Searching for “Arms”: Experimentation with Endogenous Consideration Sets

Supplement

Daniel Fershtman **Alessandro Pavan**

September 2024

Abstract

This document contains additional material. All sections, conditions, and results specific to this document have the suffix “S” to avoid confusion with the corresponding parts in the main text. Section S.1 establishes the optimality of the index policy claimed in Part (1) of Theorem 1 in the main text by means of a novel proof that exploits the recursive characterization of the search index and the categorization of the alternatives. Section S.2 contains the proof of Parts (2) and (3) of Proposition 2 in the main text. Section S.3 shows how to use Theorem 1 and Proposition 1 in the main body to arrive at a solution to the extension of Weitzman's Pandora's boxes problem with an endogenous CS. Section S.4 contains various results for the application to online consumer search. In particular, Subsection S.4.1 extends the eventual purchase theorem to a setting with an endogenous CS, Subsection S.4.2 characterizes the click-through-rates (CTRs), Subsection S.4.3 provides an example in which the CTRs, the values-per-click (VPCs), and the purchasing probabilities are derived, whereas Subsection S.4.4 shows why additional ad space may be detrimental to firms' profits. Section S.5 shows how the characterization of the optimal policy extends to certain problems with irreversible choice. Finally, Section S.6 discusses why an index policy need not be optimal in the presence of “meta” arms with associated super-processes.

S.1 Proof of Part (1) of Theorem 1 in the main text.

The proof exploits the recursive representation of the search index established in Part (2) of Theorem 1, along with the representation of the DM's payoff under the index rule established in Part (3) of Theorem 1 and an appropriate description of the state space, to verify that the DM's payoff under the index policy satisfies the Bellman equation of the corresponding dynamic program.

Proof strategy. The proof is in two steps. Step 1 uses the representation of the DM's payoff under the index rule established in Part (3) of Theorem 1 in the main text to characterize how much the DM obtains from following the index policy χ^* from the outset rather than being forced to make a different decision in the first period and then reverting to χ^* from the next period onward. Step 2 then uses the results in step 1 to establish the optimality of χ^* through dynamic programming.

Step 1. In the analysis below, we find it useful to describe changes in the composition of the CS, the evolution of the search technology, as well as all information acquired about the alternatives, entirely in terms of transitions between states. Rather than keeping track of the collection of kernels $G_\xi(\vartheta^m; \mu)$ describing the conditional distributions from which the marginal signals ϑ_{m+1} are drawn, we describe directly the evolution of each alternative's state ω^P as follows. When the DM explores an alternative currently in state ω^P, its new state $\tilde{\omega}^P$ is drawn from a distribution $H_{\omega^P} \in \Delta(\Omega^P)$ that is invariant to time.[1] When the DM explores a different alternative, or expands the CS, the alternative currently in state ω^P remains in the same state with certainty at the beginning of the next period. Similarly, each time search is conducted, given the current state of the search technology ω^S, the new state of the search technology $\tilde{\omega}^S$ is drawn from a distribution $H_{\omega^S} \in \Delta(\Omega^S)$.The distributions H_{ω^s} are time-homogeneous (i.e., the evolution of the search technology depends on past search outcomes but is invariant in calendar time), and the outcome of each new search is drawn from H_{ω^S} independently from the idiosyncratic and time-varying component θ of each alternative in the CS.

Abusing notation, then denote the state of the decision problem by a function $\mathcal{S} : \Omega \to \mathbb{N}$ that specifies, for each $\omega \in \Omega$, including $\omega \in \Omega^S$, the number of alternatives, including the search technology, that are in state ω.[2] Given this notation, for any pair of states $\mathcal{S}'$ and $\mathcal{S}''$ then define $\mathcal{S}' \vee \mathcal{S}'' \equiv (\mathcal{S}'(\omega) + \mathcal{S}''(\omega) : \omega \in \Omega)$ and $\mathcal{S}' \backslash \mathcal{S}'' \equiv (\max\{\mathcal{S}'(\omega) - \mathcal{S}''(\omega), 0\} : \omega \in \Omega)$. Any feasible state of the decision problem must specify one, and only one, state of the search technology (i.e., one state $\hat{\omega}^S$ for which $\mathcal{S}(\hat{\omega}^S) = 1$ and such that $\mathcal{S}(\omega^S) = 0$ for all $\omega^S \neq \hat{\omega}^s$). However, it will be convenient to consider fictitious (infeasible) states where search is not possible, as well as fictitious states with multiple search possibilities. If the state of the decision problem is such that either (i) the CS is empty, or (ii) there is a single alternative in the CS and the latter cannot be expanded, we will denote such a state by $e(\omega)$, where $\omega \in \Omega$ is the state of the search technology in case (i) and of the single physical alternative in case (ii).[3]

Lemma S.1. *For any $v \in \mathbb{R}$ and states $\mathcal{S}'$ and $\mathcal{S}''$, $\kappa(v|\mathcal{S}' \vee \mathcal{S}'') = \kappa(v|\mathcal{S}') + \kappa(v|\mathcal{S}'')$.*

Proof of Lemma S.1. The result follows from the fact that the state of each alternative that is not explored in a given period remains unchanged, along with the fact that the time-varying components θ of the various alternatives evolve independently of one another and of the state of the search technology, given the alternatives' categories ξ. Similarly, the state of the search technology remains unchanged in periods in which search is not conducted, and evolves independently of the

[1] Clearly, because each alternative's category ξ is fixed, given the current state $\omega^P = (\xi, \theta)$, the distribution H_{ω^P} assigns probability one to states whose category is ξ and whose signal history $\vartheta^{m+1} = (\vartheta^m, \vartheta_{m+1})$ is a "follower" of ϑ^m, meaning that it is obtained by adding a new signal realization ϑ_{m+1} to the history ϑ^m.

[2] Clearly, with this representation, there is a unique $\hat{\omega}^s \in \Omega^S$ such that $\mathcal{S}(\omega^S) = 1$ if $\omega^s = \hat{\omega}^s$ and $\mathcal{S}(\omega^S) = 0$ if $\omega^s \neq \hat{\omega}^s$. The special case where the DM does not have the option to search corresponds to the case where for all $\omega^S \in \Omega^S$, $\mathcal{S}(\omega^S) = 0$.

[3] Throughout the analysis below, we maintain the assumption that an outside option with value equal to zero is available to the DM. However, to avoid possible confusion, here we do not explicitly treat the outside option as a separate alternative.

time-varying component θ in the state of each existing alternative, given the alternatives' categories ξ. Furthermore, the index of each alternative is a function only of the alternative's state, and the index of search is a function only of the state of the search technology. Therefore, all indexes evolve independently of one another (conditional on the alternatives' categories), and evolve only when their corresponding decision (search or exploration of an alternative) is chosen. Since the decisions are taken under the index policy χ^*, the result follows from the fact that, starting from any state $\mathcal{S}$, the *total* time it takes to bring *all* indexes (that is, those of the alternatives in the CS as well as the index of search) below any value v is the sum (across alternatives in the CS and search) of the *individual* times necessary to bring each index below v in isolation. $\square$

Given the initial state $\mathcal{S}_0$, for any $\omega^P \in \{\hat{\omega}^P \in \Omega^P : \mathcal{S}_0^P(\hat{\omega}^P) > 0\}$, denote by $\mathbb{E}\left[u|\omega^P\right]$ the immediate expected payoff from exploring an alternative in state ω^P and by $\tilde{\omega}^P$ the new state of that alternative triggered by its exploration (drawn from H_{ω^P}). Let

$$V^P(\omega^P|\mathcal{S}_0) \equiv (1-\delta)\mathbb{E}\left[u|\omega^P\right] + \delta\mathbb{E}^{\chi^*}\left[\mathcal{V}\left(\mathcal{S}_0 \backslash e(\omega^P) \vee e(\tilde{\omega}^P)\right)|\omega^P\right] \tag{S.1}$$

denote the DM's payoff from starting with exploring an alternative in state ω^P and then following the index policy χ^* from the next period onward. Similarly, let

$$V^S(\omega^S|\mathcal{S}_0) \equiv -(1-\delta)\mathbb{E}\left[c|\omega^S\right] + \delta\mathbb{E}^{\chi^*}\left[\mathcal{V}\left(S_0 \backslash e(\omega^S) \vee e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S)\right)|\omega^S\right] \tag{S.2}$$

denote the DM's payoff from expanding the CS when the state of search is ω^S, and then following the index policy χ^* from the next period onward, where $\mathbb{E}\left[c|\omega^S\right]$ is the immediate expected cost from searching (when the state of the search technology is ω^S), $\tilde{\omega}^S$ is the new state of the search technology, and $W^P(\tilde{\omega}^S)$ is the state of the new alternatives brought to the CS by the current search, with c and $W^P(\tilde{\omega}^S)$ jointly drawn from the distribution H_{ω^S}.[4]

We introduce a fictitious "auxiliary option" which is available at all periods and yields a constant reward $M < \infty$ when chosen. Denote the state corresponding to this fictitious auxiliary option by ω_M^A, and enlarge Ω^P to include ω_M^A. Similarly, let $e(\omega_M^A)$ denote the state of the problem in which only the auxiliary option with fixed reward M is available. Since the payoff from the auxiliary option is constant at M, if $v \geq M$, then $\kappa(v|\mathcal{S}_0 \vee e(\omega_M^A)) = \kappa(v|\mathcal{S}_0)$, whereas if $v < M$, then $\kappa(v|\mathcal{S}_0 \vee e(\omega_M^A)) = \infty$. Hence, the representation of the DM's payoff under the index policy in Part (3) of Theorem 1 in the main text, adapted to the fictitious environment that includes the auxiliary option, implies that

$$\begin{aligned}\mathcal{V}(\mathcal{S}_0 \vee e(\omega_M^A)) &= \int_0^\infty \left(1 - \mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|\mathcal{S}_0 \vee e(\omega_M^A)\right]\right) \mathrm{d}v = M + \int_M^\infty \left(1 - \mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|\mathcal{S}_0\right]\right) \mathrm{d}v \\ &= \mathcal{V}(\mathcal{S}_0) + \int_0^M \mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|\mathcal{S}_0\right] \mathrm{d}v.\end{aligned} \tag{S.3}$$

The definition of χ^*, along with Conditions (S.1) and (S.2), then imply the following:

[4] Note that $W^P(\tilde{\omega}^S)$ is a deterministic function of the new state $\tilde{\omega}^S$ of the search technology. To see this, recall that, for any $m \in \mathbb{N}$, the function E_m in the definition of the state of the search technology counts how many alternatives of each possible state ω^P have been added to the CS, as a result of the m-th search.

Lemma S.2. *For any* (ω^S, ω^P, M),

$$\mathcal{V}(e(\omega^S) \vee e(\omega_M^A)) = \begin{cases} V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A)) & \text{if } M \leq \mathcal{I}^S(\omega^S) \\ M > V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A)) & \text{if } M > \mathcal{I}^S(\omega^S) \end{cases} \tag{S.4}$$

$$\mathcal{V}(e(\omega^P) \vee e(\omega_M^A)) = \begin{cases} V^P(\omega^P|e(\omega^P) \vee e(\omega_M^A)) & \text{if } M \leq \mathcal{I}^P(\omega^P) \\ M > V^P(\omega^P|e(\omega^P) \vee e(\omega_M^A)) & \text{if } M > \mathcal{I}^P(\omega^P). \end{cases} \tag{S.5}$$

Proof of Lemma S.2. First note that the index corresponding to the auxiliary option is equal to M. Hence, if $M \leq \mathcal{I}^S(\omega^S)$, given $e(\omega^S) \vee e(\omega_M^A)$, χ^* prescribes to start with search, implying that $\mathcal{V}(e(\omega^S) \vee e(\omega_M^A)) = V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A))$. If, instead, $M > \mathcal{I}^S(\omega^S)$, χ^* prescribes to select the auxiliary option forever, with an expected (per period) payoff of M. To see why, in this case, $M > V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A))$, observe that the payoff $V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A))$ from starting with search and then following χ^* in each subsequent period is equal to $V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A)) = \mathbb{E}_{>1}^{\chi^*}\left[(1-\delta)\sum_{s=0}^{\bar{\tau}-1}\delta^s U_s + \delta^{\bar{\tau}} M|\omega^S\right]$, where $\bar{\tau}$ is the first time at which the index of search and of all the alternatives brought to the CS by search fall weakly below M, and where the expectation is under the process that obtains starting from $e(\omega^S) \vee e(\omega_M^A)$ by searching in the first period and then following the index policy in each subsequent period (the notation $\mathbb{E}_{>1}^{\chi^*}[\cdot]$ is meant to highlight that the expectation is under such a process). This follows from the fact that, once the DM, under χ^*, opts for the auxiliary option, he will continue to select that option in all subsequent periods. By definition of $\mathcal{I}^S(\omega^S)$,

$$M > \mathcal{I}^S(\omega^S) \equiv \sup_{\pi,\tau} \frac{\mathbb{E}^\pi\left[\sum_{s=0}^{\tau-1}\delta^s U_s|\omega^S\right]}{\mathbb{E}^\pi\left[\sum_{s=0}^{\tau-1}\delta^s|\omega^S\right]} \geq \frac{\mathbb{E}_{>1}^{\chi^*}\left[\sum_{s=0}^{\bar{\tau}-1}\delta^s U_s|\omega^S\right]}{\mathbb{E}_{>1}^{\chi^*}\left[\sum_{s=0}^{\bar{\tau}-1}\delta^s|\omega^S\right]}.$$

Rearranging, $M\mathbb{E}_{>1}^{\chi^*}\left[\sum_{s=0}^{\bar{\tau}-1}\delta^s|\omega^s\right] > \mathbb{E}_{>1}^{\chi^*}\left[\sum_{s=0}^{\bar{\tau}-1}\delta^s U_s|\omega^S\right]$. Therefore,

$$\mathbb{E}_{>1}^{\chi^*}\left[(1-\delta)\sum_{s=0}^{\bar{\tau}-1}\delta^s U_s + \delta^{\bar{\tau}} M|\omega^S\right] < M\mathbb{E}_{>1}^{\chi^*}\left[(1-\delta)\sum_{s=0}^{\bar{\tau}-1}\delta^s + \delta^{\bar{\tau}}|\omega^S\right] = M.$$

Similar arguments establish Condition (S.5). □

Next, for any initial state $\mathcal{S}_0$ of the decision problem, and any state $\omega^P \in \{\hat{\omega}^P \in \Omega^P : \mathcal{S}_0(\hat{\omega}^P) > 0\}$ of the alternatives in the CS corresponding to $\mathcal{S}_0$, let $D^P(\omega^P|\mathcal{S}_0) \equiv \mathcal{V}(\mathcal{S}_0) - V^P(\omega^P|\mathcal{S}_0)$ denote the payoff differential between (a) starting by following the index rule χ^* right away and (b) exploring first one of the alternatives in state ω^P and then following χ^* thereafter. Similarly, let $D^S(\omega^S|\mathcal{S}_0) \equiv \mathcal{V}(\mathcal{S}_0) - V^S(\omega^S|\mathcal{S}_0)$ denote the payoff differential between (c) starting with χ^* and (d) starting with search in state ω^S and then following χ^*. The next lemma relates these payoff differentials to the corresponding ones in a fictitious environment with the auxiliary option.[5]

Lemma S.3. *Let* $\mathcal{S}_0$ *be the initial state of the decision problem, with* $\omega^S \in \Omega^S$ *denoting the state of*

[5]In the statement of the lemma, $\mathcal{S}_0 \setminus e(\omega^S)$ is the state of a fictitious problem where search is not possible, whereas $\mathcal{S}_0^P \setminus e(\omega^P)$ is the state of the CS obtained from $\mathcal{S}_0^P$ by subtracting an alternative in state ω^P.

the search technology, as specified in $\mathcal{S}_0$. We have that[6]

$$D^S(\omega^S|\mathcal{S}_0) = \int_0^{\mathcal{I}^*(\mathcal{S}_0^P)} D^S(\omega^S|e(\omega^S) \vee e(\omega_v^A)) d\mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|\mathcal{S}_0 \setminus e(\omega^S)\right] \\ + \mathbb{E}^{\chi^*}\left[\delta^{\kappa(0)}|S_0 \backslash e(\omega^S))\right] D^S(\omega^S|e(\omega^S) \vee e(\omega_0^A)). \tag{S.6}$$

Similarly, for any alternative in the CS in state $\omega^P \in \{\hat{\omega}^P \in \Omega^P : \mathcal{S}_0^P(\hat{\omega}^P) > 0\}$,

$$D^P(\omega^P|\mathcal{S}_0) = \int_0^{max\{\mathcal{I}^*(\mathcal{S}_0^P \backslash e(\omega^P)), \mathcal{I}^S(\omega^S)\}} D^P(\omega^P|e(\omega^P) \vee e(\omega_v^A)) d\mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|\mathcal{S}_0 \setminus e(\omega^P)\right] \\ + \mathbb{E}^{\chi^*}\left[\delta^{\kappa(0)}|S_0 \backslash e(\omega^P))\right] D^P(\omega^P|e(\omega^P) \vee e(\omega_0^A)). \tag{S.7}$$

Proof of Lemma S.3. Using Condition (S.3), we have that, given the state $\mathcal{S}_0 \vee e(\omega_M^A)$ of the decision problem, and $\omega^S \in \Omega^S$,

$$D^S(\omega^S|\mathcal{S}_0 \vee e(\omega_M^A)) = \mathcal{V}(\mathcal{S}_0) + \int_0^M \mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|\mathcal{S}_0\right] \mathrm{d}v + (1-\delta)\mathbb{E}\left[c|\omega^S\right] \\ - \delta\mathbb{E}^{\chi^*}\left[\mathcal{V}(S_0 \backslash e(\omega^S) \vee e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S)) + \int_0^M \mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|S_0 \backslash e(\omega^S) \vee e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S))\right] \mathrm{d}v|\omega^S\right], \tag{S.8}$$

where the equality follows from combining (S.2) with (S.3). Similarly,

$$D^S(\omega^S|e(\omega^S) \vee e(\omega_M^A)) = \mathcal{V}(e(\omega^S)) + \int_0^M \mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|e(\omega^S)\right] \mathrm{d}v + (1-\delta)\mathbb{E}\left[c|\omega^S\right] \\ -\delta\mathbb{E}^{\chi^*}\left[\mathcal{V}(e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S)) + \int_0^M \mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S))\right] \mathrm{d}v|\omega^S\right]. \tag{S.9}$$

Differentiating (S.8) and (S.9) with respect to M, using the independence across alternatives and search and Lemma S.1, we have that

$$\frac{\partial}{\partial M} D^S(\omega^S|\mathcal{S}_0 \vee e(\omega_M^A)) = \mathbb{E}^{\chi^*}\left[\delta^{\kappa(M)}|S_0 \backslash e(\omega^S)\right] \frac{\partial}{\partial M} D^S(\omega^S|e(\omega^S) \vee e(\omega_M^A)). \tag{S.10}$$

That is, the improvement in $D^S(\omega^S|\mathcal{S}_0 \vee e(\omega_M^A))$ that originates from a slight increase in the value of the auxiliary option M is the same as in a setting with only search and the auxiliary option, $D^S(\omega^S|e(\omega^S) \vee e(\omega_M^A))$, discounted by the expected time it takes (under the index rule χ^*) until there are no indexes with value strictly higher than M, in an environment without search where the CS is the same as the one specified in $\mathcal{S}_0$. Similar arguments imply that, for any $\omega^P \in \{\hat{\omega}^P \in \Omega^P : \mathcal{S}_0(\hat{\omega}^P) > 0\}$,

$$\frac{\partial}{\partial M} D^P(\omega^P|\mathcal{S}_0 \vee e(\omega_M^A)) = \mathbb{E}^{\chi^*}\left[\delta^{\kappa(M)}|S_0 \backslash e(\omega^P)\right] \frac{\partial}{\partial M} D^P(\omega^P|e(\omega^P) \vee e(\omega_M^A)). \tag{S.11}$$

Let $M^* \equiv \max\{\mathcal{I}^*(\mathcal{S}_0^P), \mathcal{I}^S(\omega^S)\}$. Integrating (S.10) over the interval $(0, M^*)$ of possible values

[6] Recall that $\mathcal{I}^*(\mathcal{S}_0^P)$ is the largest index of the alternatives in the CS under the state $\mathcal{S}_0$.

for the auxiliary option and rearranging, we have that

$$\begin{aligned}
D^S(\omega^S|\mathcal{S}_0 \vee e(\omega_0^A)) &= D^S(\omega^S|\mathcal{S}_0 \vee e(\omega_{M^*}^A)) - \int_0^{M^*} \mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|S_0 \backslash e(\omega^S))\right] \frac{\partial}{\partial v} D^S(\omega^S|e(\omega^S) \vee e(\omega_v^A)) \mathrm{d}v \\
&= D^S(\omega^S|\mathcal{S}_0 \vee e(\omega_{M^*}^A)) - D^S(\omega^S|e(\omega^S) \vee e(\omega_{M^*}^A)) \\
&\quad + \mathbb{E}^{\chi^*}\left[\delta^{\kappa(0)}|S_0 \backslash e(\omega^S))\right] D^S(\omega^S|e(\omega^S) \vee e(\omega_0^A)) \\
&\quad + \int_0^{M^*} D^S(\omega^S|e(\omega^S) \vee e(\omega_v^A)) \mathrm{d}\mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|S_0 \backslash e(\omega^S))\right],
\end{aligned}$$

where the second equality follows from integration by parts and from the fact that

$$\mathbb{E}^{\chi^*}\left[\delta^{\kappa(M^*)}|S_0 \backslash e(\omega^S)\right] = 1.$$

That the outside option has value normalized to zero also implies that $D^S(\omega^S|\mathcal{S}_0 \vee e(\omega_0^A)) = D^S(\omega^S|\mathcal{S}_0)$. It is also easily verified that $D^S(\omega^S|\mathcal{S}_0 \vee e(\omega_{M^*}^A)) = D^S(\omega^S|e(\omega^S) \vee e(\omega_{M^*}^A))$.[7] Therefore, we have that

$$\begin{aligned}
D^S(\omega^S|\mathcal{S}_0) = &\int_0^{M^*} D^S(\omega^S|e(\omega^S) \vee e(\omega_v^A)) \mathrm{d}\mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|S_0 \backslash e(\omega^S)\right] \\
&+ \mathbb{E}^{\chi^*}\left[\delta^{\kappa(0)}|S_0 \backslash e(\omega^S))\right] D^S(\omega^S|e(\omega^S) \vee e(\omega_0^A)).
\end{aligned} \tag{S.12}$$

Similar arguments imply that

$$\begin{aligned}
D^P(\omega^P|\mathcal{S}_0) = &\int_0^{M^*} D^P(\omega^P|e(\omega^P) \vee e(\omega_v^A)) \mathrm{d}\mathbb{E}^{\chi^*}\left[\delta^{\kappa(v)}|S_0 \backslash e(\omega^P)\right] \\
&+ \mathbb{E}^{\chi^*}\left[\delta^{\kappa(0)}|S_0 \backslash e(\omega^P))\right] D^P(\omega^P|e(\omega^P) \vee e(\omega_0^A)).
\end{aligned} \tag{S.13}$$

To complete the proof of Lemma S.3, we consider separately two cases. Case (1): given $\mathcal{S}_0$, χ^* specifies starting by exploring a physical alternative (i.e., $M^* = \mathcal{I}^*(\mathcal{S}_0^P)$). Then Condition (S.6) in the lemma follows directly from (S.12). Thus consider Condition (S.7). First observe that, for any state $\omega^P \in \Omega^P$ such that $M^* > \max\{\mathcal{I}^*(\mathcal{S}_0^P \backslash e(\omega^P)), \mathcal{I}^S(\omega^S)\}$, we have that $M^* = \mathcal{I}^P(\omega^P)$, in which case $D^P(\omega^P|\mathcal{S}_0) = D^P(\omega^P|e(\omega^P) \vee e(\omega_0^A)) = 0$ and the integrand $D^P(\omega^P|e(\omega^P) \vee e(\omega_v^A))$ in (S.13) is equal to zero over the interval $[0, \mathcal{I}^P(\omega^P)]$ and hence also over the interval $[0, \max\{\mathcal{I}^*(\mathcal{S}_0^P \backslash e(\omega^P)), \mathcal{I}^S(\omega^S)\}]$. We thus have that, in this case, Condition (S.7) clearly holds. Next observe that, for any state $\omega^P \in \Omega^P$ such that $M^* = \max\{\mathcal{I}^*(\mathcal{S}_0^P \backslash e(\omega^P)), \mathcal{I}^S(\omega^S)\}$, Condition (S.7) follows directly from (S.13).

Case (2): given $\mathcal{S}_0$, χ^* specifies starting with search (i.e., $M^* = \mathcal{I}^S(\omega^S)$). Then, for any $\omega^P \in \Omega^P$, $\max\{\mathcal{I}^*(\mathcal{S}_0^P \backslash e(\omega^P)), \mathcal{I}^S(\omega^S)\} = M^*$, in which case Condition (S.7) in the lemma follows directly from (S.13). That Condition (S.6) also holds follows from the fact that, in this case, $D^S(\omega^S|\mathcal{S}_0) = D^S(\omega^S|e(\omega^S) \vee e(\omega_0^A)) = 0$ and the integrand $D^S(\omega^S|e(\omega^S) \vee e(\omega_v^A))$ in (S.12) is equal

[7] This follows immediately from the observation that $\mathcal{V}(\mathcal{S}_0 \vee e(\omega_{M^*}^A)) = \mathcal{V}(e(\omega^S) \vee e(\omega_{M^*}^A)) = M^*$, and similarly $\mathbb{E}^{\chi^*}\left[\mathcal{V}\left(\mathcal{S}_0 \backslash e(\omega^S) \vee e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S) \vee e(\omega_{M^*}^A)\right)|\omega^S\right] = \mathbb{E}^{\chi^*}\left[\mathcal{V}\left(e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S) \vee e(\omega_{M^*}^A)\right)|\omega^S\right]$. Intuitively, under the index policy, any alternative with index strictly below M^* is never explored given the presence of an auxiliary alternative with payoff M^*.

to zero over the entire interval $\left[0, \max\{\mathcal{I}^*(\mathcal{S}_0^P \backslash e(\omega^P)), \mathcal{I}^S(\omega^S)\}\right]$. $\square$

Step 2. Using the characterization of the payoff differentials in Lemma S.3, we now establish that the average per-period payoff under χ^* solves the Bellman equation for our dynamic optimization problem. Let $\mathcal{V}^*(\mathcal{S}_0) \equiv (1-\delta)\sup_{\chi \in \mathcal{X}} \mathbb{E}^\chi \left[\sum_{t=0}^\infty \delta^t U_t | \mathcal{S}_0\right]$ denote the value function for the dynamic optimization problem.

Lemma S.4. *For any state of the decision problem $\mathcal{S}_0$, with ω^S denoting the state of the search technology as specified under $\mathcal{S}_0$,*

1. *$\mathcal{V}(\mathcal{S}_0) \geq V^S(\omega^S|\mathcal{S}_0)$, and $\mathcal{V}(\mathcal{S}_0) = V^S(\omega^S|\mathcal{S}_0)$ if and only if $\mathcal{I}^S(\omega^S) \geq \mathcal{I}^*(\mathcal{S}_0^P)$;*
2. *for any $\omega^P \in \{\hat{\omega}^P \in \Omega^P : \mathcal{S}_0(\hat{\omega}^P) > 0\}$, $\mathcal{V}(\mathcal{S}_0) \geq V^P(\omega^P|\mathcal{S}_0)$, and $\mathcal{V}(\mathcal{S}_0) = V^P(\omega^P|\mathcal{S}_0)$ if and only if $\mathcal{I}^P(\omega^P) = \mathcal{I}^*(\mathcal{S}_0^P) \geq \mathcal{I}^S(\omega^S)$.*

Hence, for any $\mathcal{S}_0$, $\mathcal{V}(\mathcal{S}_0) = \mathcal{V}^*(\mathcal{S}_0)$, and χ^* is optimal.

Proof of Lemma S.4. *Part 1*. First, use (S.4) to note that, for all $v \geq 0$, $D^S(\omega^S|e(\omega^S) \vee e(\omega_v^A)) \geq 0$, with the inequality holding as an equality if and only $v \leq \mathcal{I}^S(\omega^S)$. Therefore, from (S.6), $D^S(\omega^S|\mathcal{S}_0) \geq 0$ – and hence $\mathcal{V}(\mathcal{S}_0) \geq V^S(\omega^S|\mathcal{S}_0)$ – with the inequality holding as an equality if and only if $\mathcal{I}^*(\mathcal{S}_0^P) \leq \mathcal{I}^S(\omega^S)$.

Part 2. Similarly, use (S.5) to observe that for any $\omega^P \in \{\hat{\omega}^P \in \Omega^P : \mathcal{S}_0^P(\hat{\omega}^P) > 0\}$ and any $v \geq 0$, $D^P(\omega^P|e(\omega^P) \vee e(\omega_v^A)) \geq 0$, with the inequality holding as an equality if and only if $0 \leq v \leq \mathcal{I}^P(\omega^P)$. Therefore, from (S.7), $D^P(\omega^P|\mathcal{S}_0) \geq 0$ with the inequality holding as equality if and only if $\mathcal{I}^P(\omega^P) \geq \max\{\mathcal{I}^*(\mathcal{S}_0^P \backslash e(\omega^P)), \mathcal{I}^S(\omega^S)\}$. The result in part 2 of the lemma then follows from the fact that the last inequality holds if and only if $\mathcal{I}^P(\omega^P) = \mathcal{I}^*(\mathcal{S}_0^P) \geq \mathcal{I}^S(\omega^S)$.

Next, note that, jointly, Conditions 1 and 2 in the lemma imply that

$$\mathcal{V}(\mathcal{S}_0) = \max\left\{V^S(\omega^S|\mathcal{S}_0), \max_{\omega^P \in \{\hat{\omega}^P \in \Omega^P : \mathcal{S}_0^P(\hat{\omega}^P) > 0\}} V^P(\omega^P|\mathcal{S}_0)\right\}.$$

Hence $\mathcal{V}$ solves the Bellman equation. That $\delta^T \mathbb{E}^\chi\left[\sum_{s=T}^\infty \delta^s U_s|\mathcal{S}\right] \to 0$ as $T \to \infty$ guarantees that $\mathcal{V}(\mathcal{S}_0) = \mathcal{V}^*(\mathcal{S}_0)$, and hence the optimality of the index policyχ^*. $\square$

This completes the proof. $\blacksquare$

S.2 Proof of Parts (2) and (3) of Proposition 2 in the main text

Part (2). For simplicity, suppose that the function f^ξ is constant, that there are two categories of candidates, α and β, so that $\Xi = \{\alpha, \beta\}$, and that each search brings to the CS a single candidate from category $\xi \in \Xi$ with probability ρ^ξ, where $\sum_{\xi \in \Xi} \rho^\xi = 1$. Further assume that the committee's initial CS contains only two candidates, one of each category and that the evaluation of each candidate generates a signal $\vartheta \in \{G, B\}$, with $\vartheta = G$ denoting a "good" evaluation outcome and $\vartheta = B$ a "bad" one. Each new evaluation of a ξ-candidate yields a good outcome $\vartheta = G$ with probability $q_1^\xi \in (0, 1]$ if the candidate is qualified ($\mu = 1$) and with probability $1 - q_0^\xi$ if the candidate is unqualified ($\mu = 0$), with $q_1^\xi > 1 - q_0^\xi$. Further assume that $q_1^\xi = q_0^\xi = q^\xi$, $\xi = \alpha, \beta$, meaning that the probability that qualified candidates deliver good outcomes is the same as the probability that unqualified candidates deliver bad outcomes, in which case the posterior belief $p^\xi(\theta)$ that a ξ-candidate is qualified depends on the history θ of past evaluations only through the difference between good and bad outcomes. The evaluation of a ξ-candidate with history θ costs $\lambda^\xi(\theta) \geq 0$ to the committee.[8]

Below we show that a reduction in the threshold for the α-candidates from Ψ^α to $\hat{\Psi}^\alpha = \Psi^\alpha - \varepsilon$, $\varepsilon > 0$, may strictly reduce the ex-ante probability that an α-candidate is hired. Because f^ξ is constant, for any θ, $\hat{\mathcal{I}}^P(\alpha, \emptyset) \geq \mathcal{I}^P(\alpha, \emptyset)$. The result below thus establishes that these policies can hurt α-candidates even when they increase the indexes of the α-candidates.

Denote by $D(\theta)$ the difference between good and bad outcomes recorded in θ. We first identify conditions under which the evaluation and CS expansion dynamics take a particularly simple form, both under the original threshold Ψ^α and the new one $\hat{\Psi}^\alpha$. We then compute the ex-ante probability that an α-candidate is hired under each of the two thresholds. Finally, we identify conditions under which the ex-ante probability the position is assigned to one of the α-candidates is strictly smaller under the new threshold. The conditions are sufficient but not necessary for the result, thus revealing that the result is not knife-edge.

Suppose that initially the thresholds $(\Psi^\alpha, \Psi^\beta)$ are such that $p^\alpha(\theta) \geq \Psi^\alpha$ if and only if $D(\theta) \geq 2$ and $p^\beta(\theta) \geq \Psi^\beta$ if and only if $D(\theta) \geq 1$. Further assume that, for any $\xi \in \Xi = \{\alpha, \beta\}$, $\lambda^\xi(\theta) > v^\xi$ for each θ containing more than four elements, meaning that it is not worth evaluating any candidate more than four times. Finally, suppose that $\lambda^\beta(\theta) = 0$ for all θ of dimensionality smaller than four. For $\xi = \alpha$, instead, $\lambda^\alpha(\theta) = 0$ for all θ of dimensionality smaller than four such that $p^\alpha(\theta) \geq p^\alpha(\emptyset)$, whereas, for any θ of dimensionality smaller than four such that $p^\alpha(\theta) < p^\alpha(\emptyset)$, $\lambda^\alpha(\theta) = \lambda^\alpha$, for some $\lambda^\alpha \in \mathbb{R}_{++}$. These assumptions are not crucial for the results, but simplify the calculations below. Given these assumptions, hereafter we will confine attention to histories θ that contain no more than 4 elements and no more than 1 bad outcome.

Denote by $\varsigma_B^\xi(p^\xi) \equiv p^\xi(1 - q^\xi) + (1 - p^\xi)q^\xi$, $\varsigma_G^\xi(p^\xi) \equiv p^\xi q^\xi + (1 - p^\xi)(1 - q^\xi)$, $\varsigma_{GG}^\xi(p^\xi) \equiv p^\xi(q^\xi)^2 + (1 - p^\xi)(1 - q^\xi)^2$, and $\varsigma_{BB}^\xi(p^\xi) \equiv p^\xi(1 - q^\xi)^2 + (1 - p^\xi)(q^\xi)^2$ the probabilities that (starting with belief p^ξ about the candidate's qualification) the first outcome is bad, the first outcome is good, the first two outcomes are good, and the first two outcomes are bad, respectively.

The index of an α-candidate for which $\theta \in \{\emptyset, (B, G), (G, B)\}$ is equal to

$$\mathcal{I}^P(\alpha, \emptyset) = \mathcal{I}^P(\alpha, (G, B)) = \mathcal{I}^P(\alpha, (B, G)) = \frac{(1 - \delta)\delta^2 v^\alpha \varsigma_{GG}^\alpha(p^\alpha(\emptyset))}{1 - [\delta \varsigma_B^\alpha(p^\alpha(\emptyset)) + \delta^2 q^\alpha(1 - q^\alpha)]}. \tag{S.14}$$

To see this, recall that, by Theorem 1 in the main text, the optimal stopping time in the definition of an index is the first time at which the value of the index drops weakly below its initial value. In particular, at each of the three histories in (S.14), the optimal stopping time $\tau^{\alpha *}$ for the α-candidates

[8] These assumptions introduced above simplify some of the derivations below but are not essential for the results. Also note that the assumption that $\sum_{\xi \in \Xi} \rho^\xi = 1$ is without loss of generality. The case where search does not bring any candidate with positive probability can always be captured by letting one of the categories replicate the arrival of no new candidate.

satisfies: $\tau^{\alpha*}=1$ if the next outcome is bad; $\tau^{\alpha*}=2$ if the next two outcomes are (G,B), that is, first a good outcome and then a bad one; and $\tau^{\alpha*}=\infty$ if the next two outcomes are both good. Similarly,

$$\mathcal{I}^P(\alpha,G)=\mathcal{I}^P(\alpha,(G,B,G))=\mathcal{I}^P(\alpha,(B,G,G))=\frac{(1-\delta)\delta v^\alpha\varsigma_G^\alpha\left(p^\alpha(G)\right)}{1-\delta\varsigma_B^\alpha\left(p^\alpha(G)\right)}, \tag{S.15}$$

where $p^\xi(G)=p^\alpha(\emptyset)q^\xi/\varsigma_G^\xi(p^\alpha(\emptyset))$ is the posterior probability that the candidate is qualified when θ contains a single good outcome. This is because, at each of the three histories in (S.15), the optimal stopping time $\tau^{\alpha*}$ for the α-candidates is $\tau^{\alpha*}=1$ if the next outcome is bad and $\tau^{\alpha*}=\infty$ otherwise. For an α-candidate whose first realization is bad, instead,

$$\mathcal{I}^P(\alpha,B)=\frac{-\lambda^\alpha(1-\delta)+(1-\delta)\delta^3v^\alpha\left[p^\alpha(B)(q^\alpha)^3+(1-p^\alpha(B))(1-q^\alpha)^3\right]}{1-\delta\left[\varsigma_B^\alpha\left(p^\alpha(B)\right)+\delta q^\alpha(1-q^\alpha)\left(1+\delta\varsigma_G^\alpha\left(p^\alpha(B)\right)\right)\right]},$$

where $p^\xi(B)=p^\xi(\emptyset)(1-q^\xi)/\varsigma_B^\xi(p^\xi(\emptyset))$ is the posterior probability that the candidate is qualified when θ contains a single bad outcome. This is because, when $\theta=B$, the optimal stopping time for the α-candidates is: $\tau^{\alpha*}=1$ if the next outcome is bad; $\tau^{\alpha*}=2$ if the next two outcomes are (G,B); $\tau^{\alpha*}=3$ if the next three outcomes are (G,G,B); and $\tau^{\alpha*}=\infty$ if the next three outcomes are (G,G,G).

Next consider the β-candidates. First, observe that

$$\mathcal{I}^P(\beta,\emptyset)=\mathcal{I}^P(\beta,(B,G))=\frac{(1-\delta)\delta v^\beta\varsigma_G^\beta(p^\beta(\emptyset))}{1-\varsigma_B^\beta(p^\beta(\emptyset))\delta}. \tag{S.16}$$

This is because, at the histories in the arguments of $\mathcal{I}^P(\beta,\theta)$ in (S.16), the optimal stopping time $\tau^{\beta*}$ for the β-candidates is $\tau^{\beta*}=1$ if the next outcome is bad and $\tau^{\beta*}=\infty$ otherwise. Similarly,

$$\mathcal{I}^P(\beta,B)=\frac{(1-\delta)\delta^2v^\beta\varsigma_{GG}^\beta(p^\beta(B))}{1-\delta\left[\varsigma_B^\beta\left(p^\beta(B)\right)+\delta(1-q^\beta)q^\beta\right]}.$$

This is because, when $\theta=B$, the optimal stopping time for the β-candidates satisfies: $\tau^{\beta*}=1$ if the next outcome is bad; $\tau^{\beta*}=2$ if the next two outcomes are (G,B); and $\tau^{\beta*}=\infty$ if the next two outcomes are (G,G).

Finally, note that the index of any ξ-candidate whose history θ includes two or more bad outcomes is nonpositive, as is that of any candidate for which θ contains more than four elements.

Now suppose that the ordering of the indexes prior to the relaxation of the threshold for the α-candidates satisfies

$$\begin{aligned}&\mathcal{I}^P(\alpha,\emptyset)=\mathcal{I}^P(\alpha,(G,B))=\mathcal{I}^P(\alpha,(B,G))>\mathcal{I}^P(\beta,\emptyset)=\mathcal{I}^P(\beta,(B,G))>\mathcal{I}^P(\alpha,B)\\&>\frac{(1-\delta)\left\{-c+\delta^2\left[\rho^\alpha\left(v^\alpha\varsigma_{GG}^\alpha(p^\alpha(\emptyset))\delta\left(1+2\delta^2(1-q^\alpha)q^\alpha\right)-\lambda^\alpha\varsigma_B^\alpha(p^\alpha(\emptyset))\right)+\rho^\beta v^\beta\varsigma_G^\beta(p^\beta(\emptyset))\right]\right\}}{1-\delta^2\left[\rho^\beta\varsigma_B^\beta(p^\beta(\emptyset))+\rho^\alpha\delta\left(\varsigma_{BB}^\alpha(p^\alpha(\emptyset))+2\delta q^\alpha(1-q^\alpha)\left(\varsigma_B^\alpha(p^\alpha(\emptyset))+\delta q^\alpha(1-q^\alpha)\right)\right)\right]}\\&>\mathcal{I}^P(\beta,B).\end{aligned} \tag{S.17}$$

Then, from Part (2) of Theorem 1 and Part (6) of Proposition 1 in the main text, we have that

$$\mathcal{I}^S = \frac{(1-\delta)\left\{-c+\delta^2\left[\rho^\alpha\left(v^\alpha\varsigma^\alpha_{GG}(p^\alpha(\emptyset))\delta\left(1+2\delta^2(1-q^\alpha)q^\alpha\right)-\lambda^\alpha\varsigma^\alpha_B(p^\alpha(\emptyset))\right)+\rho^\beta v^\beta\varsigma^\beta_G(p^\beta(\emptyset))\right]\right\}}{1-\delta^2\left[\rho^\beta\varsigma^\beta_B(p^\beta(\emptyset))+\rho^\alpha\delta\left(\varsigma^\alpha_{BB}(p^\alpha(\emptyset))+2\delta q^\alpha(1-q^\alpha)\left(\varsigma^\alpha_B(p^\alpha(\emptyset))+\delta q^\alpha(1-q^\alpha)\right)\right)\right]}.$$

This is because, if the new candidate brought to the CS by search is an α-candidate, then the optimal stopping time in the definition of the search index is given by: $\tau^{\alpha*}=\infty$ after the outcome histories (G,G), (B,G,G,G), and (G,B,G,G); $\tau^{\alpha*}=3$ after the history (B,B); $\tau^{\alpha*}=4$ after the histories (B,G,B) and (G,B,B); and $\tau^{\alpha*}=5$ after the histories (B,G,G,B) and (G,B,G,B). If, instead, the candidate identified by search is a β-candidate, then the optimal stopping time in the definition of the search index is $\tau^{\beta*}=\infty$ if the first outcome is good and $\tau^{\beta*}=2$ if the first outcome is bad.

Now suppose that the third party reduces the threshold for all the α-candidates from Ψ^α to $\hat{\Psi}^\alpha=\Psi^\alpha-\varepsilon$, with $\varepsilon>0$. Suppose that, as a result of the change, $p^\alpha(\theta)\geq\hat{\Psi}^\alpha$ if and only if $D(\theta)\geq 1$. Recall that $\hat{\mathcal{I}}^P(\alpha,\theta)$ are the indexes of the α-candidates, after the reduction in the acceptance threshold. Then observe that

$$\hat{\mathcal{I}}^P(\alpha,\emptyset)=\hat{\mathcal{I}}^P(\alpha,(B,G))=\frac{(1-\delta)\delta v^\alpha\varsigma^\alpha_G(p^\alpha(\emptyset))}{1-\varsigma^\alpha_B(p^\alpha(\emptyset))\delta}. \tag{S.18}$$

This is because the optimal stopping time for the α-candidates at the histories in the arguments of the indexes $\hat{\mathcal{I}}^P(\alpha,\theta)$ in (S.18) is now given by $\tau^{\alpha*}=1$ if the next outcome is bad and $\tau^{\alpha*}=\infty$ otherwise. Similarly,

$$\hat{\mathcal{I}}^P(\alpha,B)=\frac{(1-\delta)\left[-\lambda^\alpha+\delta^2 v^\alpha\varsigma^\alpha_{GG}(p^\alpha(B))\right]}{1-\delta\left[\varsigma^\alpha_B\left(p^\alpha(B)\right)+\delta(1-q^\alpha)q^\alpha\right]}.$$

This is because the optimal stopping time for the α-candidates after one bad outcome is now given by: $\tau^{\alpha*}=1$ if the next outcome is also bad; $\tau^{\alpha*}=2$ if the next two outcomes are (G,B); and $\tau^{\alpha*}=\infty$ if the next two outcomes are (G,G).

Now suppose that $\hat{\Psi}^\alpha$ is such that the following order of the indexes holds:

$$\begin{aligned}&\hat{\mathcal{I}}^P(\alpha,\emptyset)=\hat{\mathcal{I}}^P(\alpha,(B,G))>\mathcal{I}^P(\beta,\emptyset)=\mathcal{I}^P(\beta,(B,G))\\&>\frac{(1-\delta)\left[-c+\delta^2\sum_{\xi\in\{\alpha,\beta\}}\rho^\xi v^\xi\varsigma^\xi_G(p^\xi(\emptyset))\right]}{1-\delta^2\sum_{\xi\in\{\alpha,\beta\}}\rho^\xi\varsigma^\xi_B(p^\xi(\emptyset))}>\max\left\{\hat{\mathcal{I}}^P(\alpha,B),\mathcal{I}^P(\beta,B)\right\}.\end{aligned} \tag{S.19}$$

Given the ordering in (S.19), the search index after the reduction in the threshold for the α-candidates is then given by

$$\hat{\mathcal{I}}^S=\frac{(1-\delta)\left[-c+\delta^2\sum_{\xi\in\{\alpha,\beta\}}\rho^\xi v^\xi\varsigma^\xi_G(p^\xi(\emptyset))\right]}{1-\delta^2\sum_{\xi\in\{\alpha,\beta\}}\rho^\xi\varsigma^\xi_B(p^\xi(\emptyset))}.$$

This follows directly from Part (2) of Theorem 1 along with Part (6) of Proposition 1 in the main text (in particular, the optimal stopping in the definition of the search index is now equal to ∞ if the first outcome is good and is equal to 2 if it is bad, irrespective of the type of candidate brought to the CS).

Given the results above, we now compare the ex-ante probabilities of hiring one of the α-candidates under the old and new thresholds.

Selection of α-candidates under old threshold. The order of evaluations is given by (S.17). Because the search technology is stationary, by virtue of Part (5) of Proposition 1 in the main text, when search is launched, all candidates in the CS are discarded. Denote by A_S the probability with which one of the α-candidates is hired in the continuation immediately following search. We then have that

$$A_S = \rho^\alpha \left\{\varsigma^\alpha_{GG}(p^\alpha(\emptyset))\,(1+2(1-q^\alpha)q^\alpha) + [1-\varsigma^\alpha_{GG}(p^\alpha(\emptyset))\,(1+2(1-q^\alpha)q^\alpha)]\,A_S\right\} + \rho^\beta\varsigma^\beta_B(p^\beta(\emptyset))A_S.$$

This is because an α-candidate is hired after the histories (G,G), (G,B,G,G), (B,G,G,G), whereas a β-candidate is hired if and only if its first outcome is good. Rearranging, we have that

$$A_S = \frac{\rho^\alpha\varsigma^\alpha_{GG}(p^\alpha(\emptyset))\,[1+2(1-q^\alpha)q^\alpha]}{1-\left\{\rho^\alpha\,[1-\varsigma^\alpha_{GG}(p^\alpha(\emptyset))\,(1+2(1-q^\alpha)q^\alpha)] + \rho^\beta\varsigma^\beta_B(p^\beta(\emptyset))\right\}}.$$

Therefore, the ex-ante probability one of the α-candidates is hired is equal to

$$\begin{aligned} A = {} & \varsigma^\alpha_{GG}(p^\alpha(\emptyset))\left[1+(1-q^\alpha)q^\alpha\left(1+\varsigma^\beta_B(p^\beta(\emptyset))\right)\right] \\ & + \left[\varsigma^\alpha_{BB}(p^\alpha(\emptyset)) + 2(1-q^\alpha)^2(q^\alpha)^2 + 2q^\alpha(1-q^\alpha)\varsigma^\alpha_B(p^\alpha(\emptyset))\right]\varsigma^\beta_B(p^\beta(\emptyset))A_S. \end{aligned}$$

This is because the β-candidate in the initial CS is explored if either the first evaluation of the α-candidate in the initial CS yields a bad outcome, or the history of outcomes of the α-candidate in the initial CS contains two bad outcomes. Furthermore, the α-candidate in the initial CS is hired in the following cases: (1) the first two evaluations yield two good outcomes; (2) the first four evaluations yield the outcomes (G,B,G,G);[9] (3) the first evaluation yields a bad outcome, at which point the β-candidate is evaluated, yields a bad outcome and is dismissed, and the subsequent three evaluations of the α-candidate yield three good outcomes.

Selection of α-candidates under new threshold. Denote by $\hat{A}_S$ and $\hat{A}$ the analogs of A_S and A, respectively, after the reduction in the threshold. Given the ordering in (S.19), we have that

$$\hat{A}_S = \rho^\alpha\left[\varsigma^\alpha_G(p^\alpha(\emptyset)) + \varsigma^\alpha_B(p^\alpha(\emptyset))\hat{A}_S\right] + \rho^\beta\varsigma^\beta_B(p^\beta(\emptyset))\hat{A}_S,$$

from which we obtain that

$$\hat{A}_S = \frac{\rho^\alpha\varsigma^\alpha_G(p^\alpha(\emptyset))}{1-\left[\rho^\alpha\varsigma^\alpha_B(p^\alpha(\emptyset)) + \rho^\beta\varsigma^\beta_B(p^\beta(\emptyset))\right]}.$$

[9] Note that, in cases (1) and (2), the β-candidate in the initial CS is never explored.

The ex-ante probability with which one of the α-candidates is hired is thus equal to:

$$\hat{A} = \varsigma_G^{\alpha}(p^{\alpha}(\emptyset)) + \varsigma_B^{\alpha}(p^{\alpha}(\emptyset))\varsigma_B^{\beta}(p^{\beta}(\emptyset))\hat{A}_S.$$

Note that the result follows from the fact that the α-candidate in the initial CS is evaluated first. If the evaluation yields a good outcome, the α-candidate is hired. If, instead, it yields a bad outcome, the β-candidate in the initial CS is evaluated next. If the evaluation of the β-candidate yields a positive outcome, the β-candidate is hired, bringing an end to the recruiting process. If, instead, it yields a negative outcome, search is launched, at which point the probability one of the α-candidates is hired is $\hat{A}_S$.

Comparison. The result in Part (2) in the proposition follows from observing that Conditions (S.17) and (S.19) are consistent with $A > \hat{A}$ over an open set of parameter values.

Part (3). The result follows directly from the fact that, when the CS is exogenous, the increase in the indexes $\mathcal{I}^P(\alpha, \theta)$ of the α-candidates, together with the optimal policy being the index rule, implies that the DM evaluates α-candidates at histories at which she would have evaluated one of the β-candidates. Furthermore, at any history at which the DM evaluates, or assigns the position to, one of the β-candidates, she would have done the same under the original thresholds. This implies that the probability that the position is given to one of the α-candidates is higher under the new thresholds. ■

S.3 Pandora's boxes with an endogenous CS

In this section, we show how the results in Theorem 1 and Proposition 1 in the main body also provide a solution to an extension of Weitzman's (1979) "Pandora's boxes problem" in which the set of boxes is endogenously expanded over time based on the results of past explorations. In this problem, each alternative is a "box" and belongs to a category $\xi \in \Xi$. To each category corresponds a pair (F^{ξ}, λ^{ξ}), where F^{ξ} is the distribution from which the box's prize v is drawn and λ^{ξ} is the cost of inspecting (i.e., of opening) the box. As in Weitzman's (1979) original setting, each box's prize v is drawn independently (conditional on the categories) and revealed upon the first inspection.

At each period, the DM can either (a) search for additional boxes to add to the CS, (b) open one of the boxes in the CS to learn its prize, or (c) stop and either recall the prize of one of the previously opened boxes, or take the outside option (with a value normalized to 0), with either one of the last two choices ending the decision problem. For simplicity (but also motivated by the application to online consumer search), assume that each search $m \in \mathbb{N}$ brings exactly one box, whose category ξ is drawn from Ξ according to a distribution $\rho(m) \in \Delta(\Xi)$, which may depend on the number of past searches $m-1$ but is invariant to the realizations of such past searches. The draw from each $\rho(m)$ is independent of the draw from each $\rho(l)$, $l \neq m$.

We assume that $\Xi \subset \mathbb{N}$, with higher ξ denoting superior boxes, in the sense that, for any $\xi', \xi'' \in \Xi$ with $\xi'' > \xi'$, $F^{\xi''} \succeq_{FOSD} F^{\xi'}$ and $\lambda^{\xi''} \leq \lambda^{\xi'}$ (with one of the two relationships strict). Let $\underline{\xi} \equiv \inf \Xi$ and $\overline{\xi} \equiv \sup \Xi$. The cost of the m-th expansion of the CS is $c(m)$, where $c(\cdot)$ is a positive and increasing function. In addition, we assume that, for all m, $\rho(m) \succeq_{FOSD} \rho(m+1)$; that is, the distribution $\rho(m) \in \Delta(\Xi)$ from which the category of the m-th box is drawn first-order-stochastically dominates, weakly, the distribution $\rho(m+1) \in \Delta(\Xi)$ from which the category of the $(m+1)$-th box is drawn. The combination of the assumption that $c(m)$ is weakly increasing in m and the distribution $\rho(m) \in \Delta(\Xi)$ from which the boxes are drawn "decreases" with m in a FOSD sense implies that the index of search $\mathcal{I}^S(m)$ defined below is decreasing in m and can be characterized using the same properties as when the search technology deteriorates in the sense of Definition 2 in the main body (as per Part (6) of Proposition 1 in the main body).

We denote by $\rho^{\xi}(m)$ the probability that the m-th search brings a ξ-box, with $\sum_{\xi\in\Xi}\rho^{\xi}(m) = 1$ for all m.[10] As in the baseline model, the DM discounts the future according to δ.

The setting described above is one in which the decision to walk away with the prize of an opened box, or the outside option, brings an end to the DM's problem. The framework described in Section 2 in the main body, instead, has an infinite horizon, and the DM chooses indefinitely among the alternatives. Despite this difference, the solution to this problem takes the form of an index policy akin to the one in Definition 1 in the main body.

Proposition S.1 below characterizes the optimal policy, prescribing when to search for an additional box, the order in which existing boxes should be opened, and when to stop and either recall an opened box or the outside option. The proof maps the Pandora's boxes problem with an endogenous set of boxes into an auxiliary problem that fits into the setting of Section 2 in the main body, and then uses Theorem 1 and Proposition 1 in the main body to identify the properties of the optimal policy in Proposition S.1.

Proposition S.1 (Pandora's-boxes with an endogenous set of boxes). *For any ξ-box that has not been opened yet (i.e., for which $\omega^P = (\xi, \emptyset)$ for some $\xi \in \Xi$) the reservation prize $\mathcal{I}^P(\xi,\emptyset)$ is given by the solution to:*

$$\mathcal{I}^P(\xi,\emptyset) = \frac{-\lambda^{\xi} + \delta\int_{\frac{\mathcal{I}^P(\xi,\emptyset)}{1-\delta}}^{\infty} v\,dF^{\xi}(v)}{1 + \frac{\delta}{1-\delta}\left(1 - F^{\xi}\left(\frac{\mathcal{I}^P(\xi,\emptyset)}{1-\delta}\right)\right)}. \tag{S.20}$$

For any $l \in \mathbb{R}$, let $\Xi(l) \equiv \left\{\xi \in \Xi :\ \mathcal{I}^P(\xi,\emptyset) > l\right\}$ denote the set of boxes whose reservation prize exceeds l. For any m, the reservation prize of search $\mathcal{I}^S(m)$ is given by the solution to:[11]

$$\mathcal{I}^S(m) = \frac{-c(m) + \delta\sum_{\xi\in\Xi(\mathcal{I}^S(m))}\rho^{\xi}(m)\left(-\lambda^{\xi} + \delta\int_{\frac{\mathcal{I}^S(m)}{1-\delta}}^{\infty} v\,dF^{\xi}(v)\right)}{1 + \sum_{\xi\in\Xi(\mathcal{I}^S(m))}\rho^{\xi}(m)\left[\delta + \frac{\delta^2}{1-\delta}\left(1 - F^{\xi}\left(\frac{\mathcal{I}^S(m)}{1-\delta}\right)\right)\right]}. \tag{S.21}$$

The solution to Pandora's-boxes problem with an endogenous CS takes the following form:

1. *If the highest reservation prize among all unopened boxes in the CS is greater than the reservation prize $\mathcal{I}^S(m)$ of search, and is greater than the flow value $(1-\delta)v$ of each opened box*

[10] All the results extend to the case where Ξ is infinite.

[11] Because all the relevant information about the state of the search technology is summarized in the number of past searches, we abuse notation and let $\mathcal{I}^S(m)$ denote the index for the m-th search.

and the outside-option, the DM opens one of the boxes with the highest reservation prize.

2. *If the reservation prize of search $\mathcal{I}^S(m)$ is higher than the reservation prize $\mathcal{I}^P(\xi, \emptyset)$ of any unopened box and of the flow value $(1-\delta)v$ of each opened box and the outside-option, the DM searches.*

3. *If neither of the above two situations applies, the DM stops. He then takes the prize of one of the opened boxes whose flow value $(1-\delta)v$ is the highest among the opened boxes if the latter value exceeds the outside-option, and takes the outside-option otherwise.*

As in Weitzman's problem, the reservation prizes $\mathcal{I}^P(\xi, \emptyset)$ of the boxes that have not been opened yet have the following interpretation.[12] Suppose there are only two alternatives. One is an unopened ξ-box and the other is a hypothetical box, whose prize is an annuity yielding K in each period, where K is known. The reservation prize is the value of K for which the DM is indifferent between taking the hypothetical box (yielding a continuation payoff of $K/(1-\delta)$) and inspecting the ξ-box while maintaining the option to recall the hypothetical box once the prize v of the ξ-box is discovered.

The reservation prize $\mathcal{I}^S(m)$ of search extends this interpretation as follows. Suppose there are two options: the hypothetical box with known value K described above, and the option of expanding the CS. The reservation prize of search is the value K for which the DM is indifferent between taking the hypothetical box right away, and expanding the CS, maintaining the option to take the hypothetical box either (a) once the category ξ of the newly discovered box is discovered and $\mathcal{I}^P(\xi, \emptyset) \leq K$, or (b), in case $\mathcal{I}^P(\xi, \emptyset) > K$, after the prize v of the newly discovered ξ-box is learned and $v \leq K/(1-\delta)$.

Proof of Proposition S.1. Consider a relaxed problem in which the DM gets a flow payoff equal to $(1-\delta)v$ each time she selects an opened box with value v, and can revert her decision at any period. The solution to such a problem is the index policy of Theorem 1 in the main text and has the property that, once an opened box is selected, it continues to be selected in all subsequent periods. The index policy for such a problem is thus feasible (and hence optimal) also in the primitive problem.

To see that the index of a ξ-box that has not been opened yet is given by (S.20), note that the index of an opened box is equal to $(1-\delta)v$. Because the optimal stopping time τ^* in the definition

$$\mathcal{I}^P(\xi, \emptyset) \equiv \sup_{\tau>0} \frac{\mathbb{E}\left[\sum_{s=0}^{\tau-1} \delta^s u_s | \xi, \emptyset\right]}{\mathbb{E}\left[\sum_{s=0}^{\tau-1} \delta^s | \xi, \emptyset\right]}, \tag{S.22}$$

of the index $\mathcal{I}^P(\xi, \emptyset)$ is the first time at which the value of the index drops below its value $\mathcal{I}^P(\xi, \emptyset)$

[12] Weitzman defines the reservation prize $\hat{\mathcal{I}}^P(\omega^P)$ for $\omega^P = (\xi, \emptyset)$ as the solution to $\lambda^\xi = \delta \int_{\hat{\mathcal{I}}^P(\omega^P)}^{\infty} (v - \hat{\mathcal{I}}^P(\omega^P)) \mathrm{d}F^\xi(v) - (1-\delta)\hat{\mathcal{I}}^P(\omega^P)$, which yields
$\hat{\mathcal{I}}^P(\omega^P) = [-\lambda^\xi + \delta \int_{\hat{\mathcal{I}}^P(\omega^P)}^{\infty} v \mathrm{d}F^\xi(v)]/[1 - F^\xi(\hat{\mathcal{I}}^P(\omega^P))]$. The reservation prizes in (S.20) are thus equal to those in Weitzman (1979) multiplied by $(1-\delta)$, that is, $\mathcal{I}^P(\omega^P) = (1-\delta)\hat{\mathcal{I}}^P(\omega^P)$.

at the time the index is computed, we then have that $\tau^* = 1$ if $(1-\delta)v \leq \mathcal{I}^P(\xi, \emptyset)$ and $\tau^* = \infty$ otherwise.

Turning to the index for search, the combination of the assumption that $c(m)$ is weakly increasing in m with the assumption that the distribution $\rho(m) \in \Delta(\Xi)$ from which the boxes are drawn "decreases" with m in a FOSD sense implies that the optimal stopping-time τ^* in

$$\mathcal{I}^S(m) \equiv \sup_{\pi,\tau} \frac{\mathbb{E}^\pi\left[\sum_{s=0}^{\tau-1} \delta^s U_s | m\right]}{\mathbb{E}^\pi\left[\sum_{s=0}^{\tau-1} \delta^s | m\right]}, \tag{S.23}$$

is equal to (a) $\tau^* = \infty$ if the box identified at the m-th search has a reservation prize $\mathcal{I}^P(\xi, \emptyset) > \mathcal{I}^S(m)$ and its realized flow payoff satisfies $v(1-\delta) > \mathcal{I}^S(m)$, (b) $\tau^* = 1$ if $\mathcal{I}^P(\xi, \emptyset) \leq \mathcal{I}^S(m)$, and (c) $\tau^* = 2$ if $\mathcal{I}^P(\xi, \emptyset) > \mathcal{I}^S(m)$ and $v(1-\delta) \leq \mathcal{I}^S(m)$. ■

S.4 Online consumer search

As indicated in the main body, the results in Theorem 1 and Proposition 1 in the main text (along with their adaptation to the Pandora's boxes problem introduced above) can be used in the application to online consumer search to endogenize the probability with which the consumer reads the ads, clicks on them, and finalizes her purchases. Furthermore, one can derive a structural relationship between the various positions and their click-through-rates (CTRs), accounting for the uncertainty that the consumer faces about the ads displayed at the various positions – a feature that the model with exogenous CSs does not capture.

S.4.1 Eventual purchase theorem with an endogenous CS

We start by showing that the properties of Proposition 1 in the main text permit one to express the consumer's eventual purchase decisions in terms of comparisons of the products' discovery values. In a similar setting, but with an exogenous CS, Choi, Dai and Kim (2018) – and, independently, Armstrong (2017) – derive a static condition characterizing eventual purchasing decisions based on a comparison of "effective values." Such a characterization extends to search problems with an endogenous CS. Let v_m denote the value to the consumer for the product sold by the firm advertising at the m-th position. For all $m \geq 1$, let $w_m \equiv \min\{\mathcal{I}_m, v_m(1-\delta)\}$ be the "effective value" of the product advertised at the m-th position (for brevity, product m) when the product is already in the consumer's CS, and $d_m \equiv \min\{w_m, \mathcal{I}^S(m)\}$ the product's "discovery value," when the product must be brought to the CS before it can be explored (that is, before the consumer learns the product's ad type $\xi(m)$). Let product $m = 0$ correspond to the consumer's outside option, with $w_0 = d_0 = 0$. Note that w_m and d_m are learned by the consumer only after reading the m-th ad (which reveals its type $\xi(m)$) and clicking on it which reveal its value v_m.

Proposition S.2 (**Eventual purchases**). *The consumer purchases product m if, for all $l \in \mathbb{N} \cup \{0\}$, $l \neq m$, $d_l < d_m$ (and only if $d_l \leq d_m$, for all $l \neq m$).*

As in Choi, Dai and Kim (2018), purchasing decisions are thus determined by a static comparison of the products' values, as in canonical discrete-choice models. Contrary to Choi, Dai and Kim (2018), however, such values account for the uncertainty the consumer faces over the products occupying the various positions (equivalently, over each product's ad type $\xi(m)$ prior to reading the ad displayed in the m-th position). Allowing for such an uncertainty is important. When all products are already in the consumer's CS, positions do not play any specific role and there is no reason why downstream positions should be expected to receive fewer clicks than upstream ones.

In contrast, when the consumer faces uncertainty about the ads occupying the various positions and chooses how to alternate between reading new ads and clicking on the links of those ads she read already, the model delivers useful structural relationships linking the positions' CTRs to the primitives of the problem. In particular, the characterization of eventual purchase with endogenous CS implies that, when the reading cost $c(\cdot)$ is non-decreasing and the probability of finding "attractive" ads declines (weakly) with the positions, all other things equal, the further down a product is on the list, the lower the ex-ante probability the product is purchased (and hence its ex-ante demand), a property often assumed, but not micro-founded, in the models considered in the pertinent literature. The characterization follows from the fact that the optimal policy is an index rule, along with the fact that the search index $\mathcal{I}^S(m)$ declines with m. Heuristically, if a consumer reads the m-th ad, it must be that the reservation prizes $\mathcal{I}_l$ of all products $l < m$ already in her CS, as well as the discovered values $v_l(1-\delta)$ of those products $l < m$ that have been inspected already, are no greater than $\mathcal{I}^S(m)$. When the search index $\mathcal{I}^S(\cdot)$ is non-increasing, $\mathcal{I}^S(m+1) \leq \mathcal{I}^S(m)$. Hence, if after reading the m-th ad, $\mathcal{I}_m \geq \mathcal{I}^S(m)$, the consumer necessarily clicks on the m-th ad, thus learning product m's value v_m. Once v_m is learned, if $(1-\delta)v_m \geq \mathcal{I}_m$, the consumer stops the search and purchases product m. The above properties imply that the purchasing decisions are indeed determined by the comparisons of the discovery values defined above. The formal proof follows.

Proof of Proposition S.2. Since product 0 corresponds to the outside option, a product is always purchased. Let $l \neq m$ be such that $d_l < d_m$. We show that product l will not be purchased.

Case 1: $l > m$ (i.e., l is read after m is read). First, suppose that $d_l = \mathcal{I}^S(l)$. Because $\mathcal{I}^S(l) \leq \mathcal{I}^S(m)$ and because $\min\{\mathcal{I}_m, (1-\delta)v_m\} \geq d_m > \mathcal{I}^S(l)$, under the index policy of Theorem 1 in the main text, product l is read only after product m is clicked upon. Once m is clicked, however, because $(1-\delta)v_m > \mathcal{I}^S(l)$, l is never read. Hence, l will not be purchased. Next suppose that $d_l = \mathcal{I}_l$. Then, $\min\{\mathcal{I}_m, (1-\delta)v_m\} \geq d_m > \mathcal{I}_l$. Thus, product l is clicked only after m is clicked. But again, once m is clicked, because $(1-\delta)v_m > \mathcal{I}_l$, l is never clicked, implying that l is not purchased. Finally, suppose $d_l = (1-\delta)v_l$. Then, because $\min\{\mathcal{I}_m, (1-\delta)v_m\} \geq d_m > (1-\delta)v_l$, m must be clicked before l is purchased. Because $v_m > v_l$, l is not purchased after m's value is learned.

Case 2: $l < m$ (i.e., l is read before m is read). Because

$$\mathcal{I}^S(m) \geq d_m > d_l \equiv \min\{\mathcal{I}_l, (1-\delta)v_l, \mathcal{I}^S(l)\},$$

and because $\mathcal{I}^S(m) \leq \mathcal{I}^S(l)$, it must be that $d_l = \min\{\mathcal{I}_l, (1-\delta)v_l\}$ and hence

$$\min\{\mathcal{I}_l, (1-\delta)v_l\} < d_m \leq \min\{\mathcal{I}_m, (1-\delta)v_m\}. \tag{S.24}$$

Furthermore, because the search technology is non-improving, $\mathcal{I}^S(l+1) \geq ... \geq \mathcal{I}^S(m-1) \geq \mathcal{I}^S(m)$. Along with the fact that $d_l = \min\{\mathcal{I}_l, (1-\delta)v_l\} < d_m \leq \mathcal{I}^S(m)$, this implies that $\min\{\mathcal{I}_l, (1-\delta)v_l\} < \mathcal{I}^S(k)$ for all $(l+1) \leq k \leq m$. This last property in turn implies that either clicking on l, or purchasing l, is dominated by reading any product k, with $(l+1) \leq k \leq m$. If m is read, then (S.24) implies that l will not be purchased (the arguments are similar to those for case 1). If, instead, m is not read, it must be that another product $k \neq l, m$ is purchased. In either case, product l is not purchased. ■

S.4.2 Click-Through-Rates (CTRs)

The results in Theorem 1 and Proposition 1 in the main body can also be used to characterize the positions' click-through rates (hereafter, CTRs), i.e., the fraction of ads at each position that, once read, are clicked upon. Formally, for each position m, the corresponding CTR is equal to

$$CTR(m) \equiv \Pr\left(m\text{'s ad is clicked}\right).$$

Depending on the problem of interest, the information used to compute the above probability may contain the type of the firms advertising at the different positions (as when firms know the attractiveness of each others' ads at the bidding stage and the above probability is computed by the firms given the induced allocation) or only the knowledge of the rules used by the search engine to assign the ads to the various positions (as when the probability is computed by a platform that does not know the attractiveness of the firms' ads, or by a firm that also lacks such information).

The next proposition relates the CTRs to the effective and discovery values introduced above.

Proposition S.3 (**Click-through rates**). *For each position $m \geq 1$, the click-through-rate is given by*[13]

$$CTR(m) = Pr\left(\mathcal{I}^S(m) \geq max_{l<m}\{w_l\} \cap \mathcal{I}_m \geq \{max_{l<m}\{w_l\}, max_{l>m}\{d_l\}\}\right).$$

To understand the formula, note that, in order for the ad in position m to be read, it must be that $\mathcal{I}^S(m) \geq \max_{l<m}\{w_l\}$, for otherwise the consumer selects one of the products advertised in one of the preceding positions before reading the ad displayed in the m-th position. Once product m is read, in order for it to be clicked upon, it must be that its index $\mathcal{I}_m$ exceeds the effective value of each product brought to the consumer's CS prior to m, but also the discovery value of all products

[13]For simplicity, the formula in the proposition assumes that, in case of indifference, the consumer favors position m (both when it comes to reading and clicking it). This is what justifies the weak inequalities in the formula. The proof discusses how alternative ways of breaking the indifferences must be accounted for if one were to compute bounds for such probabilities.

advertised further down the list, for otherwise the consumer selects one of the other products before clicking on m. Note that the property that $\mathcal{I}^S(l)$ is weakly decreasing in l is important here. It implies that, if for some position $l > m$, $d_l > \mathcal{I}_m$, then for all $j = m+1, ..., l$, $\mathcal{I}^S(j) > \mathcal{I}_m$, meaning that the consumer will necessarily read the ad of any product displayed between position m and position l before clicking on m. If for any of such product the discovery value exceeds $\mathcal{I}_m$, the consumer purchases one of these products before clicking on m, and never clicks on the m-th product.

Take the perspective of an observer (e.g., a platform, a firm, or a savvy consumer) knowing the rules used to assign the ads to the positions but not the attractiveness of the ads. While the probability each ad is read is decreasing in m, $\Pr\left(\mathcal{I}_m \geq \max_{l>m}\{d_l\}\right)$ need not be decreasing in m. Hence, from the observer's perspective, CTRs need not be monotone in positions, consistently with what has been noticed in the empirical literature.

Proof of Proposition S.3. The proof is in two steps. Step 1 shows that $\mathcal{I}^S(m) \geq \max_{l<m}\{w_l\}$ is necessary for product m to be read and that $\mathcal{I}^S(m) > \max_{l<m}\{w_l\}$ implies that product m is necessarily read. Step 2 shows that product m is clicked only if

$$\mathcal{I}^S(m) \geq \max_{l<m}\{w_l\} \quad \text{and} \quad \mathcal{I}_m \geq \max\left\{\max_{l>m}\{d_l\}, \max_{l<m}\{w_l\}\right\} \tag{S.25}$$

and that, when both of the above inequalities are strict, product m is necessarily clicked. The result in the proposition then follows directly from the above properties.

Step 1. To see that $\mathcal{I}^S(m) \geq \max_{l<m}\{w_l\}$ is necessary for product m to be read, suppose that, for some $l < m$, $w_l > \mathcal{I}^S(m)$. That is, both the index corresponding to clicking on product l, $\mathcal{I}_l$, and the one corresponding to purchasing product l, $(1-\delta)v_l$, are strictly greater than $\mathcal{I}^S(m)$. Because product l is read before product m is read, by Theorem 1 in the main text, m is never read.

Next, we show that, when $\mathcal{I}^S(m) > \max_{l<m}\{w_l\}$, product m is always read. To see this, note that since the search cost $c(\cdot)$ is increasing, $\mathcal{I}^S(1) \geq ... \geq \mathcal{I}^S(m-1) \geq \mathcal{I}^S(m)$. Therefore, $\mathcal{I}^S(m) > \max_{l<m}\{w_l\}$ implies that, for any $1 \leq l \leq m$, $\mathcal{I}^S(l) > w_{l-1} = \min\{\mathcal{I}_{l-1}, (1-\delta)v_{l-1}\}$. Hence, by Theorem 1 in the main text, for any $1 \leq l \leq m$, it cannot be that product $l-1$ is purchased before product l is read. Repeatedly applying this argument for all $1 \leq l \leq m$ implies product m must be read before any product $l < m$ is purchased.

Step 2. To see that both inequalities in (S.25) must hold for product m to be clicked, first observe that we already established in Step 1 that the first inequality in (S.25) is necessary for product m to be read. Thus assume that such inequality holds. To see that the second inequality in (S.25) must also hold, suppose that $\mathcal{I}_m < \max\left\{\max_{l>m}\{d_l\}, \max_{l<m}\{w_l\}\right\}$. Then either there exists a product $l < m$ such that $w_l > \mathcal{I}_m$, or a product $l > m$ such that $d_l > \mathcal{I}_m$, or both. Suppose there is a product $l < m$ such that $w_l > \mathcal{I}_m$. Then product m cannot be clicked, because product l is necessarily read before m and, because both $\mathcal{I}_l$ and $(1-\delta)v_l$ are strictly greater than $\mathcal{I}_m$, product l is purchased before m is clicked. Next, suppose that there exists a product $l > m$ such that $d_l = \min\{\mathcal{I}^S(l), \mathcal{I}_l, (1-\delta)v_l\} > \mathcal{I}_m$. By the monotonicity of the search indexes, $\mathcal{I}^S(m) \geq \mathcal{I}^S(m+1) \geq ... \geq \mathcal{I}^S(l)$. That $\mathcal{I}^S(l) > \mathcal{I}_m$, then implies that $\mathcal{I}^S(k) > \mathcal{I}_m$ for any

$k = m, m+1, ..., l$. In turn, this last property implies that clicking on m is dominated by reading product k, for any $k = m+1, ..., l$. If product l is read, because both $\mathcal{I}_l$ and $(1-\delta)v_l$ are strictly greater than $\mathcal{I}_m$, product m is not clicked. If, instead, product l is not read, it must be that another product $k \neq l, m$, with $k \in \{m+1, ..., l-1\}$, is purchased. In either case, product m is not clicked. Hence, the two inequalities in (S.25) are necessary for product m to be clicked.

Next, we show that when the two inequalities in (S.25) are strict, product m is necessarily clicked. We already established in Step 1 that, when the first inequality in (S.25) is strict, product m is read. Now suppose that the second inequality is also strict. That $\mathcal{I}_m > \max_{l<m}\{w_l\}$ implies that, for each product $l < m$, either $\mathcal{I}_l$ or $(1-\delta)v_l$ are strictly smaller than $\mathcal{I}_m$. Because product m is read, by Theorem 1 in the main text, it cannot be that any product $l < m$ is purchased before product m is clicked. Similarly, that $\mathcal{I}_m > \max_{l>m}\{d_l\}$ implies that, for each $l > m$, either $\mathcal{I}^S(l)$, or $\mathcal{I}_l$, or $(1-\delta)v_l$ are strictly smaller than $\mathcal{I}_m$, which again guarantees that no product $l > m$ can be purchased before product m is clicked. Since one of the products is necessarily purchased (product 0 representing the outside option), it must be that product m is clicked. Hence, we conclude that when the two inequalities in (S.25) are both strict, product m is necessarily clicked. ■

S.4.3 Parametric example

Consider a market with two firms, each of which advertises a single product. Let z denote the profit a firm derives from selling its product and assume that the two firms' profits are drawn independently from a distribution Z. Suppose that the product of each firm can either be highly attractive to the consumer ($\xi = H$) or less attractive ($\xi = L$), with the types drawn independently from $\Xi = \{H, L\}$, with $Pr(\xi = H) = q^H$. A highly-attractive product yields an utility to the consumer, net of the purchasing price, drawn uniformly from $[0, 1+\alpha]$, where $\alpha > 0$. A less-attractive product, instead, yields a net utility drawn uniformly from $[0, 1]$. The consumer learns the attractiveness of a firm's product by reading the firm's ad but discovers her net value for the product only by clicking on the ad and being directed to the firm's webpage. For simplicity, assume that $\lambda \equiv c \equiv 0$, so that the only cost is discounting.

The two firms advertise their products on a platform using the ascending-clock version of the generalized second-price auction to allocate the two ad positions. The firm dropping out first is allocated the second position and pays nothing, whereas the other firm is allocated the first position and pays to the platform the price at which the other firm dropped out per click.

Using the formula for the reservation prize in Proposition S.1 characterizing the optimal policy in Pandora's-boxes problem with an endogenous CS, the "clicking index" of an L-type firm is given by

$$\mathcal{I}_L^P \equiv \mathcal{I}^P(L, \emptyset) = \frac{\frac{\delta}{2(1-\delta)^2}\left((1-\delta)^2 - (\mathcal{I}^P(L,\emptyset))^2\right)}{1 + \frac{\delta}{(1-\delta)^2}\left((1-\delta) - \mathcal{I}^P(L,\emptyset)\right)},$$

which, solving for $\mathcal{I}^P(L,\emptyset)$, yields

$$\mathcal{I}_L^P = \frac{1-\delta}{\delta}\left(1-\sqrt{1-\delta^2}\right).$$

Similarly, the clicking index of an H-type firm is

$$\mathcal{I}_H^P \equiv \mathcal{I}^P(H,\emptyset) = (1+\alpha)\left(\frac{1-\delta}{\delta}\right)\left(1-\sqrt{1-\delta^2}\right).$$

Because the consumer always reads the first ad (as there are no direct cost of reading and the outside option is equal to 0), the initial search index (corresponding to the decision to read the first ad) plays no role in the analysis and hence we do not provide its characterization. The following lemma characterizes in closed form the index for reading the second ad which, which is a function of ρ^L and ρ^H. The probabilities ρ^L and ρ^H, may of course depend on the type of ad encountered in the first position.

Lemma S.5. *Let ρ^L and ρ^H represent the probabilities that the consumer assigns to finding an L or H firm in the second position. The index $\mathcal{I}^S$ for the decision to read the second ad is equal to*

$$\mathcal{I}^S(\rho^L,\rho^H) = \frac{(1-\delta)(1+\alpha)\left[1-\sqrt{1-\frac{\delta^4}{1+\alpha}\left(1+\alpha\rho^L\right)\left(1+\alpha\rho^H\right)}\right]}{\delta^2(1+\alpha\rho^L)}$$

if

$$\frac{\delta\left(1-\sqrt{1-\delta^2}\right)\left(1+\alpha\rho^L\right)}{1+\alpha} > 1-\sqrt{1-\frac{\delta^4}{1+\alpha}\left(1+\alpha\rho^L\right)\left(1+\alpha\rho^H\right)}$$

and otherwise,

$$\mathcal{I}^S(\rho^L,\rho^H) = \frac{(1-\delta)(1+\alpha)\left[1-\delta+\delta\rho^H-\sqrt{\left(1-\delta+\delta\rho^H\right)^2-\delta^4\left(\rho^H\right)^2}\right]}{\delta^2\rho^H}.$$

Proof of Lemma S.5. Recall that the index of search in the Pandora's boxes problem with endogenous CS is given by

$$\mathcal{I}^S(\rho^L,\rho^H) = \frac{\delta^2\sum_{\xi\in\Xi(\mathcal{I}^S(m))}\rho^\xi(m)\left(\int_{\frac{\mathcal{I}^S(m)}{1-\delta}}^{\infty} v\mathrm{d}F^\xi(v)\right)}{1+\sum_{\xi\in\Xi(\mathcal{I}^S(m))}\rho^\xi(m)\left[\delta+\frac{\delta^2}{1-\delta}\left(1-F^\xi\left(\frac{\mathcal{I}^S(m)}{1-\delta}\right)\right)\right]},$$

where $\Xi(l) \equiv \left\{\xi \in \Xi : \ \mathcal{I}^P(\xi,\emptyset) > l\right\}$. Since there are no direct costs,

$$\mathcal{I}^S(\rho^L,\rho^H) = \frac{-\delta^2 \sum_{\xi\in\Xi(\mathcal{I}^S(m))} \rho^\xi(m) \left(\int_{\frac{\mathcal{I}^S(m)}{1-\delta}}^{\infty} v \mathrm{d}F^\xi(v)\right)}{1+\sum_{\xi\in\Xi(\mathcal{I}^S(m))} \rho^\xi(m)\left[\delta + \frac{\delta^2}{1-\delta}\left(1-F^\xi\left(\frac{\mathcal{I}^S(m)}{1-\delta}\right)\right)\right]}.$$

Note that since $\mathcal{I}^P(\emptyset,H) > \mathcal{I}^P(\emptyset,L)$, it cannot be that $\Xi(\mathcal{I}^S(\rho^L,\rho^H)) = \{L\}$. Furthermore, it cannot be that $\Xi(\mathcal{I}^S(\rho^L,\rho^H)) = \emptyset$, as in this case no category has an index greater than search, which means stopping (in the definition of search index) occurs immediately after search is carried out, yielding $\mathcal{I}^S(2) = 0$, a contradicting. Hence, there are two feasible cases to consider: (i) $\Xi(\mathcal{I}^S(\rho^L,\rho^H)) = \{H,L\}$ and (ii) $\Xi(\mathcal{I}^S(\rho^L,\rho^H)) = \{H\}$. We derive the index of search for each of these cases.

Denote $\overline{\xi} \equiv \rho^L + \rho^H(1+\alpha)$, $\hat{\xi} \equiv \rho^L + \frac{\rho^H}{1+\alpha}$.

Case (i) - $\Xi(\mathcal{I}^S(\rho^L,\rho^H)) = \{H,L\}$. In this case, the index of search is

$$\mathcal{I}^S(\rho^L,\rho^H) = \frac{-\delta\left[\rho^L\left(\delta\int_{\frac{\mathcal{I}^S(m)}{1-\delta}}^{\infty} v\mathrm{d}F^L(u)\right) + \rho^H\left(\delta\int_{\frac{\mathcal{I}^S(m)}{1-\delta}}^{\infty} v\mathrm{d}F^H(u)\right)\right]}{1+\left(\rho^L\left[\delta+\frac{\delta^2}{1-\delta}\left(1-F^L\left(\frac{\mathcal{I}^S(m)}{1-\delta}\right)\right)\right]+\rho^H\left[\delta+\frac{\delta^2}{1-\delta}\left(1-F^H\left(\frac{\mathcal{I}^S(m)}{1-\delta}\right)\right)\right]\right)},$$

which in the current example, after some algebra, can be rewritten as

$$\mathcal{I}^S(\rho^L,\rho^H) = \frac{-(1-\delta)^2\left(-\frac{\delta^2}{2}\overline{\xi}\right) - \frac{\delta^2}{2}\left(\mathcal{I}^S(m)\right)^2\hat{\xi}}{1-\delta-\delta^2\hat{\xi}\mathcal{I}^S(m)}.$$

Solving for $\mathcal{I}^S(\rho^L,\rho^H)$, we have that in case (i),

$$\begin{aligned}\mathcal{I}^S(\rho^L,\rho^H) &= \frac{1-\delta}{\delta^2\hat{\xi}}\left(1-\sqrt{1-\delta^4\hat{\xi}\overline{\xi}}\right)\\ &= \frac{(1-\delta)(1+\alpha)}{\delta^2\left(1+\alpha\rho^L\right)}\left(1-\sqrt{1-\delta^4\frac{\left(1+\alpha\rho^L\right)\left(1+\rho^H\alpha\right)}{1+\alpha}}\right)\end{aligned}$$

Case (ii) - $\Xi(\mathcal{I}^S(\rho^L,\rho^H)) = \{H\}$. In this case,

$$\mathcal{I}^S(\rho^L,\rho^H) = \frac{\delta^2\rho^H\left(\int_{\frac{\mathcal{I}^S(m)}{1-\delta}}^{\infty} v\mathrm{d}F^H(u)\right)}{1+\rho^H\left[\delta+\frac{\delta^2}{1-\delta}\left(1-F^H\left(\frac{\mathcal{I}^S(m)}{1-\delta}\right)\right)\right]},$$

which, after some algebra, can be written as

$$\mathcal{I}^S(m) = \frac{\frac{1}{2}\delta^2(1-\delta)^2\rho^H(1+\alpha) - \frac{\delta^2\rho^H}{2(1+\alpha)}\left(\mathcal{I}^S(m)\right)^2}{(1-\delta)^2+\rho^H\left(\delta(1-\delta)-\delta^2\left(\frac{\mathcal{I}^S(m)}{1+\alpha}\right)\right)}.$$

Solving for $\mathcal{I}^S(\rho^L, \rho^H)$, we have that in case (ii),

$$\mathcal{I}^S(m) = \frac{(1-\delta)(1+\alpha)}{\delta^2 \rho^H}\left(1-\delta+\delta\rho^H - \sqrt{(1-\delta+\delta\rho^H)^2 - \delta^2(\rho^H)^2}\right).$$

Now, case (i), $\Xi(\mathcal{I}^S(\rho^L, \rho^H)) = \{H, L\}$, is the relevant case if and only if

$$\mathcal{I}^P(\emptyset, L) = \frac{1-\delta}{\delta}\left(1-\sqrt{1-\delta^2}\right) > \frac{1-\delta}{\delta^2 \hat{\xi}}\left(1-\sqrt{1-\delta^4 \hat{\xi}\bar{\xi}}\right),$$

where recall that the RHS of the latter inequality is the index of search in case (i). The latter inequality can equivalently be written as

$$\delta\left(\frac{1+\alpha\rho^L}{1+\alpha}\right)\left(1-\sqrt{1-\delta^2}\right) > 1-\sqrt{1-\frac{\delta^4}{1+\alpha}\left(1+\alpha\rho^L\right)\left(1+\alpha\rho^H\right)}. \tag{S.26}$$

We have therefore shown that the search index is equal to

$$\frac{(1-\delta)(1+\alpha)}{\delta^2(1+\alpha\rho^L)}\left(1-\sqrt{1-\frac{\delta^4}{1+\alpha}\left(1+\alpha\rho^L\right)\left(1+\alpha\rho^H\right)}\right)$$

if (S.26) holds, and otherwise it is equal to

$$\frac{(1-\delta)(1+\alpha)}{\delta^2 \rho^H}\left(1-\delta+\delta\rho^H - \sqrt{(1-\delta+\delta\rho^H)^2 - \delta^4(\rho^H)^2}\right).$$

■

Given the lemma above, the CTRs, VPCs, and eventual purchase decisions, are pinned down by the results in the subsections above.

S.4.4 Detrimental effect of additional ad space - comparative statics

Consider the setting described in Section 4.2 in the main body. The following result illustrates how the increase in the probability that search brings an additional product by firm ξ may reduce the index of search, inducing the consumer to visit the website of one of firm ξ's competitors before searching for new products. When strong enough, such an effect may reduce the probability that one of firm ξ's product is selected, and hence firm ξ's profits.

Proposition S.4. *Consider the environment described at the end of Section 4.2 in the main body. An increase in the probability ρ^ξ that search brings an additional product from firm ξ may reduce firm ξ's ex-ante expected profits.*

Proof of Proposition S.4. Suppose that each F^ξ is a Bernoulli distribution assigning probability p^ξ

to $v = \hat{v}^\xi$ and $(1-p^\xi)$ to $v = 0$, with $\hat{v}^\xi \in \mathbb{R}_{++}$.[14] Each firm makes equal profits on each of its two products. Hence, each firm's ex-ante total profits are equal to the total probability with which one of its two products is selected. To keep things simple, suppose the consumer incurs no cost for inspecting any product other than the time cost of postponing the final purchase: that is, $\lambda^\xi = 0$ for $\xi = A, B, C$. The consumer's discount factor is δ.

Exogenous CS. Suppose the identity of the firm receiving the additional slot is determined ex-ante, i.e., before the consumer starts the exploration process. Given the composition of the CS, the consumer then sequentially decides which product to inspect and when to stop, at which point she either chooses one of the inspected products or her outside option (whose value is normalized to zero). As shown in the main text, the reservation prize for each ξ's product, before the latter is inspected, is equal to

$$\mathcal{I}(\xi, \emptyset) = \frac{(1-\delta)\delta p^\xi \hat{v}^\xi}{1-\delta+\delta p^\xi},$$

whereas the reservation prize of each ξ's product after it is inspected is equal to $\mathcal{I}(\xi, v) = (1-\delta)v$, with $v \in \{\hat{v}^\xi, 0\}$. The optimal policy is to inspect products in descending order of their reservation prizes, stopping when the remaining reservation prizes are all smaller than the maximal realized value among the inspected products. Clearly, in this environment, each firm benefits from an increase in the probability it is given a second slot.

Endogenous CS. Now suppose that the consumer's initial CS consists of three products, one from each firm $\xi = A, B, C$, and that the CS can be expanded only once, with the expansion bringing an additional product drawn from Ξ according to ρ, with $\rho^\xi \geq 0$, $\xi = A, B, C$, and with $\sum_\xi \rho^\xi = 1$. The result in the corollary then follows from Claim S.1 below.

Claim S.1. Suppose that $\mathcal{I}(A, \emptyset) > \mathcal{I}(B, \emptyset) > \mathcal{I}(C, \emptyset)$. There exist parameter values consistent with the above inequalities such that an increase in ρ^B, together with a reduction by the same amount in ρ^A, leads to a decrease in the overall probability that one of firm B's products is sold (and hence in its ex-ante expected profits).

Proof. We establish the claim above by showing that an increase in the probability that search brings an extra B-product (along with a reduction by the same amount in the probability that it brings an A-product) may reduce the attractiveness of search thus inducing the consumer to inspect firm C's product before expanding the CS. We show that this effect may imply a drop in firm B's ex-ante profits.

It is easy to verify that the index for search is equal to

$$\mathcal{I}^S = \delta^2 \max_{k \in \{A,B,C\}} \left\{ \frac{\sum_{\xi \in \{\xi' \in \Xi : \mathcal{I}(\xi', \emptyset) \geq \mathcal{I}(k, \emptyset)\}} \rho^\xi p^\xi \hat{v}^\xi}{1 + \sum_{\xi \in \{\xi' \in \Xi : \mathcal{I}(\xi', \emptyset) \geq \mathcal{I}(k, \emptyset)\}} \rho^\xi \delta \left(1 + \frac{p^\xi \delta}{1-\delta}\right)} \right\}. \tag{S.27}$$

[14] One can think of $\hat{v}^\xi$ as the value (net of price) to the consumer in case the product is a good match, and p^ξ as the probability of such an event.

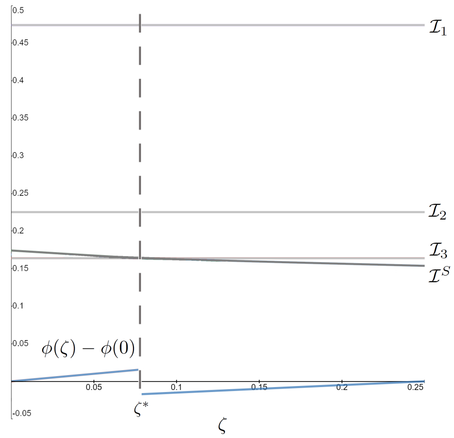

Figure 1: The change $\phi(\zeta) - \phi(0)$ in the probability with which a product of firm B is selected, as a function of ζ (in blue).

For concreteness, let $\delta = 0.9$ and suppose that $(\hat{v}^A, p^A) = (10, \frac{1}{10})$, $(\hat{v}^B, p^B) = (3, \frac{1}{3})$, and $(\hat{v}^C, p^C) = (2, \frac{1}{2})$. Note that the distributions F^ξ from which the consumer's values for the firms' products are drawn have the same mean, but are mean preserving spreads of one another; hence $\mathcal{I}(A, \emptyset) > \mathcal{I}(B, \emptyset) > \mathcal{I}(C, \emptyset)$. Suppose that, initially, $\rho^A = \rho^B = \frac{1}{4}$, and $\rho^C = \frac{1}{2}$. It is easily verified that $\mathcal{I}(A, \emptyset) = 0.473$, $\mathcal{I}(B, \emptyset) = 0.225$, $\mathcal{I}(C, \emptyset) = 0.163$, and $\mathcal{I}^S = 0.174$. Also note that $\mathcal{I}(C, \emptyset) < \mathcal{I}^S < \mathcal{I}(B, \emptyset)$, so that $\mathcal{I}^S$ does not take into account the benefits from inspecting firm C's additional product, in case search brings a second product by firm C.

Now suppose ρ^B is increased by $\zeta \in [0, 0.25]$ while ρ^A is reduced by the same amount. Let $\phi(\zeta)$ denote the probability that one of firm B's products is ultimately chosen when the probability that search brings a B-product is $\rho^B + \zeta$. Figure 1 depicts the change $\phi(\zeta) - \phi(0)$ in the probability that one of firm B's products is selected as a function of ζ, where $\phi(0) = (1 - p^A)(p^B + (1 - p^B)\rho^B p^B) = 0.35$. The horizontal gray lines correspond to the indices $\mathcal{I}(A, \emptyset)$, $\mathcal{I}(B, \emptyset)$, and $\mathcal{I}(C, \emptyset)$, whereas the dark gray curve depicts $\mathcal{I}^S$, as a function of ζ. Note that $\mathcal{I}^S$ is decreasing in ζ, since $\mathcal{I}(A, \emptyset) > \mathcal{I}(B, \emptyset)$. Hence, an increase in ζ implies a lower index for search. $\mathcal{I}^S$ starts out above $\mathcal{I}(C, \emptyset)$, and intersects $\mathcal{I}(C, \emptyset)$ at an interior ζ (smaller than 0.25), denoted by ζ^* (the vertical dashed line). For $\zeta < \zeta^*$, $\mathcal{I}(C, \emptyset) < \mathcal{I}^S < \mathcal{I}(B, \emptyset)$, whereas for $\zeta > \zeta^*$, $\mathcal{I}^S < \mathcal{I}(C, \emptyset)$. The function $\mathcal{I}^S(\zeta)$ has a kink at $\zeta = \zeta^*$. For $\zeta \in [0, \zeta^*)$, the CS is expanded before firm C's product is inspected, whereas for $\zeta \in (\zeta^*, 0.25]$ the opposite is true. Therefore, the probability that one of firm B's products is chosen is equal to $\phi(\zeta) = (1 - p^A)(p^B + (1 - p^B)(\rho^B + \zeta)p^B)$ for $\zeta \in [0, \zeta^*)$ and is equal to $\phi(\zeta) = (1 - p^A)\left(p^B + (1 - p^B)(1 - p^C)(\rho^B + \zeta)p^B\right)$ for $\zeta \in (\zeta^*, 0.25]$, with a downward discontinuity at $\zeta = \zeta^*$ equal to $(1 - p^A)(1 - p^B)p^B p^C(\rho^2 + \zeta^*)$. Furthermore, the downward drop in $\phi(\zeta)$ at $\zeta = \zeta^*$ makes $\phi(\zeta) - \phi(0)$ negative over $(\zeta^*, 0.25]$, thus establishing the claim above. $\square$

S.5 Irreversible choice

Consider the following amendment to the general model of Section 2 in the main text. At any period t, in addition to exploring an alternative in the CS or expanding the latter, the DM can *irreversibly commit* to any alternative in the CS, provided that the alternative has been explored at least M_ξ times (with ξ denoting the alternative's category).[15] Once the DM irreversibly commits to an alternative, there are no further decisions to be made. Irreversibly committing to an alternative yields a flow payoff to the DM from that moment onward, the value of which may be only imperfectly known to the DM at the time the irreversible decision is made. In particular, denote by $R(\omega^P)$ the *expected flow value* from irreversibly committing to an alternative whose current state is $\omega^P = (\xi, \theta)$. Note that $R(\omega^P)$ admits two equivalent interpretations: (i) the DM obtains an immediate expected payoff equal to $R(\omega^P)/(1-\delta)$ after which there are no further payoffs; (ii) payoffs continue to accrue at all subsequent periods after the irreversible choice is made, with each expected flow payoff equal to $R(\omega^P)$.

For any $\omega^P = (\xi, \theta)$ and $\hat{\omega}^P = (\hat{\xi}, \hat{\theta})$, say that $\hat{\omega}^P$ "follows" ω^P if and only if $\hat{\xi} = \xi$, $\theta = (\vartheta_1, ..., \vartheta_m)$, for some m, and $\hat{\theta} = (\vartheta_1, ..., \vartheta_m, ..., \vartheta_{\hat{m}})$ for some $\hat{m} \geq m$. Denote this relation by $\hat{\omega}^P \succeq \omega^P$.

Condition 1. A category-ξ alternative has the *better-later-than-sooner property* if, for any $\omega^P = (\xi, \theta)$ such that $\theta = (\vartheta_1, ..., \vartheta_m)$, with $m \geq M_\xi$, and any $\hat{\omega}^P \succeq \omega^P$, either $R(\hat{\omega}^P) \geq R(\omega^P)$, or $R(\hat{\omega}^P), R(\omega^P) \leq 0$.

The following environments are examples of settings satisfying Condition 1.

Example S.1 (**Weitzman's generalized problem**). Consider the following extension of Weitzman's original problem: (i) The set of boxes is endogenous; (ii) each category-ξ box requires M_ξ explorations before the box's value is revealed; (iii) the DM can irreversibly commit (i.e., select) a box only if its value has been revealed, i.e., only after M_ξ explorations, where M_ξ can be stochastic; (iv) the flow payoff from exploring a box without committing to it is equal to the cost of exploring the box (with the latter evolving stochastically based on the number of past explorations) and is equal to zero for any exploration $t > M_\xi$; (v) the payoff $R(\omega^P)$ from irreversibly committing to a box whose value has been revealed (i.e., after the M_ξ-th exploration) remains constant after the M_ξ-th exploration and is equal to the box's prize.

Example S.2 (**Purchase/Lease problem**). In each period, an apartment owner either chooses one of the real-estate agents she knows to lease her apartment, or searches for new agents. In addition, the owner can use one of the agents to sell the apartment. The decision to sell the apartment is irreversible. Once the apartment is sold, the owner's problem is over. The (expected) flow value u_{jt} the owner assigns to leasing the apartment through agent j of category ξ in state $\omega^P = (\xi, \theta)$ is a function of the information $\theta = (\vartheta_1, ..., \vartheta_m)$ the owner has accumulated over time about agent j's ability to deal with all sorts of problems related to tenants. The (expected) value $R(\omega^P)$ the owner

[15] If $M_\xi = 0$, the DM can irreversibly commit to any ξ-alternative without first exploring it.

assigns to selling the apartment through the same agent may also depend on the agent's expertise with tenant-related problems but is primarily a function of the familiarity the agent has with the apartment, which is determined by the number of times m the agent has been hired by the owner in the past. If the agent has no or little past experience selling apartments, $R(\omega^P) \leq 0$. Else, for any $\theta = (\vartheta_1, ..., \vartheta_m)$ and $\hat{\theta} = (\vartheta_1, ..., \vartheta_m, ..., \vartheta_{\hat{m}})$ such that $\hat{m} \geq m$, $R(\xi, \hat{\theta}) \geq R(\xi, \theta)$. Contrary to Weitzman's generalized problem above, the DM may derive a higher (expected) value from using an alternative without irreversibly committing to it (i.e, from leasing instead of selling) for an arbitrary long, possibly infinite, number of periods.

To accommodate for irreversible choice, we need to modify the definition of the index of each alternative in state $\omega^P \in \Omega^P$ as follows:

$$\mathcal{I}^P(\omega^P) \equiv \sup_{\pi,\tau} \frac{\mathbb{E}^\pi\left[\sum_{s=0}^{\tau-1} \delta^s U_s | \omega^P\right]}{\mathbb{E}^\pi\left[\sum_{s=0}^{\tau-1} \delta^s | \omega^P\right]}, \tag{S.28}$$

where τ is a stopping time, and where π is a rule specifying whether the DM explores the alternative, or irreversibly commits to it. Similarly, modify the index of search $\mathcal{I}^S(\omega^S)$ by letting the rule π now specify not only whether the DM keeps searching or explores one of the alternatives brought to the CS through search, but also whether or not she irreversibly commits to one of the alternatives that the new search brought to the CS.

Next, amend the definition of the index policy χ^* as follows. At each period $t \geq 0$, given the state $\mathcal{S}_t$ of the decision problem, the policy specifies to (a) search if $\mathcal{I}^S$ is greater than the index $\mathcal{I}^P$ of any alternative in the CS and the expected "retirement" value R of each alternative in the CS; (b) experiment with an alternative in state ω^P if its index $\mathcal{I}^P$ is greater than its expected retirement value R, as well as the index of search, and both the index and the expected retirement value of any other alternative in the CS; (c) choose (i.e., irreversibly commit to) an alternative in state ω^P if its retirement value R is greater than its index $\mathcal{I}^P$, as well as the index of search and both the index and the expected retirement value of any other alternative in the CS.

We then have the following result:

Theorem S.1 (Indexability with irreversible choice)**.** *Suppose Condition 1 is satisfied for all $\xi \in \Xi$. The conclusions in Theorem 1 in the main text apply to the problem with irreversible choice under consideration. However, the stopping time τ^* in the characterization of the index of search is now the first time (strictly above the one at which the index is computed) at which $\mathcal{I}^S$, all the indexes of the alternatives brought to the CS by search, and all retirement values of such alternatives fall below the value $\mathcal{I}^S(\omega^S)$ of the search index when the latter is computed.*

The result is established by considering a fictitious problem without irreversible choice in which, each time the DM experiments with an alternative and changes its state to ω^P, an "auxiliary alternative" with constant flow payoff equal to $R(\omega^P)$ is added to the CS and remains available in all subsequent periods, irrespectively of possible changes in the state of the alternative that generated it. The better-later-than-sooner property of Condition 1 guarantees that, if the DM ever selects one of these auxiliary alternatives, she necessarily picks the one corresponding to the latest exploration of the alternative that generated it. This last property in turn implies that both (a) the non-perishability of the auxiliary alternatives and (b) the reversibility of choice in the fictitious problem play no role, which in turn implies that the optimal policy in the fictitious problem coincides with the one in the primitive problem.

Proof of Theorem S.1. To ease the notation, assume the initial CS is empty. It will be evident from the arguments below that the optimality of χ^* does not hinge on this assumption. Consider first an environment where $M_\xi = 0$ for all ξ. It will also become evident from the arguments below that the result easily extends to environments where $M_\xi > 0$, as well as to environments where M_ξ is stochastic and learned over time.

Consider the following *fictitious environment*, where all choices are *reversible*. Whenever an alternative of category ξ is brought to the CS, an additional *auxiliary* alternative is also introduced into the CS, yielding a fixed flow payoff of $R(\xi, \emptyset)$.[16] Furthermore, whenever a non-auxiliary alternative in state ω^P is explored, a new auxiliary alternative yielding a fixed payoff of $R(\tilde{\omega}^P)$ is also added to the CS, where $\tilde{\omega}^P$ denotes the new state of the explored alternative drawn from H_{ω^P}, as in the baseline model.[17] We say that an auxiliary alternative *corresponds to a (non-auxiliary) alternative in state* ω^P if it has been introduced to the CS as the result of either search (in which case $\theta = \emptyset$) or the exploration of an alternative in state ω^P. In this auxiliary environment, define the index of search as in the main text, with the rule π specifying whether to keep searching or exploring one of the alternatives introduced through search, including the auxiliary alternatives brought to the CS by search or by the explorations of the alternatives brought to the CS through search. For each alternative in state ω^P, define its new index as in (S.28), with the rule π in the definition of the index specifying for each period prior to stopping whether to explore the alternative itself or one of the auxiliary alternatives introduced as the result of the alternative's current and future explorations (i.e., following the period at which the index is computed; importantly, π excludes any auxiliary alternative introduced in periods prior to the one in which the index is computed). Finally, let the index of any auxiliary alternative coincide with the alternative's retirement value, as specified by the function R.

It is easy to see that the same steps as in the proof of Theorem 1 in the main text imply that, in this auxiliary environment, the index policy based on the above new indices is optimal.[18] It is also easy to see that the DM's problem in the auxiliary environment is a relaxation of the problem in the primitive environment in which (a) all decisions are reversible, and (b) alternatives can be retired also in states that are not feasible any more due to the subsequent explorations of the same alternative. Hereafter, we argue that the DM's payoff in the primitive environment under the proposed index policy is the same as under the corresponding index policy in the fictitious

[16] Recall that $R(\xi, \emptyset)$ is the retirement value of a physical alternative of category ξ that has never been explored.

[17] If $M_\xi > 0$, the introduction of the auxiliary alternative as the result of the exploration of an alternative in state $\omega^P = (\xi, \theta)$ occurs only if $\theta = (\vartheta_1, ..., \vartheta_s)$ with $s \geq M_\xi$, that is, only if the alternative has been explored at least M_ξ times.

[18] The proof must be adjusted to accommodate for the auxiliary alternatives introduced as the result of the DM exploring the physical alternatives. Since all the steps are virtually the same, the proof is omitted.

environment. To see this, first observe that, in the fictitious environment, once the DM explores an auxiliary alternative, she continues to do so in all subsequent periods, since the indexes $R(\omega^P)$ of the auxiliary alternatives do not change. This implies that the reversibility of choice in the fictitious environment plays no role. Next, observe that Condition 1 implies that, in the fictitious environment, if the DM selects an auxiliary alternative, she always picks the one corresponding to the "newest" state of the corresponding non-auxiliary alternative that created it; this is because the latest has the highest expected value R among all the auxiliary alternatives corresponding to the same non-auxiliary alternative. This implies that the non-perishability of the older versions of the auxiliary alternatives in the fictitious environment also plays no role. The same condition also guarantees that the policy π in the definition of the index of the non-auxiliary alternatives in the fictitious problem coincides with the one in (S.28) where the selection π is restricted to be over the exploration of the non-auxiliary alternative under consideration and the retirement of the latter in its most recent state.

Finally, note that the proof immediately extends to settings in which $M_\xi > 0$ by assuming that, in the fictitious environment, an auxiliary alternative is introduced into the CS only when its corresponding non-auxiliary alternative has been explored more than M_ξ times, with M_ξ possibly stochastic and learned over time (in this latter case, the time-varying component of an alternative's state, θ, may also contain information about M_ξ). ■

S.6 Sub-optimality of index policies with "meta arms"

In this section, we briefly illustrate, by means of an example, why multi-armed bandit problems in which alternatives take the form of "meta arms", i.e., sub-decision problems with their own sub-decisions, typically do not admit an index solution. This is so even if each sub-problem is independent from the others, and even if one knows the solution to each independent sub-problem. In the same vein, dependence or correlation between alternatives typically precludes an index solution. This is the case even if a subset of dependent alternatives evolves independently of all other alternatives, and even if one knows how to optimally choose among the dependent alternatives in each given subset in isolation.

Consider the following extension of the environment described in the main text. There are $k \in \mathbb{N}$ sets of arms, $K_1, ..., K_k$. Arms from different sets evolve independently of one another, but the state of each arm within a set may depend on the state of other arms from the same set. More generally, suppose that each arm corresponds to a "meta arm", the activation of which involves decisions other than when to stop using it. Each meta arm has its own decision process which is independent of the other meta arms.

Clearly, the model in the main body of the paper is a special case of this richer setting. Suppose that, for each set of arms K_i, one can compute the optimal sequence of pulls, independently of the other sets of arms. Equivalently, suppose that for each "meta arm" one can compute the optimal sequence of decisions that define the usage of that arm, independently of the solution to the other meta arms' problems. It is tempting to conjecture that one may then assign an independent index to each set of arms K_i (alternatively, to each "meta arm") and that the optimal policy for the overall problem reduces to an index policy, whereby the meta arm with the highest index is selected in each period.

Perhaps surprisingly, the optimal policy for this enriched problem does not admit an index representation. When arms are not defined as in the canonical multi-armed bandit problem, but rather feature a more complicated internal decision problem (preserving the independence across arms), the optimal policy need not be an index policy. The following example illustrates.

Example S.3. There are two arms. Arm 1 yields a reward of 1,000 when it is first pulled. In all subsequent pulls, it yields a reward λ, where λ is initially unknown and may be either 1 or 10, with equal probability. After the first pull of arm 1, λ is perfectly revealed and is fully persistent. Arm 2 is a "meta arm" corresponding to the following decision problem. When the decision maker pulls arm 2 for the first time, she must also choose *how* to pull it. There are two ways to pull this arm, 2(A) and 2(B). If the decision maker selects $2(A)$, the arm yields a reward of 100 for a single period, followed by no rewards thereafter. If, instead, the decision maker selects $2(B)$, the arm yields a reward equal to 11 in each of its subsequent pulls. The choice of which version of arm 2 to use must be made the first time that arm 2 is pulled and can not be reversed.

Assume $\delta = 0.9$. The optimal policy for this problem is the following. In period 1, arm 1 is pulled. If $\lambda = 10$, then arm 2 in version $2(A)$ is then pulled for a single period, followed by arm 1 again in all subsequent periods. If, instead, $\lambda = 1$, arm 2 is then pulled in version $2(B)$ in all subsequent periods. Note that, under the optimal policy, the decision of how to use arm 2 depends on the realization of arm 1's first pull. It is then evident that the optimal policy is not an index policy, no matter how one defines the indices. This is because an index policy requires that both the index of each arm and its utilization (when an arm can be used in different versions, as in the case of "meta arm" 2 in this example) be invariant in the results of the activation of all other arms.