Proposition 5 (Optimal search strategy with N identical boxes) *Given search history $\boldsymbol{\vartheta}_k$ with any $k \in \{0, 1, ..., N-1\}$, the agent will continue searching if and only if:*

[10] We assume all boxes are ex-ante identical to avoid discussion on search order, which can be excessively complicated according to Section 5.2.

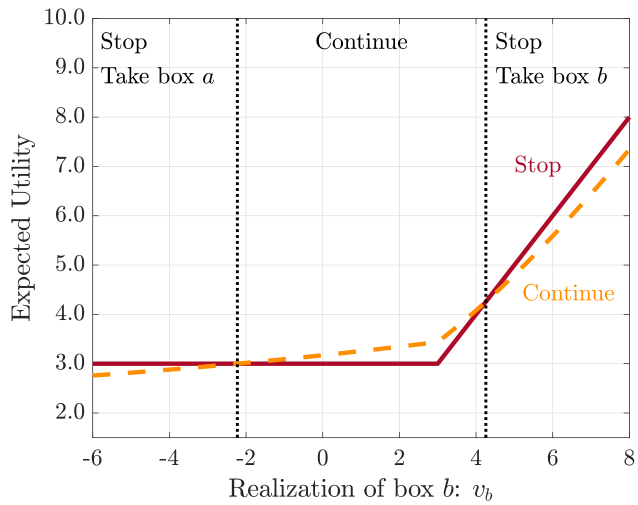

Figure 3: Three-box example: optimal decision after sampling two boxes

Notes: Denote a, b, and c as the first, second, and last box in the agent's exogenous search order. In this simulation, we set parameters at $\mu_X = 2.5$, $\sigma_X = 5$, $\{\sigma_a, \sigma_b, \sigma_c\} = \{1, 2, 3\}$, $\mu_a = \mu_b = \mu_c = 0$, and $s = 0.8$. The realization of the first-sampled box is assumed to be 3. The red, solid line indicates the expected utility when she decides to stop sampling after the second box b, while the orange, dashed line indicates the expected utility when she deicides to continue sampling the last box c. The two thresholds are shown by the two black, dotted lines.

1. *search cost is small:* $s < g_k\left(\left\{\boldsymbol{\vartheta}_{k-1}, \max\{0, \boldsymbol{\vartheta}_{k-1}\}\right\}\right)$*; and*

2. $\underline{v}_k \leq v_k \leq \overline{v}_k$*, where* $\underline{v}_k$ *and* $\overline{v}_k$ *are defined such that* $g_k\left(\{\boldsymbol{\vartheta}_{k-1}, \underline{v}_k\}\right) = g_k\left(\{\boldsymbol{\vartheta}_{k-1}, \overline{v}_k\}\right) = s.$

Otherwise, she will stop sampling and take the best box that has been opened (or nothing), which may need to recall the previous boxes.

The optimal search strategy includes the possibility of recall even before sampling all the boxes. The intuition is again the two cutoff values for the optimal stopping rule. Given the observed value of previously sampled boxes, if the realization for the current box is too low, the agent will form a belief that the common attribute is too bad for her to keep sampling. As a result, she will stop opening new boxes and return to the best one (or the outside option) that she has sampled. On the other hand, if this realization is good enough, she will stop sampling and directly take this box, which is consistent with the existing literature.

Recall can also occur in a more general setting with heterogenous boxes. Figure 3 illustrates this insight in a three-box example. Let the search order be $\{a, b, c\}$. Given that the realization of the first box $v_a = 3$, the utility of stopping is $\max\{3, v_b\}$ (the red, solid line). Meanwhile, the realization of box b will also impact the expected utility from continuing sampling via belief updating, which is shown by the orange, dashed line. The two thresholds divide the agent's best response into three cases: (1) when $v_b < \underline{v}_b$, she will stop sampling and return to take the first

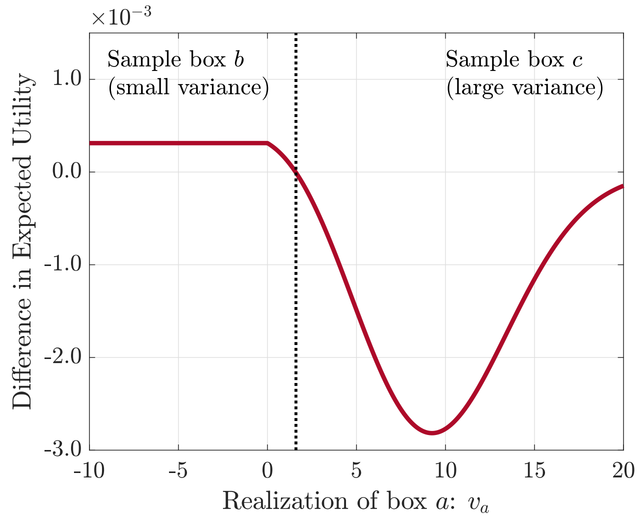

Figure 4: Three-box example: optimal search order as a function of the first realization

Notes: Denote a, b, and c as the first, second, and last box in the agent's exogenous search order. In this simulation, we set parameters at $\mu_X = 2.5$, $\sigma_X = 5$, $\sigma_a \to +\infty$, $\sigma_b = 1$, $\sigma_c = 2$, $\mu_a = \mu_b = \mu_c = 0$, and $s = 0.1$. The red, solid line indicates the difference in expected utility between the two search orders that sample box b or c next, as a function of the realization v_a. The black, dotted line indicates the v_a that makes the agent indifferent between choosing the two boxes as the next one to inspect.

box a; (2) when $\underline{v}_b \leq v_b \leq \overline{v}_b$, she will continue sampling; and (3) when $v_b \geq \overline{v}_b$, she will stop sampling and take the second box b.

5.2 History-dependent search order

In the three-box model, we also find that the agent's optimal search order depends on her search history. Depending on the realization of the first inspection, she may prefer to sample a smaller-variance box next for better information, or a larger-variance box for a better chance of realizing a high match value. This finding complements the traditional literature where consumers usually search with a fixed order.

Proposition 6 (History-dependence) *The optimal search order is history-dependent.*

To see this case, we propose a special example where the first box a's private variance $\sigma_a \to +\infty$. Because the value of box a is extremely noisy, the first inspection is not informative on the value of the common attribute. Typically, if $v_a < 0$, her optimization problem will reduce to the two-box model and, according to Proposition 4, she will sample the smaller-variance box when the ex-ante variance of the common value is larger enough. However, when $v_a > 0$, a better realization of box a will increase the outside option of the agent, which motivates her to take more risk by sampling the larger-variance box next. Therefore, the agent will choose the lower-variance box to search when v_a is small, and the higher-variance one when v_a is large, as is shown by Figure 4. This example hence suggests a history-dependent search order in common value search.

6 Conclusion

In this paper, we propose a search framework with interdependent values. Introducing this interdependence adds up to each inspection an additional value of information, which reshapes the agent's optimal search strategy. We show that there exists a new, lower threshold for her to stop sampling, and the optimal search order now depends on the trade-off between the traditional luck effect and the cockroach effect. Furthermore, we show that this framework can account for many new observations in the real world, including recalling previously sampled items as well as the history-dependent search order.

Although our paper is built upon the standard framework of consumer search, our insights can be generalized to all kinds of search problem. In the labor market, interviewers will update their belief on other interviewees graduating from the same college after an interview; an investor can know more about the aggregate shock in the economy after checking some financial assets, which will change his/her future behavior in the search. The key idea is that agents, by searching one item, can learn about the values of other items. This feature leads to the cockroach effect and therefore determines their optimal search strategies.

Due to tractability, our discussion is mainly restricted to a two-box theory and some simulation results. Similar to the correlated bandits problem, a more general model is notoriously hard to solve, but it is desirable for completing this discussion and further testing the theory by the data.

References

ANDERSON, S. P. AND R. RENAULT (1999): "Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model," The RAND Journal of Economics, 719–735.

ARBATSKAYA, M. (2007): "Ordered search," The RAND Journal of Economics, 38, 119–126.

ARMSTRONG, M. (2017): "Ordered consumer search," Journal of the European Economic Association, 15, 989–1024.

ARMSTRONG, M., J. VICKERS, AND J. ZHOU (2009): "Prominence and consumer search," The RAND Journal of Economics, 40, 209–233.

BOLTON, P. AND C. HARRIS (1999): "Strategic experimentation," Econometrica, 67, 349–374.

BRANCO, F., M. SUN, AND J. M. VILLAS-BOAS (2012): "Optimal search for product information," Management Science, 58, 2037–2056.

CHEN, Y. AND T. ZHANG (2018): "Entry and welfare in search markets," The Economic Journal, 128, 55–80.

DIAMOND, P. A. (1971): "A model of price adjustment," Journal of economic theory, 3, 156–168.

DING, Y. AND T. ZHANG (2018): "Price-directed consumer search," International Journal of Industrial Organization, 58, 106–135.

DOVAL, L. (2018): "Whether or not to open Pandora's box," Journal of Economic Theory, 175, 127–158.

EECKHOUT, J. AND X. WENG (2015): "Common value experimentation," Journal of Economic Theory, 160, 317–339.

GONZALEZ, F. M. AND S. SHI (2010): "An equilibrium theory of learning, search, and wages," Econometrica, 78, 509–537.

HAAN, M. A. AND J. L. MORAGA-GONZÁLEZ (2011): "Advertising for attention in a consumer search model," The Economic Journal, 121, 552–579.

KE, T. T. AND S. LIN (2020): "Informational complementarity," Management Science, 66, 3699–3716.

KE, T. T., Z.-J. M. SHEN, AND J. M. VILLAS-BOAS (2016): "Search for information on multiple products," Management Science, 62, 3576–3603.

KE, T. T., W. TANG, J. M. VILLAS-BOAS, AND Y. P. ZHANG (2022): "Parallel search for information in continuous time—optimal stopping and geometry of the pde," Applied Mathematics & Optimization, 85, 1–25.

KLABJAN, D., W. OLSZEWSKI, AND A. WOLINSKY (2014): "Attributes," Games and Economic Behavior, 88, 190–206.

KOGUT, C. A. (1990): "Consumer search behavior and sunk costs," Journal of Economic Behavior & Organization, 14, 381–392.

MILGROM, P. R. (1981): "Good news and bad news: Representation theorems and applications," The Bell Journal of Economics, 380–391.

PERRY, M. AND A. WIGDERSON (1986): "Search in a known pattern," Journal of Political Economy, 94, 225–230.

SCHOTTER, A. AND Y. M. BRAUNSTEIN (1981): "Economic search: an experimental study," Economic inquiry, 19, 1–25.

WEITZMAN, M. L. (1979): "Optimal Search for The Best Alternative," Econometrica, 641–654.

WOLINSKY, A. (1986): "True Monopolistic Competition as A Result of Imperfect Information," The Quarterly Journal of Economics, 101, 493–511.

ZHOU, J. (2011): "Ordered search in differentiated markets," International Journal of Industrial Organization, 29, 253–262.

Appendix

Appendix A Proof of Lemma 1

We prove this lemma by showing that:

$$\frac{\partial}{\partial t}\Big[\mathbb{E}_{v_a}[U_1(v_a;a)]\Big] = -\Big[F_{V_a}(\overline{v}_a(t)) - F_{V_a}(\underline{v}_a(t))\Big]. \tag{A.1}$$

Recall from Equation (9) that:

$$\mathbb{E}_{v_a}[U_1(v_a;a)] = \mathbb{E}_{v_a}\Big[\max\{0, v_a\}\Big] + \int_{\underline{v}_a}^{\overline{v}_a}\Big[g_a(v_a) - s\Big]\,dF_{V_a}(v_a). \tag{A.2}$$

Taking the differentiation of $\mathbb{E}_{v_a}[U_1(v_a;a)]$ regarding search cost t, we get:

$$\begin{aligned}
\frac{\partial}{\partial t}\Big[\mathbb{E}_{v_a}[U_1(v_a;a)]\Big] &= \frac{\partial}{\partial t}\overline{v}_a \cdot \Big[g_a(\overline{v}_a) - s\Big] - \frac{\partial}{\partial t}\underline{v}_a \cdot \Big[g_a(\underline{v}_a) - s\Big] + \int_{\underline{v}_a}^{\overline{v}_a}(-1)\,dF_{V_a}(v_a) \\
&= \frac{\partial}{\partial t}\overline{v}_a \times 0 - \frac{\partial}{\partial t}\underline{v}_a \times 0 - \Big[F_{V_a}(\overline{v}_a(t)) - F_{V_a}(\underline{v}_a(t))\Big] \\
&= -\Big[F_{V_a}(\overline{v}_a(t)) - F_{V_a}(\underline{v}_a(t))\Big].
\end{aligned} \tag{A.3}$$

Integrating Equation (A.3) gives us Equation (10) in Lemma 1.

Appendix B Proof of Proposition 3

Case I. Cockroach Effect. We first show that first sampling box 2 is optimal for the cockroach effect. Based on Equation (12), when $\sigma_1 = \sigma_2$, this statement is equivalent to the claim that $\underline{v}_2(t) - \mu_2 > \underline{v}_1(t) - \mu_1$ for any given search cost $t > 0$.

Given $\underline{v}_2 < 0$ solves $g_2(v_2) = t$, we have:

$$\begin{aligned}
g_1(\underline{v}_2 - \mu_2 + \mu_1) &= \mathbb{E}_{V_2}\Big[\max\{0, X + E_2\}\Big| V_1 = \underline{v}_2 - \mu_2 + \mu_1\Big] \\
&= \mathbb{E}_{V_2}\Big[\max\{0, X + (E_2 - \mu_2) + \mu_2\}\Big| V_1 - \mu_1 = \underline{v}_2 - \mu_2\Big] \\
&= \mathbb{E}_{V_2}\Big[\max\{0, X + (E_1 - \mu_1) + \mu_2\}\Big| V_2 - \mu_2 = \underline{v}_2 - \mu_2\Big] && \text{(B.1)} \\
&> \mathbb{E}_{V_2}\Big[\max\{0, X + E_1\}\Big| V_2 = \underline{v}_2\Big] \\
&= t, && \text{(B.2)}
\end{aligned}$$

where (B.1) is due to the fact that $V_1 - \mu_1$ and $V_2 - \mu_2$ have the same prior, and so do $E_1 - \mu_1$ and $E_2 - \mu_2$. The intuition is that, after adjusting by the idiosyncratic mean, the two boxes have the same ex-ante distribution, and thus have the same amount of information. Furthermore, when $v_1 < 0$, $g_1(v_1)$ is increasing over v_1 according to Equation (6). Combining with inequality (B.2), this monotonicity gives that:

$$\underline{v}_1 < \underline{v}_2 - \mu_2 + \mu_1 \quad \Leftrightarrow \quad \underline{v}_2 - \mu_2 > \underline{v}_1 - \mu_1. \tag{B.3}$$

Case II. Luck Effect. To show that first sampling box 2 also dominates in the luck effect, we need to prove $\overline{v}_1(t) - \mu_1 > \overline{v}_2(t) - \mu_2$ for any given $t > 0$. Similarly, given $\overline{v}_1 > 0$ that solves $g_1(\overline{v}_1) = t$, we have:

$$\begin{aligned}
g_2(\overline{v}_1 - \mu_1 + \mu_2) &= \mathbb{E}_{E_1}\left[\max\{0, E_2 - E_1\}\middle| V_2 = \underline{v}_1 - \mu_1 + \mu_2\right] \\
&= \mathbb{E}_{E_1}\left[\max\{0, (E_2 - \mu_2) - (E_1 - \mu_1) + (\mu_2 - \mu_1)\}\middle| V_2 - \mu_2 = \underline{v}_1 - \mu_1\right] \\
&= \mathbb{E}_{E_2}\left[\max\{0, (E_1 - \mu_1) - (E_2 - \mu_2) + (\mu_2 - \mu_1)\}\middle| V_1 - \mu_1 = \underline{v}_1 - \mu_1\right] && \text{(B.4)} \\
&> \mathbb{E}_{E_2}\left[\max\{0, E_1 - E_2\}\middle| V_1 = \overline{v}_1\right] \\
&= t. && \text{(B.5)}
\end{aligned}$$

As Equation (8) shows that $g_2(v_2)$ is declining on v_2 when $v_2 > 0$, we know that:

$$\overline{v}_2 < \overline{v}_1 - \mu_1 + \mu_2 \quad \Leftrightarrow \quad \overline{v}_1 - \mu_1 > \overline{v}_2 - \mu_2. \tag{B.6}$$

Appendix C Proof of Proposition 4

Recall that in this case we have $\mu_1 = \mu_2 = 0$. For transparency, we will divide the proof into different claims and prove them step by step. We first rewrite the optimal stopping rule by implementing the normal distribution assumption.

Appendix C.1 Optimal Stopping Rule

Suppose the agent first samples box a. Following Weitzman (1979), we first define the reservation price z_a as:

$$z_a(v_a) := \arg\left\{ s = \int_{z_a}^{+\infty} (v_b - z_a)\, dF_{V_b|V_a=v_a}(v_b) \right\}, \tag{C.1}$$

where $F_{V_b|V_a=z_a}(\cdot)$ is the posterior of V_b after seeing $V_a = v_a$. Different from Weitzman (1979), the reservation price here is endogenous on v_a because of the history-dependent belief updating. Note that from Lemma 2, we can write Equation (C.1) into:

$$s = \int_{z_a - \frac{\sigma_X^2 (v_a - \mu_X)}{\sigma_X^2 + \sigma_a^2}}^{+\infty} \left(v_b - \left[z - \frac{\sigma_X^2 (v_a - \mu_X)}{\sigma_X^2 + \sigma_a^2} \right] \right) d\Phi \left(\frac{v_b - \mu_X}{\sqrt{\frac{\sigma_X^2 \left(\sigma_a^2 + \sigma_b^2\right) + \sigma_a^2 \sigma_b^2}{\sigma_X^2 + \sigma_a^2}}} \right). \tag{C.2}$$

Defining $\widetilde{z}_a := z_a - \frac{\sigma_X^2 (v_a - \mu_X)}{\sigma_X^2 + \sigma_a^2}$, we find that $\widetilde{z}_a$ is history-independent, that is, can be pinned down by Equation (C.2) regardless of v_a. Therefore, we can write $z_a = \widetilde{z}_a + \frac{\sigma_X^2 (v_a - \mu_X)}{\sigma_X^2 + \sigma_a^2}$. Furthermore, the assumption that $g_a(\mu_X) > s$ gives:

$$\int_{\mu_X}^{+\infty} (v_b - \mu_X)\, dF_{V_b|V_a=\mu_X}(v_b) > s$$

$$\int_{\mu_X}^{+\infty} (v_b - \mu_X)\, d\Phi \left(\frac{v_b - \mu_X}{\sqrt{\frac{\sigma_X^2 \left(\sigma_a^2 + \sigma_b^2\right) + \sigma_a^2 \sigma_b^2}{\sigma_X^2 + \sigma_a^2}}} \right) > s$$

$$\widetilde{z}_a > \mu_X \tag{C.3}$$

She will continue sampling box b if and only if the expected gain exceeds search costs, which is equivalent to:

$$\max\{0, v_a\} \leq \widetilde{z}_a + \frac{\sigma_X^2 (v_a - \mu_X)}{\sigma_X^2 + \sigma_a^2}. \tag{C.4}$$

Similar to Section 3, she will continue sampling iff $\underline{v}_a \leq v_a \leq \overline{v}_a$. We discuss two possible cases based on the realization v_a.

When $v_a \leq 0$. In this situation, Condition (C.4) can be rewritten as:

$$0 \leq \widetilde{z}_a + \frac{\sigma_X^2 (v_a - \mu_X)}{\sigma_X^2 + \sigma_a^2}, \tag{C.5}$$

which gives us the lower threshold:

$$\underline{v}_a = \mu_X - \left(1 + \frac{\sigma_a^2}{\sigma_X^2}\right) \widetilde{z}_a < 0. \tag{C.6}$$

When $v_a \geq 0$. In this case, the condition becomes:

$$v_a \leq \widetilde{z}_a + \frac{\sigma_X^2 (v_a - \mu_X)}{\sigma_X^2 + \sigma_a^2}, \tag{C.7}$$

which induces a upper threshold:

$$\overline{v}_a = \left(1 + \frac{\sigma_X^2}{\sigma_a^2}\right) \widetilde{z}_a - \frac{\sigma_X^2}{\sigma_a^2} \mu_X > 0. \tag{C.8}$$

Appendix C.2 Optimal Search Order

We now investigate the optimal search order based on the $\Delta U := \mathbb{E}_{v_2}[U_1(v_2; 2)] - \mathbb{E}_{v_1}[U_1(v_1; 1)]$ in Equation (12). Plugging the $\underline{v}_a$ and $\overline{v}_a$ in, we get:

$$\begin{aligned} \Delta U = & \overbrace{\int_0^s \left[\Phi\left(\frac{\sqrt{\sigma_X^2 + \sigma_1^2}}{\sigma_1^2} \left(\widetilde{z}_1(t) - \mu_X\right) \right) - \Phi\left(\frac{\sqrt{\sigma_X^2 + \sigma_2^2}}{\sigma_2^2} \left(\widetilde{z}_2(t) - \mu_X\right) \right) \right] dt}^{\Delta \text{ Luck Effect}} \\ & + \underbrace{\int_0^s \left[\Phi\left(-\frac{\sqrt{\sigma_X^2 + \sigma_2^2}}{\sigma_X^2} \widetilde{z}_2(t) \right) - \Phi\left(-\frac{\sqrt{\sigma_X^2 + \sigma_1^2}}{\sigma_X^2} \widetilde{z}_1(t) \right) \right] dt}_{\Delta \text{ Cockroach Effect}}. \end{aligned} \tag{C.9}$$

We will first discuss these two effects separately, then formalize the trade-off between them.

Luck Effect. We first prove that searching box 2 first dominates in terms of the luck effect, i.e., in Equation (C.9) we have Δ Luck Effect > 0. Notice that the definition of $\widetilde{z}_a$, Equation (C.2), can be further written into:

$$\sqrt{\sigma_a^2 + \sigma_X^2} s = \int_{\sqrt{\sigma_a^2 + \sigma_X^2}(\widetilde{z}_a - \mu_X)}^{+\infty} \left[x - \sqrt{\sigma_a^2 + \sigma_X^2} \left(\widetilde{z}_a - \mu_X\right) \right] d\Phi \left(\frac{x}{\sqrt{\sigma_X^2 \left(\sigma_a^2 + \sigma_b^2\right) + \sigma_a^2 \sigma_b^2}} \right). \tag{C.10}$$

Since $\sigma_2^2 > \sigma_1^2$, Equation (C.10) induces:

$$\begin{aligned} & \sqrt{\sigma_1^2 + \sigma_X^2} \left(\widetilde{z}_1 - \mu_X\right) > \sqrt{\sigma_2^2 + \sigma_X^2} \left(\widetilde{z}_2 - \mu_X\right) \\ \Rightarrow \quad & \frac{\sqrt{\sigma_1^2 + \sigma_X^2}}{\sigma_1^2} \left(\widetilde{z}_1 - \mu_X\right) > \frac{\sqrt{\sigma_2^2 + \sigma_X^2}}{\sigma_2^2} \left(\widetilde{z}_2 - \mu_X\right) \end{aligned} \tag{C.11}$$

for any given search cost t. Combining with the fact that $\widetilde{z}_a > \mu_X, \forall a$, we derive that

$$\int_0^s \left[\Phi\left(\frac{\sqrt{\sigma_X^2+\sigma_1^2}}{\sigma_1^2} \left(\widetilde{z}_1(t) - \mu_X\right) \right) - \Phi\left(\frac{\sqrt{\sigma_X^2+\sigma_2^2}}{\sigma_2^2} \left(\widetilde{z}_2(t) - \mu_X\right) \right) \right] dt > 0.$$

Cockroach Effect. Next, we want to show that sampling box 1 first is favorable in terms of the cockroach effect. We need to compare $\sqrt{\sigma_X^2+\sigma_a^2}\widetilde{z}_a$ in this case. From Equation (C.10), we have:

$$\left(\sqrt{\sigma_2^2+\sigma_X^2} - \sqrt{\sigma_1^2+\sigma_X^2}\right) s = \left[\sqrt{\sigma_1^2+\sigma_X^2}\left(\widetilde{z}_1 - \mu_X\right) - \sqrt{\sigma_2^2+\sigma_X^2}\left(\widetilde{z}_2 - \mu_X\right)\right] \left[1 - \Phi\left(\frac{\sqrt{\sigma_1^2+\sigma_X^2}\left(\widetilde{z}_1 - \mu_X\right)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right) + \sigma_1^2\sigma_2^2}} \right)\right]$$
$$+ \int_{\sqrt{\sigma_2^2+\sigma_X^2}(\widetilde{z}_2 - \mu_X)}^{\sqrt{\sigma_1^2+\sigma_X^2}(\widetilde{z}_1 - \mu_X)} \left[x - \sqrt{\sigma_2^2+\sigma_X^2}\left(\widetilde{z}_2 - \mu_X\right)\right] d\Phi\left(\frac{x}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right) + \sigma_1^2\sigma_2^2}} \right),$$

which gives us:

$$\sqrt{\sigma_X^2+\sigma_2^2}\widetilde{z}_2 - \sqrt{\sigma_X^2+\sigma_1^2}\widetilde{z}_1 = \frac{\Omega}{1 - \Phi\left(\frac{\sqrt{\sigma_1^2+\sigma_X^2}(\widetilde{z}_1 - \mu_X)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right) + \sigma_1^2\sigma_2^2}} \right)}, \tag{C.12}$$

where

$$\Omega = \left(\sqrt{\sigma_2^2+\sigma_X^2} - \sqrt{\sigma_1^2+\sigma_X^2}\right) \left\{ \mu_X \left[1 - \Phi\left(\frac{\sqrt{\sigma_1^2+\sigma_X^2}\left(\widetilde{z}_1 - \mu_X\right)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right) + \sigma_1^2\sigma_2^2}} \right)\right] - s \right\}$$
$$+ \int_{\sqrt{\sigma_2^2+\sigma_X^2}(\widetilde{z}_2 - \mu_X)}^{\sqrt{\sigma_1^2+\sigma_X^2}(\widetilde{z}_1 - \mu_X)} \left[x - \sqrt{\sigma_2^2+\sigma_X^2}\left(\widetilde{z}_2 - \mu_X\right)\right] d\Phi\left(\frac{x}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right) + \sigma_1^2\sigma_2^2}} \right). \tag{C.13}$$

Therefore, we only need to show $\Omega > 0$. Because $\sqrt{\sigma_1^2+\sigma_X^2}\,(\tilde{z}_1-\mu_X) > \sqrt{\sigma_2^2+\sigma_X^2}\,(\tilde{z}_2-\mu_X)$, a sufficient condition is:

$$
\begin{aligned}
& \mu_X > \frac{s}{1-\Phi\left(\frac{\sqrt{\sigma_1^2+\sigma_X^2}(\tilde{z}_1-\mu_X)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}}\right)} \\
\Leftrightarrow \quad & \mu_X > \frac{\int_0^{+\infty} x d\Phi\left(\frac{\sqrt{\sigma_1^2+\sigma_X^2}(x+\tilde{z}_1-\mu_X)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}}\right)}{1-\Phi\left(\frac{\sqrt{\sigma_1^2+\sigma_X^2}(\tilde{z}_1-\mu_X)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}}\right)} \\
\Leftrightarrow \quad & \mu_X > \mathbb{E}[R|R>0], \quad \text{where} \quad R \sim \mathcal{N}\left(\mu_X-\tilde{z}_1, \sqrt{\frac{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}{\sigma_1^2+\sigma_X^2}}\right).
\end{aligned}
$$

Furthermore, the truncated conditional expectation of normal distribution R can be computed as:

$$
\begin{aligned}
\mathbb{E}[R|R>0] &= \mu_X-\tilde{z}_1+\sqrt{\frac{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}{\sigma_1^2+\sigma_X^2}}\overbrace{\left[\frac{\phi\left(\frac{\sqrt{\sigma_1^2+\sigma_X^2}(\tilde{z}_1-\mu_X)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}}\right)}{1-\Phi\left(\frac{\sqrt{\sigma_1^2+\sigma_X^2}(\tilde{z}_1-\mu_X)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}}\right)}\right]}^{\text{Hazard ratio}} \\
&< \mu_X-\tilde{z}_1+\sqrt{\frac{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}{\sigma_1^2+\sigma_X^2}}\left[\frac{\sqrt{\sigma_1^2+\sigma_X^2}\,(\tilde{z}_1-\mu_X)}{\sqrt{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}}+1\right] \\
&= \sqrt{\frac{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}{\sigma_1^2+\sigma_X^2}}.
\end{aligned}
\tag{C.14}
$$

Therefore, the assumption that $\mu_X > \sqrt{\frac{\sigma_X^2\left(\sigma_1^2+\sigma_2^2\right)+\sigma_1^2\sigma_2^2}{\sigma_1^2+\sigma_X^2}}$ makes sure $\sqrt{\sigma_X^2+\sigma_2^2}\tilde{z}_2 > \sqrt{\sigma_X^2+\sigma_1^2}\tilde{z}_1$, i.e., box 1 dominates in the cockroach effect.

Trade-off. Based on the discussion above, we know that the optimal search order depends on the trade-off between these two opposing effects. When σ_X^2 is small enough, the difference in the luck effect will be dominating, so sampling box 1 first is optimal. The exact point is showed in Weitzman (1979), so we skip the proof here. When $\sigma_X^2 \to +\infty$, however, the cockroach effect becomes dominating and sampling box 2 is optimal.

To see this, recall that:

$$\lim_{\sigma_c \to +\infty} s = \lim_{\sigma_c \to +\infty} \int_{\widetilde{z}_a}^{+\infty} (x - \widetilde{z}_a) \, d\Phi \left(\frac{x - \mu_X}{\sqrt{\frac{\sigma_X^2 \left(\sigma_1^2 + \sigma_2^2\right) + \sigma_1^2 \sigma_2^2}{\sigma_a^2 + \sigma_X^2}}} \right)$$

$$s = \int_{\widetilde{z}_a}^{+\infty} (x - \widetilde{z}_a) \, d\Phi \left(\frac{x - \mu_X}{\sqrt{\sigma_1^2 + \sigma_2^2}} \right) \tag{C.15}$$

which gives us $\lim_{\sigma_X \to +\infty} \widetilde{z}_a = \widetilde{z}, \forall a$. Therefore, we have:

$$\lim_{\sigma_X^2 \to +\infty} \Phi \left(\frac{\sqrt{\sigma_X^2 + \sigma_a^2}}{\sigma_a^2} \left(\widetilde{z}_a - \mu_X \right) \right) = \Phi(+\infty) \quad \text{and} \quad \lim_{\sigma_X^2 \to +\infty} \Phi \left(- \frac{\sqrt{\sigma_X^2 + \sigma_a^2}}{\sigma_X^2} \widetilde{z}_a \right) = \Phi(0).$$

Due to the exponentiation of $\Phi(\cdot)$, we know that the difference in the luck effect is a higher-order infinitesimal of the difference in the cockroach effect, i.e., the latter effect dominates the former one.