

Searching for Arms: Experimentation with Endogenous Consideration Sets*

Daniel Fershtman[†] Alessandro Pavan[‡]

March 2021

Abstract

We study the problem of a decision maker alternating between exploring existing alternatives in the consideration set and searching for new ones. We characterize the optimal policy and its key properties, and identify implications for search and exploration dynamics. When the search technology is stationary, or improves over time, search is equivalent to replacement. With deteriorating technologies, instead, alternatives are revisited after search is launched and each expansion is treated as if it were the last one. A key consequence of the endogeneity of the consideration set is that an improvement in the desirability of a category of alternatives may lead to a reduction in the exploration of the category as well as in the eventual selection of an alternative from that category. We apply the analysis to clinical trials, experimentation toward regulatory approval, and consumer search.

*The paper supersedes previous versions circulated under the titles “Searching for Arms” and “Sequential Learning with Endogenous Consideration Sets.” We thank the editor and referees for excellent comments, as well as Dirk Bergemann, Eddie Dekel, David Dillenberger, Laura Doval, Kfir Eliaz, Stephan Laueremann, Benny Moldovanu, Philip Reny, Eran Shmaya, Rani Spiegler, Bruno Strulovici, Asher Wolinsky, Jidong Zhou, and seminar participants at University of Bonn, Cornell University, Hebrew University of Jerusalem, HEC Paris, Paris Dauphine, Tel-Aviv University, Toulouse School of Economics, University of Pennsylvania, Yale University, Warwick Economic Theory Workshop, 3rd Columbia Conference on Economic Theory (CCET), and the 30th Stony Brook International Conference on Game Theory for helpful discussions and comments. We also thank Matteo Camboni and Tuval Danenberg for excellent research assistance. Fershtman gratefully acknowledges funding by the German Research Foundation (DFG) through CRC TR 224 (Project B1) and from ISF grant #1202/20. The usual disclaimer applies.

[†]Eitan Berglas School of Economics, Tel Aviv University. Email: danielfer@tauex.tau.ac.il

[‡]Department of Economics, Northwestern University. Email: alepavan@northwestern.edu

1 Introduction

An important implication of limited attention is that a decision maker (hereafter, DM) can consider only a subset of all the alternatives available to her. This scarcity gives rise to the concept of consideration sets (hereafter, CS), which has been applied to a range of economic problems.¹ Classic models of sequential experimentation or learning typically involve a DM exploring a *fixed* set of alternatives with unknown returns. Yet, a ubiquitous feature of many dynamic decision problems is that the set of alternatives a DM can explore is expanded over time, in response to the information gathered by exploring the alternatives already in the CS.

In this paper, we study the tradeoff between the exploration of alternatives already in the CS and the expansion of the latter through search for additional alternatives. A key difference between exploration and expansion is the direct-vs-indirect nature of the two activities. When an alternative is in the CS, the DM can “point to it,” that is, she can choose to explore that particular alternative instead of others. When, instead, an alternative is outside the CS, the DM cannot point to it, meaning that she cannot choose to explore that specific alternative instead of others.² This inability may reflect natural randomness in the search process, which may uncover alternatives different from those the DM was hoping for. Alternatively, search may bring more than a single alternative and such batching may have important implications for the decision to expand the CS in the first place. Finally, DM may have limited knowledge about what is outside her CS, and/or her ability to bring new alternatives to it.

To study the tradeoff between exploration of alternatives already in the CS and expansion of the latter, we consider a generalization of the classic multi-armed bandit problem in which the set of “arms” is endogenous.³ Exploring an alternative already in the CS (pulling an arm) yields a flow payoff and generates information (for example, about the distribution from which the flow payoff is drawn). Search for new “arms” (that is, choosing to expand the CS) is costly and brings a new set of alternatives (i.e., of arms) according to a process that controls for the “state of the search technology.”

We characterize the solution to the above problem and show that it takes the form of an “index” policy. Each alternative in the CS is assigned a history-dependent number that

¹The concept of consideration sets is also central to the marketing literature.

²Likewise, the DM cannot choose to bring a specific alternative from outside of the CS into the CS: If she could, there would be no distinction between exploring alternatives inside and outside the CS, making the latter irrelevant.

³See [Bergemann and Välimäki \(2008\)](#) for an overview of the multi-armed bandit literature and its applications to economics.

is a function only of the state of that alternative. This number (the arm’s “index”) is the same as in [Gittins and Jones’ \(1974\)](#) original work on bandit problems with an exogenous set of arms. Search (that is, the decision to expand the CS) is also assigned an index, which depends only on the state of the search technology. Crucially, the index of search does not depend on the information generated by the exploration of any of the alternatives already in the CS. It also differs from the value the DM attaches to the expansion of the CS but is linked to the indices of the new alternatives the DM expects to find through the current and future searches. The optimal policy consist in selecting at any period the alternative for which the index is the highest. We establish the optimality of such a policy through a novel proof based on a recursive characterization of the index for search that also permits us to identify various properties of the exploration and expansion dynamics under the optimal policy.⁴ We then show how such a characterization facilitates novel comparative statics relevant for applications.⁵

At any point in time, the decision to expand the CS depends on the composition of the current CS only through (a) the state of the alternative with the highest index, and (b) the state of the search technology.⁶ Similarly, conditional on forgoing search in a given period, the decision of which alternative to explore in the current CS is independent of the state of the search technology, despite the fact that search may bring alternatives that are more similar to certain alternatives currently in the CS than others.⁷ If the search technology is stationary or improving, in a sense made precise below, then alternatives in the CS at the time of its expansion never receive attention in the future, and hence are effectively discarded once the CS is expanded. Each search is then equivalent to replacement of the current CS with a new one. When, instead, the search technology deteriorates over time (e.g., because the DM becomes pessimistic about the possibility of finding attractive new alternatives), the alternatives in the current CS are put on hold and may be revisited after

⁴The reader may ask what makes expansion of the CS different from the exploration of a “meta” arm that comprises all the physical arms that can be brought to the CS by search. The difference is that the evaluation of such a “meta” arm requires knowing how to subsequently explore the arms that search brings to the CS, which is what is investigated in the first place. Furthermore, dynamic decision problems with “meta” arms rarely admit an index solution. We come back to this point in due course.

⁵When the optimal policy does not admit an index representation, it can still be computed using dynamic programming. However, this typically requires strong assumptions on the primitives necessary to overcome the “curse of dimensionality” that many of such problems typically suffer from.

⁶This result holds despite the fact that the alternatives in the CS may share similarities with those that are expected to be found through the expansion of the CS (e.g., through the specification of the categories such alternatives belong to), and despite the fact that the opportunity cost of searching for new alternatives (which is linked to the value of continuing with the current CS) depends on the entire composition of the CS and not only on the state of the alternative with the highest index.

⁷These properties can be seen as a generalization of the IIA (independence of irrelevant alternatives) property of classic multi-armed bandit problems to a problem with an endogenous set of alternatives.

the CS is expanded. Furthermore, in this case, the decision to expand the CS is made as if there will be no further expansion after the current one.

The analysis can be applied to a broad class of experimentation and sequential learning problems. We consider three of them in the paper: clinical trials, experimentation toward regulatory approval, and consumer search.

Clinical trials. One of the most prominent applications of the multi-armed bandit model is clinical trials.⁸ A DM (e.g., a physician, or medical facility) must choose in each period which medical treatment to administer, trading off the well-being of the patients receiving the treatments with the value of generating information about the treatments' efficacy and safety useful to treat future patients. Each treatment belongs in a different category with treatments from the same category sharing the same core characteristics. We enrich this problem by giving the DM the possibility to expand the set of treatments that can be administered (e.g., by conducting basic research in-house, or by commissioning it to external research labs).

We study the effects of an improvement in a category of treatments on the usage of the treatments from that category. The improvement may take the form of (a) an increase in the ex-ante probability that treatments from that category are effective, (b) an increase in the DM's payoff from administering a successful treatment from that category, or (c) an increase in the informativeness of the outcomes of the treatments from that category. Any of such improvements increases the index of each treatment from the improved category, without affecting the index of the treatments from other categories. If the CS were exogenous, such improvements would unambiguously increase the usage of the treatments from the category that has become more attractive. We show that this need not be the case when the CS is endogenous. The reason is that such improvements also contribute to an increase in the desirability of search, as measured by its index. This new effect, when strong enough, may crowd out the usage of the treatments that become more attractive.

That search may crowd out the usage of treatments from the improved category may be surprising given that the improvement in such treatments is the only reason for the increase in the index of search in the first place. Such crowding out is possible due to the fact that the increase in the index of the treatments that have become more attractive is not uniform across histories. On the other hand, the increase in the index of search hinges on "averaging" over the different histories at which the treatments from the improved category are administered. This implies that there can be histories at which the increase

⁸See, e.g., [Bergemann and Välimäki \(2008\)](#), [Diskstein \(2014\)](#), as well as [Katehakis and Derman \(1986\)](#), [Berry \(2006\)](#), [Villar, Bowden and Wason \(2015\)](#), and the FDA's "Guidance for the Use of Bayesian Statistics in Medical Device Clinical Trials".

in the index of search is larger than the increase in the index of the treatments that have become more attractive. In such circumstances, the DM may favor search over further exploration of such treatments. The expansion of the CS induced by search may re-balance the composition of the CS in favor of categories other than the one that experienced an increase in attractiveness. This may occur, for example, because search brings primarily treatments from these other categories. Alternatively, it may occur because the initial CS contained primarily treatments from the improved category, whereas each new search brings (on average) a treatment from each category with equal probability. In either case, search contributes to a change in the composition of the CS in favor of categories other than the one whose attractiveness improved. When this new effect is strong enough, it can lead to a *reduction* in the usage of the treatments from the improved category, not only relative to treatments from other categories, but overall.

Experimentation towards regulatory approval. Consider a pharmaceutical company seeking to persuade the FDA, or an equivalent regulatory authority, to approve one of its drugs. The safety of such drugs is unknown ex-ante, both to the firm and to the regulatory authority. At each period, the firm can conduct a public experiment on one of its products, or invest in R&D with the hope to discover new products to test in front of the regulator for approval. The approval of a product requires a large enough number of tests with a satisfactory outcome. Once a product is approved, the firm can bring it to the market and experimentation with that product is brought to an end. The firm's objective is to maximize the expected discounted payoff from the sale of its products, net of all the experimentation and search costs.

We show that a reduction in regulatory standards of approval for a given category of products may have the (perhaps unintended) effect of reducing the number of approved products from that category. As in the case of clinical trials, this is due to the fact that a reduction in a category's standard of approval increases the index of all products from that category, but also the index of search (provided of course that search brings products from that category with positive probability). We show that when the tests on the category of products whose standard of approval has been reduced are sufficiently precise, the increase in the index of search may outweigh the increase in the index of those products whose early tests yielded bad outcomes. When this is the case, the reduction in the standard of approval may lead to a reduction in the ex-ante probability with which a product from that category is approved.

Consumer search. Our third application is an extension of [Weitzman's \(1979\)](#) classic "Pandora's boxes problem" in which, at each period, the DM can expand the set

of boxes to explore, based on what she has found by opening the boxes already in the CS. As in Weitzman’s original model, it takes a single exploration of each alternative to learn its value. We derive an “eventual-purchase theorem” in the spirit of [Choi, Dai and Kim \(2018\)](#) (see also [Armstrong and Vickers, 2015](#) and [Armstrong, 2017](#)) that relates the probability with which each alternative is eventually selected to the primitives of the problem (realized values and search technology), but accounting for the endogeneity of the CS. We then apply the results to the problem of a consumer searching online for new products, alternating between (a) reading new ads (which amounts to bringing new products to the CS), (b) clicking on ads of products already in the CS (which amounts to exploring the products, given that clicking typically brings the consumer to a vendor’s website where the value of the product to the consumer is revealed), and (c) finalizing a purchase with one of the visited vendors. The results permit us to endogenize the click-through rates (CTRs) associated with the various ad positions—that is, the probability that a displayed ad is eventually clicked upon. We show that CTRs need not be monotone in the positions of the ads, even when the purchasing probabilities. Finally, we show that, when the consumer’s CS is endogenous, a firm that is already advertising on a platform may suffer from receiving additional ad space for one of its products, despite the extra ad space being unambiguously profitable when brought exogenously to the consumer’s CS.

The rest of the paper is organized as follows. The remainder of this section briefly discusses the most pertinent literature. Section 2 introduces the model. Section 3 characterizes the optimal policy and identifies key properties for the dynamics of experimentation and expansion of the CS. Sections 4, 5, and 6 contain the applications to clinical trials, experimentation toward regulatory approval, and consumer search, respectively. Section 7 discusses a few extensions. All proofs omitted in the main text are either in the Appendix at the end of the document or in the Online Supplementary Material.⁹

1.1 Related literature

The paper is part of a fast-growing literature on CSs. See, e.g., [Hauser and Wernerfelt, 1990](#) and [Roberts and Lattin, 1991](#) for the marketing literature. In economics, [Eliaz and Spiegler \(2011\)](#) study implications of different CSs on firms’ behavior, assuming such

⁹Specifically, the latter contains (a) the formal proof of the negative effects of a reduction in regulatory standards on the approval of products from a given category (formally stated in Proposition 2 in the main text), (b) the formal proof of the characterization of CTR (Proposition 4) (c) the formal illustration of the detrimental effects of additional ad space on firms’ profits (Proposition 5), (d) an extension of the main theorem to a class of problems with irreversible choice, and (e) a discussion of the suboptimality of index policies in general bandit problems with meta-arms corresponding to independent sub-problems.

sets are exogenous. [Masatlioglu, Nakajima, and Ozbay \(2012\)](#) and [Manzini and Mariotti \(2014\)](#), instead, identify CSs from choice behavior. [Caplin, Dean, and Leahy \(2018\)](#) provide necessary and sufficient conditions for rationally-inattentive agents to focus on a subset of all available choices, thus endogenizing the CSs. [Simon \(1955\)](#) considers a sequential search model, in which alternatives are examined until a “satisfying” alternative is found. [Caplin, Dean, and Martin \(2011\)](#) show that the rule in [Simon \(1955\)](#) can be viewed as resulting from an optimal procedure when there are information costs. Our analysis complements the one in this literature by providing a dynamic micro-foundation for endogenous CSs. Rather than committing to a CS up front and proceeding to evaluate its alternatives, the DM gradually expands the CS, in response to the results obtained from the exploration of the alternatives in the set.

The paper also belongs to the recent literature on sequential learning in settings in which the DM explores one alternative at a time. Most closely related are [Ke, Shen, and Villas-Boas \(2016\)](#), [Austen-Smith and Martinelli \(2018\)](#), [Ke and Villas-Boas \(2019\)](#) and [Gossner, Steiner and Stewart \(2019\)](#).¹⁰ Related are also [Che and Mierendorff \(2019\)](#), who study the optimal sequential allocation of attention to two different signal sources biased towards alternative actions, and [Liang, Mu, and Syrgkanis \(2019\)](#), who study a DM’s dynamic acquisition of information about an unknown Gaussian state, given access to multiple information sources. In these papers, the set of alternatives is fixed ex-ante. In our model, instead, the DM expands the CS over time in response to the information she collects about the alternatives already in it.

The application to experimentation towards regulatory approval in [Section 5](#) is related to [Gossner, Steiner and Stewart \(2019\)](#) and [Henry and Ottaviani \(2019\)](#). [Gossner, Steiner and Stewart \(2019\)](#) study the effect of attention manipulation on the ultimate selection among a fixed set of alternatives, in a model in which a DM is forced to focus initially on a particular alternative. The DM dynamically allocates her attention among the various alternatives until her posterior belief about the value of one of the products exceeds an exogenous threshold. Our work is complementary to theirs and brings two novel aspects to the problem: (i) the endogeneity of the CS and (ii) the analysis of the effects of variations in the regulatory standards of approval for one category of products on the approval of products from that category.

Related is also [Garfagnini and Strulovici \(2016\)](#), who study how successive (forward-looking) agents experiment with endogenous technologies. Trying “radically” new tech-

¹⁰[Ke and Villas-Boas \(2019\)](#) study optimal search before choice in a setting where the optimal policy is not indexable and identify various properties of the optimal policy.

nologies reduces the cost of experimenting with similar technologies, which effectively expands the set of affordable technologies.¹¹ [Schneider and Wolf \(2019\)](#) study the time-risk tradeoff of an agent who wishes to solve a problem before a deadline, and allocates her time between implementing a given method and developing (and then implementing) a new one. While, at a high level, the problems examined in these papers bear a resemblance to ours in that they also consider environments in which the set of alternatives expands over time, both the models and the questions addressed are fundamentally different.

[Fershtman and Pavan \(2020\)](#) study the effects of “soft” affirmative action on minority recruitment, in a setting in which the candidate pool is endogenous. The model in that paper is a special version of the one in the present paper and builds on the general analysis developed in the present paper. While that paper studies the effects of changes in the search technology on the selection of alternatives from a given category, the analysis in Sections 4 and 5 in the present paper focuses on the effects of variations in a category’s attractiveness on the usage (or ultimate selection) of alternatives from that category.

The application in Section 6 is an extension of [Weitzman’s \(1979\)](#) to a consumer search problem with an endogenous set of boxes.¹² In independent work, [Greminger \(2020\)](#) also considers such an extension. His problem is a special version of ours in which payoffs are additively separable in an observable and an unobservable component. While the analysis in [Greminger \(2020\)](#) focuses on the comparison between direct and indirect search, we focus on the implications of search for the click-through-rates of different ads’ positions in online search, as well as the detrimental effects of additional ad space on firms’ profits.

Our characterization of the optimality of an index policy is related to the branching-arm literature (e.g., [Weiss, 1988](#), [Weber, 1992](#), and [Keller and Oldale, 2003](#)). This literature studies decision problems in which arms branch into new ones. Under appropriate assumptions, the problem in the present paper is a special case of the one considered in that literature. To the best of our knowledge, our proof of indexability is new and uncovers a novel recursive representation of the index of search which is key to the predictions about exploration and expansion dynamics and for the applications we consider.

¹¹Technologies are interdependent in their environment. In particular, a radically new technology is informative about the value of similar technologies.

¹²Despite its many applications, relatively few extensions of Weitzman’s problem have been studied in the literature. Notable exceptions include [Olszewski and Weber \(2015\)](#), [Choi and Smith \(2016\)](#), and [Doval \(2018\)](#). In these papers, though, the set of boxes is fixed.

2 Model

In this section we introduce a general model of sequential experimentation with an endogenous CS. The model is abstract, but admits many applications, some of which we consider in the subsequent sections.

In each period $t = 0, 1, 2, \dots$, the DM chooses between exploring one of the alternatives within her CS and expanding the CS by searching for additional alternatives. Exploring an alternative yields a flow payoff and generates information about it. Expanding the CS yields a stochastic set of new alternatives, which are added to the CS and can be explored in subsequent periods.

Decisions: exploration and expansion. Denote by $C_t \equiv (0, \dots, n_t)$ the period- t CS, with $n_t \in \mathbb{N}$. The latter comprises all alternatives $i = 0, \dots, n_t$ that can be explored in period t , with the initial set $C_0 \equiv (0, \dots, n_0)$ specified exogenously and with alternative $i = 0$ corresponding to the selection of the DM's outside option, yielding a payoff normalized to zero.

Let $x_{it} \in \{0, 1\}$ denote the decision to explore alternative i in period t , with $x_{it} = 1$ if the alternative is explored and $x_{it} = 0$ otherwise. Let $x_t \equiv (x_{jt})_{j=0}^{\infty}$ denote a complete description of which alternatives are explored in period t and $x \equiv (x_s)_{s=0}^{\infty}$ a complete sequence of exploration decisions. Given C_t , expansion of the CS in period t (that is, search) brings a stochastic set of new alternatives $C_{t+1} \setminus C_t = (n_t + 1, \dots, n_{t+1})$ which are added to the current CS and expand the latter from C_t to C_{t+1} . Expansion involves a direct cost $c_t \geq 0$. Let $y_t \in \{0, 1\}$ denote the decision to expand the CS in period t , with $y_t = 1$ if the DM searches for new alternatives and $y_t = 0$ otherwise. Denote by $y \equiv (y_s)_{s=0}^{\infty}$ a complete sequence of expansion decisions. The period- t complete description of exploration/expansion decisions is summarized by $d_t \equiv (x_t, y_t)$, whereas a complete sequence of exploration/expansion decisions is denoted by $d \equiv (d_s)_{s=0}^{\infty}$. A sequence of decisions d is *feasible* if, and only if, for all periods $t \geq 0$, (i) $x_{jt} = 1$ only if $j \in C_t$, and (ii) $\sum_{j=0}^{\infty} x_{jt} + y_t = 1$.

Categories, learning, payoffs, and states. Each alternative belongs to a fixed (i.e., time-invariant) *category* $\xi \in \Xi$ that is observable to the DM and that completely characterizes the alternative's experimentation technology and payoff process. Let $\mu \in \mathbb{R}$ denote a fixed (i.e., time-invariant) dimension about the alternative that the DM is learning about, with μ drawn from a distribution Γ_ξ . When the DM explores the alternative, she observes a signal realization about μ and updates her beliefs about μ using Bayes' rule. Let $\theta^m \equiv (\theta_s)_{s=1}^m$ denote a history of signal realizations of length $m \in \mathbb{N}$ and θ a

generic sequence of signal realizations of arbitrary length.

Denote by $\omega^P = (\xi, \theta) \in \Omega^P$ an alternative’s “state”.¹³ The initial state of each alternative from category ξ (before the DM explores it) is (ξ, \emptyset) , where \emptyset denotes the null history of signal realizations. Hence, the component ξ of an alternative’s “state” (its category) is fixed and summarizes all information that the DM possesses from the outset about the alternative, prior to its exploration. The component θ , instead, is time-variant and summarizes all the information that the DM has received about the alternative over time, through its past explorations.¹⁴ Instead of introducing a collection of kernels describing the conditional distributions from which the marginal signals are drawn, given the past signals realization, we find it more parsimonious (but equally flexible) to assume that each alternative’s state ω^P evolves 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.¹⁵ When, instead, 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. Importantly, signal realizations are drawn independently across alternatives, given the alternatives’ categories.

The flow payoff u that the DM obtains from exploring an alternative currently in state ω^P is also drawn from a distribution that is fully determined by the alternative’s current state, ω^P , and that is time-homogeneous. We do not introduce any notation for such distributions because they are not directly used in the analysis below.

Example. To fix ideas, consider the canonical multi-armed-bandit problem with a fixed and exogenous CS. Each alternative is an “arm” and its category ξ summarizes all the arm’s time-invariant information. Whenever the arm is pulled, it generates a flow payoff (e.g., a cash flow) u drawn from a distribution that depends on both the arm’s category ξ and a parameter μ unknown to the DM. For example, the distribution could be a Normal distribution of known variance σ_ξ and unknown mean μ , with the latter drawn from the distribution Γ_ξ . In this example, the signals the DM receives about μ coincide with the flow payoffs, but this need not be the case in other applications.

The search (expansion) technology. The state of the search technology is sum-

¹³The superscript P is meant to highlight the fact that this is the state of a “physical” alternative in the CS, not the state of the search technology, or the overall state of the decision problem, defined below.

¹⁴As shown below, when the CS is endogenous, the outcome of each CS’s expansion may depend on the categories of the alternatives added to the CS through previous expansions. If the CS was exogenous, without loss of generality, one could always take each alternative’s category to coincide with its “name”.

¹⁵Clearly, because each alternative’s category ξ is time-invariant, given the current state $\omega^P = (\xi, \theta)$, the distribution H_{ω^P} assign probability one to states whose category is ξ and whose signal history $\theta' \succeq \theta$ is a “follower” of θ , meaning that is obtained by adding a new signal realization to the history θ .

marized by the history $\omega^S = ((c_0, E_0), (c_1, E_1), \dots, (c_m, E_m))$ of past search outcomes, with each outcome (c_k, E_k) , $k = 0, \dots, m$, consisting of the cost $c_k \in \mathbb{R}$ incurred at the k -th search and the number of alternatives $E_k = (n_k(\xi) : \xi \in \Xi)$ of each category ξ added to the CS at the k -th search. Here $m \in \mathbb{N}$ denotes the number of times search has been carried out in the past, and, for each category ξ , the component $n_k(\xi) \in \mathbb{N}$ of the vector E_k represents the number of category- ξ alternatives found as the result of the k -th search.¹⁶

Denote the set of possible states of search by Ω^S . 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 $H_{\omega^S} \in \Delta(\Omega^S)$. Note that this formulation allows the search technology to depend in flexible ways on the results of previous searches. The formulation also permits us to keep track of all the results of past searches in a concise way. The key assumptions are that 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.

The state of the decision problem. The state of the decision problem is given by the pair $\mathcal{S} \equiv (\omega^S, \mathcal{S}^P)$, where $\mathcal{S}^P : \Omega^P \rightarrow \mathbb{N}$ is a counting function describing, for each $\omega^P \in \Omega^P$, the number of alternatives in the CS in state ω^P . Hence, \mathcal{S}^P can be viewed as the overall state of the current CS. Let $\Omega \equiv \Omega^P \cup \Omega^S$.¹⁷ Abusing notation, we will sometimes find it useful to denote the state of the decision problem by a function $\mathcal{S} : \Omega \rightarrow \mathbb{N}$ 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 . We will then denote by \mathcal{S}_t the state of the decision problem at the beginning of period t .¹⁸ This representation of the state of the decision problem keeps track of all relevant information in a parsimonious way and, as will become clear below, greatly facilitates the analysis.

Policies. A *policy* χ for the decision problem described above is a rule specifying, for each period t and each period- t state \mathcal{S}_t , whether to experiment with one of the alternatives in the CS or expand the latter. Denote by u_{jt} the flow period- t payoff from alternative $j = 0, \dots, \infty$ and by $U_t \equiv \sum_{j=0}^{\infty} x_{jt}u_{jt} - c_t y_t$ the realized period- t flow payoff (including the cost of search) given the period- t decision $d_t = (x_t, y_t)$. A policy χ is *optimal* if it maximizes the expected discounted sum $\mathbb{E}^\chi [\sum_{t=0}^{\infty} \delta^t U_t | \mathcal{S}_0]$ of the flow payoffs,

¹⁶The initial state is (c_0, E_0) , where c_0 can be taken arbitrarily as it plays no role in the analysis (the cost of the first search is c_1), and with E_0 describing the composition of the initial CS, C_0 .

¹⁷Note that $\Omega^P \cap \Omega^S = \emptyset$.

¹⁸Clearly, with this representation, at each period t , there is a unique $\hat{\omega}^s \in \Omega^S$ such that $\mathcal{S}_t(\omega^S) = 1$ if $\omega^s = \hat{\omega}^s$ and $\mathcal{S}_t(\omega^S) = 0$ 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 periods t and all $\omega^S \in \Omega^S$, $\mathcal{S}_t(\omega^S) = 0$.

where $\delta \in (0, 1)$ denotes the discount factor, and $\mathbb{E}^\chi[\cdot|\mathcal{S}_0]$ denotes the expectation under the endogenous process for the flow payoffs U_t obtained by starting from state \mathcal{S}_0 and following the policy χ at each period. Clearly, because the entire decision problem is time-homogeneous, so is the optimal policy (that is, for any two periods t and t' such that $\mathcal{S}_t = \mathcal{S}_{t'}$, the decisions specified by the optimal policy for the two periods are the same).

To guarantee that the process of the expected payoffs is well behaved, we assume that, for any state \mathcal{S} and policy χ , $\delta^T \mathbb{E}^\chi [\sum_{s=T}^\infty \delta^s U_s | \mathcal{S}] \rightarrow 0$ as $T \rightarrow \infty$.¹⁹

Remark 1. In certain applications, θ may have interpretations other than signals about a fixed unknown parameter μ . In the Appendix, we describe a more general environment which imposes fewer assumptions on the states of the alternatives and their evolution as a result of experimentation.

Remark 2. The model above describes an experimentation environment in which payoffs are accumulated alongside learning, and the problem need not necessarily end with a particular decision made by the DM. In Sections 5-6 (and more generally in the Online Supplementary Material), we also consider applications in which the DM sequentially decides between learning about alternatives in her CS and expanding the CS, until a final choice is made among the alternatives in the CS, ending the decision problem.

3 Optimal Policy and Key Implications

In this section, we characterize the optimal policy and its implications for experimentation and search dynamics.

3.1 Optimal Policy

For each $\omega^P \in \Omega^P$, let²⁰

$$\mathcal{I}^P(\omega^P) \equiv \sup_{\tau > 0} \frac{\mathbb{E} \left[\sum_{s=0}^{\tau-1} \delta^s u_s | \omega^P \right]}{\mathbb{E} \left[\sum_{s=0}^{\tau-1} \delta^s | \omega^P \right]}, \quad (1)$$

denote the “index” of an alternative in the CS currently in state ω^P , where τ denotes a stopping time (that is, a rule prescribing when to stop, as a function of the observed signal realizations), and where u_s denotes the flow payoff from the alternative’s s -th utilization.

¹⁹This property guarantees that the solution to the Bellman equation of the above dynamic program coincides with the true value function; it is immediately satisfied if payoffs and costs are uniformly bounded.

²⁰The expectation here is under the process obtained under a rule that selects the alternative under consideration in all periods.

The definition in (1) is equivalent to the definition in [Gittins and Jones \(1974\)](#) and [Gittins \(1979\)](#). As is well known, the optimal stopping rule in the definition of the index is the first time at which the index falls below the value at the time the index was computed (see, e.g., [Mandelbaum, 1986](#)).

Given each state $\mathcal{S} = (\omega^S, \mathcal{S}^P)$ of the decision problem, denote the maximal index among the alternatives within the CS by $\mathcal{I}^*(\mathcal{S}^P) = \max_{\omega^P \in \{\hat{\omega}^P \in \Omega^P: \mathcal{S}^P(\hat{\omega}^P) > 0\}} \mathcal{I}(\omega^P)$.

We now define an index for search (i.e., for the option to expand the CS). This index is independent of the state of each alternative in the CS, conditional on the state of the search technology ω^S .²¹ Analogously to the indices defined above, the index for search is defined as the maximal expected average discounted net payoff, per unit of expected discounted time, obtained between the current period and an optimal stopping time. Contrary to the standard indices, however, the maximization is not just over the stopping time, but also over the rule governing the selection among the new alternatives brought to the CS by search and further searches. Denote by τ a stopping time, and by π a rule prescribing, for any period s between the current one and the stopping time τ , either the selection of one of the new alternatives brought to the CS by search or further search. Importantly, π selects only among search and alternatives that were not already in the CS when the decision to search was made.²²

Formally, given the state of the search technology $\omega^S \in \Omega^S$, the index for search is defined by

$$\mathcal{I}^S(\omega^S) \equiv \sup_{\pi, \tau} \frac{\mathbb{E}^\pi \left[\sum_{s=0}^{\tau-1} \delta^s U_s | \omega^S \right]}{\mathbb{E}^\pi \left[\sum_{s=0}^{\tau-1} \delta^s | \omega^S \right]}. \quad (2)$$

Definition 1. Let χ^* be the *index policy* selecting at each period t the option with the greatest index given the overall state \mathcal{S}_t of the decision problem: search if $\mathcal{I}^S(\omega^S) \geq \mathcal{I}^*(\mathcal{S}^P)$, and an arbitrary alternative with index $\mathcal{I}^*(\mathcal{S}^P)$ if $\mathcal{I}^S(\omega^S) < \mathcal{I}^*(\mathcal{S}^P)$.²³ Ties between alternatives are broken arbitrarily. In order to maintain consistency throughout the analysis, we assume that, when $\mathcal{I}^S(\omega_t^S) = \mathcal{I}^*(\mathcal{S}_t^P)$, search is carried out.

To characterize the optimal policy, we first introduce the following notation. Let $\kappa(v) \in \mathbb{N} \cup \{\infty\}$ denote the first time at which, when the DM follows the index policy χ^* , (a) the search technology reaches a state in which its index is no greater than v , and

²¹That is, the index depends on the state of each alternative in the CS only through the information that the latter state contains for the state ω^S of the search technology.

²²Suppose the index for search is computed in period t when the state of the search technology is ω^S . Then, for each period $t < s < \tau$, π selects between further search and the selection of alternatives in the CS at period s that were not in the CS in period t .

²³Recall that $\mathcal{I}^*(\mathcal{S}^P)$ is the largest index among the alternatives in the CS.

(b) *all* alternatives in the CS – regardless of when they were introduced into it – have an index no greater than v . That is, $\kappa(v)$ is the minimal number of periods until all indices are weakly below v . In case this event never occurs, $\kappa(v) = \infty$.²⁴

Let $\mathcal{V}^*(\mathcal{S}_0) = (1 - \delta) \sup_{\chi} \mathbb{E}^{\chi} [\sum_{t=0}^{\infty} \delta^t U_t | \mathcal{S}_0]$ denote the maximal expected per-period payoff the DM can attain across all feasible policies χ , given the initial state \mathcal{S}_0 .

Theorem 1. (i) *The index policy χ^* is optimal in the sequential experimentation problem with endogenous CS.*

(ii) *The index for search, as defined in (2), admits the following recursive representation. For any $\omega^S \in \Omega^S$,*

$$\mathcal{I}^S(\omega^S) = \frac{\mathbb{E}^{\chi^*} \left[\sum_{s=0}^{\tau^*-1} \delta^s U_s | \omega^S \right]}{\mathbb{E}^{\chi^*} \left[\sum_{s=0}^{\tau^*-1} \delta^s | \omega^S \right]}, \quad (3)$$

where τ^* is the first time $s \geq 1$ at which \mathcal{I}^S and all the indexes of the alternatives brought to the CS by search fall weakly below the value $\mathcal{I}^S(\omega^S)$ of the search index when search was launched, and where the expectations are under the process induced by χ^* .

(iii) *The DM's expected (per-period) payoff under the policy χ^* is equal to*

$$\int_0^{\infty} \left(1 - \mathbb{E}^{\chi^*} \left[\delta^{\kappa(v)} | \mathcal{S}_0 \right] \right) dv. \quad (4)$$

Proof. See the Appendix. □

As in the classic multi-armed bandit problem with exogenous CS, independence across alternatives is the key assumption behind the optimality of the index policy established in part (i). That is, the payoffs from the various alternatives are drawn independently across the alternatives, given the latter's categories, and the categories of the new alternatives brought to the CS at each expansion only depend on the categories brought to the CS in the past. Under such assumptions, the theorem above establishes a generalization of the Gittins-index Theorem, according to which selecting in each period the alternative, or search, with the highest index is fully optimal for the DM.²⁵ Part (ii) further characterizes

²⁴Note that between the current period and the first period at which all indices are weakly below v , if the DM searches, new alternatives are added to the CS, in which case the evolution of their indices must also be taken into account in the calculation of $\kappa(v)$.

²⁵The reason why indexability of the optimal policy is not obvious is that search is a “meta-arm” bringing alternatives whose returns are correlated at the time search is launched (through the alternatives' categories) and that one needs to process optimally. Problems in which alternatives correspond to “meta arms”, i.e., to sub-problems with their own sub-decisions, typically do not admit an index solution, 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

the optimal stopping time in the index of search. Such recursive representation facilitates an explicit characterization of the index in applications, permits us to identify various properties of the dynamics of experimentation and search, as we show below, and can be used for comparative statics, as illustrated in the applications in the next sections. Finally, part (iii) offers a convenient representation of the DM’s payoff under the optimal rule that can be used, for example, to price access to a search technology with limited knowledge about the details of the environment, as we discuss below.

3.2 Implications for Exploration and Expansion Dynamics

We now highlight several properties of the dynamics of exploration and expansion of the CS, as a function of the search technology.

Corollary 1 (Invariance of expansion to CS composition). *At any period, the decision to expand the CS is invariant to the composition of the CS, conditional on the value $\mathcal{I}^*(\mathcal{S}^P)$ of the alternative with the highest index, and the state ω^S of the search technology.*

The corollary is an immediate implication of the optimal policy being an index policy. However, note that the result is by no means trivial, because the *opportunity cost* of expanding the CS (i.e., the value of continuing with the current CS) typically depends on the entire composition of the CS, beyond the information contained in $\mathcal{I}^*(\mathcal{S}^P)$ and ω^S .

Corollary 2 (Independence of Irrelevant Alternatives). *At any period t , for any pair of alternatives $i, j \in C_t$ with $i \neq j$, the choice between exploring alternative i or exploring j is invariant to the period- t state ω^S of the search technology.*

Corollary 2 is also an immediate implication of Theorem 1. Starting with each period t , the *relative* amount of time the DM spends on each pair of alternatives in the period- t CS is invariant to what the DM expects to find by expanding the CS. This is true despite the fact that further expansions of the CS may bring alternatives that are more similar to one alternative than the other.

Corollary 3 (Possible irrelevance of improvements in search technology). *An improvement in the search technology increasing the probability of finding alternatives of positive expected value (vis-a-vis the outside option) need not affect the decision to expand the CS even at histories at which, prior to the improvement, the DM is indifferent between expanding the CS and exploring one of the alternatives already in it.*

indexability. This is so even if each subset of dependent alternatives evolves independently of all other subsets, and even if one knows how to optimally choose among the dependent alternatives in each subset in isolation. We provide an example illustrating such difficulties in the online Supplementary Material.

The result follows from the fact that improvements in the search technology need not imply an increase in the index of search. This is because, as shown in part (ii) of Theorem 1, the optimal stopping time in the index of search is the first time at which the index of search and the indices of all alternatives brought to the CS by search fall weakly below the value of the index of search at the time search was launched. As a result, any improvement in the search technology affecting only those alternatives whose index at the time of arrival is below the value of the search index at the time search was launched does not affect the value of the search index, and hence the decision to expand the CS.

Definition 2. (i) A search technology is *stationary* if, given any two states of the search technology, $\omega^S = ((c_0, E_0), (c_1, E_1), \dots, (c_m, E_m))$ and $\hat{\omega}^S = ((c_0, E_0), (c_1, E_1), \dots, (c_k, E_k))$, the marginal distribution over (c_{m+1}, E_{m+1}) under H_{ω^S} is the same as the marginal distribution over (c_{k+1}, E_{k+1}) under $H_{\hat{\omega}^S}$. (ii) A search technology is *deteriorating* if, given any $\omega^S = ((c_0, E_0), (c_1, E_1), \dots, (c_m, E_m))$ and $\hat{\omega}^S = ((c_0, E_0), (c_1, E_1), \dots, (c_k, E_k))$, with $\hat{\omega}^S = ((c_0, E_0), (c_1, E_1), \dots, (c_m, E_m), \dots, (c_{m+s}, E_{m+s}))$ for some $s \in \mathbb{N}$, the marginal distribution over $(-c_{m+s+1}, E_{m+s+1})$ under $H_{\hat{\omega}^S}$ is first-order stochastically dominated by the marginal distribution over $(-c_{m+1}, E_{m+1})$ under H_{ω^S} . (iii) A search technology is *improving* if, for each pair of states ω^S and $\hat{\omega}^S$ as in part (ii), the marginal distribution over $(-c_{m+s+1}, E_{m+s+1})$ under $H_{\hat{\omega}^S}$ first-order stochastically dominates the marginal distribution over $(-c_{m+1}, E_{m+1})$ under H_{ω^S} .²⁶

Corollary 4 (Stationary value function). *If the search technology is stationary, for any two states $\mathcal{S}, \mathcal{S}'$ at which the DM expands the CS, $\mathcal{V}^*(\mathcal{S}) = \mathcal{V}^*(\mathcal{S}')$.*

The corollary says that the continuation value when search is launched is invariant to the state of the CS. The result follows from the fact that, without loss of optimality, the DM never comes back to any alternative in the CS after search is launched. The same property holds in case of improving search technologies, as reported in the next corollary.

Corollary 5 (Stationary replacement). *If the search technology is stationary or improving and search is carried out at period t , without loss of optimality, the DM never comes back to any alternative in the CS at period t .*

Since the state of an alternative changes only when the DM selects it, if, in period t , $\mathcal{I}^S \geq \mathcal{I}^*(\mathcal{S}^P)$, under a stationary or improving search technology, the same inequality

²⁶That is, the search technology is deteriorating if, regardless of the outcome of past searches, for any k and any upper set $Z \subset \mathbb{R} \times \mathbb{N}^{|\Xi|}$ (that is, any set $Z \subset \mathbb{R} \times \mathbb{N}^{|\Xi|}$ such that for each $z_1, z_2 \in \mathbb{R} \times \mathbb{N}^{|\Xi|}$ with $z_2 \geq z_1$, $z_2 \in Z$ if $z_1 \in Z$), one has that $\Pr((-c_{k+1}, E_{k+1}) \in Z) \leq \Pr((-c_k, E_k) \in Z)$. This definition is quite strong. In more specific environments, where there is an order on the set of categories Ξ , weaker definitions are consistent with the results in the corollaries below.

remains true in all subsequent periods. In this case, search corresponds to disposal of all alternatives in the current CS. Each time the DM searches, she starts fresh.

Corollary 6 (Single search ahead). *If the search technology is stationary or deteriorating, at any history, the decision to expand the CS is the same as in a fictitious environment in which the DM expects she will have only one further opportunity to search.*

The result follows again from the recursive characterization of the optimal stopping time for the index of search, as per part (ii) of Theorem 1. This time coincides with the first time at which the index of any physical alternative brought to the CS by the current or future searches, and the index of search itself, drop below the value of the search index at the time the current search was launched. If the search technology is stationary, or deteriorating, the index of search falls (weakly) below its current value immediately after search is launched. Hence, $\mathcal{I}^S(\omega^S)$ is invariant to the outcome of any search following the current one, conditional on ω^S .

Corollary 7 (Pricing formula). *Suppose the DM does not have the option to search (i.e., \mathcal{S}_0 is such that $\mathcal{S}_0(\omega^S) = 0$ for all $\omega^S \in \Omega^S$). Let $\hat{\mathcal{S}}_0$ denote a state of the decision problem that coincides with \mathcal{S}_0 except for the fact that $\hat{\mathcal{S}}_0(\hat{\omega}^S) = 1$ for one, and only one, $\hat{\omega}^S \in \Omega^S$. The DM's willingness-to-pay to have access to a search technology in state $\hat{\omega}^S$ is equal to*

$$\mathcal{P}^*(\mathcal{S}_0; \hat{\omega}^S) = \int_0^\infty \left(\mathbb{E} \left[\delta^{\kappa(v)} | \mathcal{S}_0 \right] - \mathbb{E} \left[\delta^{\kappa(v)} | \hat{\mathcal{S}}_0 \right] \right) dv.$$

The result in Corollary 7, which follows directly from part (iii) in Theorem 1, can be used to price access to a search technology with limited knowledge about the details of the environment. To see this, suppose that the econometrician, the analyst, or a search engine, have enough data about the average time it takes for an agent with an exogenous outside option equal to $v \in \mathbb{R}_+$ to exit and take the outside option, both when search is available and when it is not. Then by integrating over the relevant values of the outside option one can compute $\mathcal{P}^*(\mathcal{S}_0; \hat{\omega}^S)$ and hence the maximal price that the DM is willing to pay to access the search technology.

4 Clinical Trials

In each period, a physician must choose which medical treatment to administer among those in her CS, or whether to expand the latter. Whenever the physician administers a treatment, she observes the outcome on the patient that receives it. The outcome yields a payoff to the physician (which may be linked to the well-being of the patients receiving

the treatment) and is informative about the treatment's efficacy (which may be valuable for the good of future patients).

Suppose there are two possible categories of treatments, indexed by $\xi \in \Xi = \{\alpha, \beta\}$. Ex-ante, treatments from the same category are identical. The efficacy $\mu^\xi \in \{0, 1\}$ of each ξ -treatment is unknown ex-ante with $p^\xi(\emptyset) = \Pr(\mu^\xi = 1)$ denoting the ex-ante probability that a ξ -treatment is effective, and with each μ^ξ drawn independently across treatments of the same category. Using a treatment generates information about the treatment's efficacy. Specifically, each time a treatment is used, an outcome $\vartheta \in \{G, B\}$ is observed, with $\vartheta = G$ denoting a "good" outcome, and $\vartheta = B$ a "bad" outcome. If the treatment is effective (i.e., if $\mu^\xi = 1$), the outcome is good with probability $q^\xi = \Pr(\vartheta = G | \mu^\xi = 1) \in (0, 1]$. If the treatment is ineffective (i.e., if $\mu^\xi = 0$), the outcome is bad with certainty.²⁷

The physician's flow payoff from administering a ξ -treatment is v^ξ if the outcome is good and 0 otherwise. Given a history $\theta = (\vartheta_1, \dots)$ of a treatment's outcomes, denote by $p^\xi(\theta)$ the posterior probability that the specific ξ -treatment is effective.

Each search costs the physician $c \geq 0$ and results in the discovery of a new ξ -treatment with probability ρ^ξ , with $\rho^\alpha + \rho^\beta = 1$.²⁸ Each treatment's state $\omega^P = (\xi, \theta)$ consists of the treatment's category ξ along with its history θ of previous outcomes. Letting $\Lambda^\xi(\theta) \equiv p^\xi(\theta)q^\xi$ denote the posterior probability that the treatment yields a good outcome, we have, for any $\omega^P = (\xi, \theta)$,

$$\mathcal{I}^P(\omega^P) = \frac{(1 - \delta + \delta q^\xi) \Lambda^\xi(\theta) v^\xi}{1 - \delta + \delta \Lambda^\xi(\theta)}. \quad (5)$$

Recall that the optimal stopping time in the index definition in (1) is the first time at which the index drops below its initial value (i.e., its value at the time the index is calculated). In this application, this event occurs the first time at which the posterior belief that the treatment is effective drops below its value $p^\xi(\theta)$ at the time the index is computed. The formula in (5) then uses the fact that a good outcome perfectly reveals that the treatment is effective, in which case $\tau^* = \infty$ in (1), whereas a single bad outcome suffices to reduce the posterior belief that the treatment is effective below the value at the time the index was computed, implying that $\tau^* = 1$.

Next, consider the index of search. Using the recursive representation in part (ii) of

²⁷All the results below extend to the case where ineffective treatments also yield good outcomes with positive probability, but smaller than effective ones.

²⁸In particular, the search technology is stationary. After each history $\omega^s = ((c, (E_0(\xi))_{\xi \in \Xi}), \dots, (c, (E_m(\xi))_{\xi \in \Xi}))$, H_{ω^s} assigns probability ρ^α to $((c, (E_0(\xi))_{\xi \in \Xi}), \dots, (c, (E_m(\xi))_{\xi \in \Xi}), (c, (1, 0)))$ and probability $1 - \rho^\alpha = \rho^\beta$ to $((c, (E_0(\xi))_{\xi \in \Xi}), \dots, (c, (E_m(\xi))_{\xi \in \Xi}), (c, (0, 1)))$. That each search brings one and only one new treatment is not essential for the results below.

Theorem 1 together with Corollary 6, we have that the index for search is invariant to ω^S and equal to²⁹

$$\mathcal{I}^S = \frac{(1 - \delta) \left\{ \sum_{\xi \in \{\alpha, \beta\}} \rho^\xi \mathbb{E} \left[\sum_{s=0}^{\tau^{\xi^*} - 1} \delta^s u_s \mid (\xi, \emptyset) \right] - c \right\}}{1 - \sum_{\xi \in \{\alpha, \beta\}} \rho^\xi \mathbb{E} \left[\delta^{\tau^{\xi^*}} \mid (\xi, \emptyset) \right]}, \quad (6)$$

where τ^{ξ^*} is the first time at which the value of the index of the new ξ -treatment brought to the CS by search drops weakly below \mathcal{I}^S ($\tau^{\xi^*} = \infty$ if this event never occurs), and where u_s denotes the flow payoff from the s -th administration of the treatment. To see that the index for search has the above structure recall that the search technology is stationary in this problem. Corollary 6 then implies that the index is the same as in a fictitious environment with a single opportunity to search. Part (ii) of Theorem 1 in turn implies that the optimal stopping time in (2) is the first time at which the posterior belief about the newly added treatment's efficacy is such that the treatment's index drops below the index of search when the latter was launched.

4.1 Detrimental effects of a category's improvement on its usage

The endogeneity of the CS may lead to experimentation dynamics fundamentally different than those that arise when the CS is exogenously fixed ex-ante. To illustrate this possibility, suppose the desirability of one category of alternatives improves. If the CS is fixed, such an improvement typically increases the relative usage of the improved category under the optimal policy. When, instead, the CS is endogenous, this need not be the case: A category's improvement may lead to a reduction in its usage (not only relative to other categories, but overall). While these effects hold quite generally, for concreteness, we illustrate them here in the context of the clinical-trials application introduced above.

To keep things simple, suppose that initially there are two treatments in the physician's CS, one of each category. Because the posterior belief $p^\xi(\theta)$ that a ξ -treatment is effective is equal to one if $\theta \neq \emptyset$ contains at least one good outcome and else depends on $\theta \neq \emptyset$ only through the number s of bad outcomes recorded in θ , with an abuse of notation, hereafter we simplify the formulas for $p^\xi(\theta)$, $\Lambda^\xi(\theta)$, and $\mathcal{I}^P(\xi, \theta)$, by replacing any vector $\theta = (B, B, \dots, B)$ containing only bad outcomes with the number s of bad outcomes in the vector.³⁰ We continue to denote by $p^\xi(\emptyset)$ and $\mathcal{I}^P(\xi, \emptyset)$ the prior belief a ξ -treatment is effective and the index of a ξ -treatment that has never been tested, respectively, and

²⁹The expectations in the formula in (6) are under a rule selecting in each period the ξ -treatment brought to the CS by search.

³⁰Clearly, if θ contains one or more good outcomes, then $p^\xi(\theta) = 1$, in which case $\mathcal{I}^P(\omega^P) = [(1 - \delta + \delta q^\xi) q^\xi v^\xi] / [1 - \delta + \delta q^\xi] = q^\xi v^\xi$.

then let $\Lambda^\xi(\emptyset) \equiv p^\xi(\emptyset)q^\xi$.

Suppose that the following order applies

$$\begin{aligned}
\mathcal{I}^P(\alpha, \emptyset) &> \mathcal{I}^P(\beta, \emptyset) > \mathcal{I}^P(\alpha, 1) \\
&> \frac{-c(1-\delta) + \delta \{ \rho^\alpha \Lambda^\alpha(\emptyset) [1 - \delta^2(1 - q^\alpha)^2] v^\alpha + \rho^\beta \Lambda^\beta(\emptyset) [1 - \delta(1 - q^\beta)] v^\beta \}}{1 - \delta^2 \{ \rho^\alpha \delta [1 - 2\Lambda^\alpha(\emptyset) + q^\alpha \Lambda^\alpha(\emptyset)] + \rho^\beta [1 - \Lambda^\beta(\emptyset)] \}} \\
&> \max\{\mathcal{I}^P(\beta, 1), \mathcal{I}^P(\alpha, 2)\}.
\end{aligned} \tag{7}$$

We then argue that the formula for the index of search in (6) simplifies to

$$\mathcal{I}^S = \frac{-c(1-\delta) + \delta \{ \rho^\alpha \Lambda^\alpha(\emptyset) [1 - \delta^2(1 - q^\alpha)^2] v^\alpha + \rho^\beta \Lambda^\beta(\emptyset) [1 - \delta(1 - q^\beta)] v^\beta \}}{1 - \delta^2 \{ \rho^\alpha \delta [1 - 2\Lambda^\alpha(\emptyset) + q^\alpha \Lambda^\alpha(\emptyset)] + \rho^\beta [1 - \Lambda^\beta(\emptyset)] \}}.$$

Once again, the result follows from part (ii) of Theorem 1, along with Corollary 5. In particular, the order in (7) implies that the optimal stopping time in (6) is equal to: (i) $\tau^{\xi^*} = \infty$ if either the new treatment is an α -treatment and a good outcome is observed in one of the treatment's first two administrations, or the new treatment is a β -treatment and a good outcome is observed after the treatment's first administration; (ii) $\tau^{\xi^*} = 2$ if the new treatment is a β -treatment and the outcome of its first administration is bad; (iii) $\tau^{\xi^*} = 3$ if the newly identified treatment is an α -treatment and each of its first two administrations yielded a bad outcome.³¹

Now suppose that the α -treatments improve. Such an improvement may take the form of (1) an increase in the ex-ante probability $p^\alpha(\emptyset)$ that each α -treatment is effective, (2) an increase in the payoff v^α the physician derives from a good outcome delivered through an α -treatment, or (3) an increase in the informativeness of the outcomes of the α -treatments. Formally, let $\hat{p}^\alpha(\emptyset) = p^\alpha(\emptyset) + \varepsilon_p$, $\hat{v}^\alpha = v^\alpha + \varepsilon_v$, and $\hat{q}^\alpha = q^\alpha + \varepsilon_q$, and then let $\hat{\Lambda}^\alpha(\emptyset) \equiv \hat{p}^\alpha(\emptyset)\hat{q}^\alpha$. Next, suppose that

$$\begin{aligned}
\hat{\mathcal{I}}^P(\alpha, \emptyset) &> \mathcal{I}^P(\beta, \emptyset) > \frac{-c(1-\delta) + \delta \{ \rho^\alpha \hat{\Lambda}^\alpha(\emptyset) [1 - \delta(1 - q^\alpha)] \hat{v}^\alpha + \rho^\beta \Lambda^\beta(\emptyset) [1 - \delta(1 - q^\beta)] v^\beta \}}{1 - \delta^2 [1 - \rho^\alpha \hat{\Lambda}^\alpha(\emptyset) - \rho^\beta \Lambda^\beta(\emptyset)]} \\
&> \max\{\hat{\mathcal{I}}^P(\alpha, 1), \mathcal{I}^P(\beta, 1)\},
\end{aligned} \tag{8}$$

where the hat on the indices $\hat{\mathcal{I}}$ indicates the indices are computed after the improvement in the α -treatments. We then argue that the index for search after the improvement is

³¹Recall that search itself occupies one period.

equal to

$$\hat{\tau}^S = \frac{-c(1-\delta) + \delta \left\{ \rho^\alpha \hat{\Lambda}^\alpha(\emptyset) [1 - \delta(1 - q^\alpha)] \hat{v}^\alpha + \rho^\beta \Lambda^\beta(\emptyset) [1 - \delta(1 - q^\beta)] v^\beta \right\}}{1 - \delta^2 [1 - \rho^\alpha \hat{\Lambda}^\alpha(\emptyset) - \rho^\beta \Lambda^\beta(\emptyset)]}.$$

The result follows again from part (ii) of Theorem 1 along with Corollary 5. Given (8), the optimal stopping time in the definition of the index of search is $\tau^{\xi*} = 2$ if the new treatment brought to the CS by search yields a bad outcome after its first administration, and $\tau^{\xi*} = \infty$ otherwise.

Next, compare the ex-ante expected discounted number of times the physician administers an α -treatment before the improvement and after. The ordering in (7) implies that, before the improvement, the physician starts by administering the α -treatment in the CS. If such a treatment yields a bad outcome, she then administers the β -treatment in the CS. If the latter also yields a bad outcome, the physician administers again the α -treatment in the CS that yielded the initial bad outcome. If this latter treatment yields a second bad outcome, the physician then searches for new treatments. If, at any point, the administered treatment yields a good outcome, because the treatment is revealed effective, the physician then administers it in all subsequent periods, thus bringing the experimentation de facto to a halt. Because the search technology is stationary, by virtue of Corollary 5, all treatments in the CS are effectively discarded once each search for new treatments is carried out. Therefore, the expected discounted number of times the physician administers an α treatment after each search is carried out is given by

$$A_S = \rho^\alpha \left[1 + \delta + \frac{\Lambda^\alpha(\emptyset)(2 - q^\alpha)\delta^2}{1 - \delta} + (1 - \Lambda^\alpha(\emptyset)(2 - q^\alpha))\delta^3 A_S \right] + \rho^\beta (1 - \Lambda^\beta(\emptyset))\delta^2 A_S.$$

Solving for A_S , we have that

$$A_S = \frac{\frac{\rho^\alpha}{1-\delta} (1 - \delta^2 + \delta^2 \Lambda^\alpha(\emptyset) (2 - q^\alpha))}{1 - \delta^2 (\rho^\alpha (1 - \Lambda^\alpha(\emptyset) (2 - q^\alpha)) \delta + \rho^\beta (1 - \Lambda^\beta(\emptyset)))}.$$

From an ex-ante standpoint, the overall expected discounted number of times an α treatment is administered is therefore equal to

$$A = 1 + \frac{\delta \Lambda^\alpha(\emptyset)}{1 - \delta} + (1 - \Lambda^\alpha(\emptyset))(1 - \Lambda^\beta(\emptyset))\delta^2 \left[1 + \frac{\delta \Lambda^\alpha(1)}{1 - \delta} + (1 - \Lambda^\alpha(1))\delta^2 A_S \right], \quad (9)$$

where $\Lambda^\xi(1) = p^\xi(1)q^\xi = (1 - q^\xi)\Lambda^\xi(\emptyset)/(1 - \Lambda^\xi(\emptyset))$ is the probability of a good outcome from a ξ -treatment that yielded a bad outcome at its first administration.

Now let \hat{A}_S and \hat{A} be the analogs of A_S and A , respectively, after the improvement in the α -treatments. Under the order in (8), the physician first administers the α -treatment

in the CS. If the latter yields a bad outcome, the physician then administers the β -treatment in the CS. If the latter also yields a bad outcome, the physician then searches for new treatments.³² Then

$$\hat{A}_S = \rho^\alpha \left(1 + \frac{\delta \hat{\Lambda}^\alpha(\emptyset)}{1 - \delta} + \delta^2 \hat{A}_S (1 - \hat{\Lambda}^\alpha(\emptyset)) \right) + \rho^\beta \left(1 - \Lambda^\beta(\emptyset) \right) \delta^2 \hat{A}_S.$$

Solving for \hat{A}_S , we have that

$$\hat{A}_S = \frac{\rho^\alpha \left(1 + \frac{\delta \hat{\Lambda}^\alpha(\emptyset)}{1 - \delta} \right)}{1 - \delta^2 + \delta^2 \left(\rho^\alpha \hat{\Lambda}^\alpha(\emptyset) + \rho^\beta \Lambda^\beta(\emptyset) \right)}.$$

Therefore, the ex-ante expected discounted number of times an α -treatment is administered when the α -treatments are improved is equal to

$$\hat{A} = 1 + \frac{\delta \hat{\Lambda}^\alpha(\emptyset)}{1 - \delta} + (1 - \hat{\Lambda}^\alpha(\emptyset))(1 - \Lambda^\beta(\emptyset))\delta^3 \hat{A}_S. \quad (10)$$

It can be verified that Conditions (7) and (8) are consistent with

$$A > \hat{A} \quad (11)$$

over an open set of parameter values such that $\varepsilon_p \geq 0$, $\varepsilon_q \geq 0$, and $\varepsilon_v \geq 0$, with at least one inequality strict. We thus have the following result:

Proposition 1. *Suppose Conditions (7), (8), and (11) hold. An improvement in the α -treatments leads to an ex-ante reduction in the expected discounted number of times such treatments are administered. This can happen irrespective of whether the improvement originates in an increase in the ex-ante probability such treatments are effective ($\varepsilon_p > 0$), in the informativeness of such treatments ($\varepsilon_q > 0$), or in the value the physician derives from a good outcome delivered by such treatments ($\varepsilon_v > 0$).*

Proof. The proof follows from the arguments preceding the proposition. \square

As anticipated above, the result is due to the endogeneity of the CS. If the latter were exogenous, an improvement in the α -treatments would never lead to an increase in the administration of such treatments. When the CS is endogenous, however, an improvement in the α -treatments also increases the value of expanding the CS. Because the increase in

³²Again, if at any point a treatment yields a good outcome, because it is revealed effective, it is then administered in each subsequent period. Furthermore, because the search technology is stationary, all treatments in the current CS are effectively discarded when search is carried out.

the index of the α -treatments is not homogeneous across histories, there can be histories at which the increase in the index of search is larger than the increase in the index of the α -treatments. At such histories, the DM may choose to expand the CS instead of further exploring one of the α -treatments, despite the increase in the attractiveness of the latter. Once the expansion is carried out, it may favor the β -treatments more than the α -ones. For example, when the initial CS contains predominately α -treatments and $\rho^\beta \cong \rho^\alpha$, then search is likely to re-balance the composition of the CS in favor of the β -treatments. A similar re-balancing may occur when the initial CS contains an equal number of α - and β -treatments but $\rho^\beta \gg \rho^\alpha$. In either case, the re-balancing can reduce the overall usage of the α -treatments. What is interesting about the result in the proposition is that such re-balancing may be detrimental to the usage of the α -treatment despite the fact that the only reason why search has become more attractive is that it brings α -treatments whose attractiveness have improved. Importantly, the crowding out effect that search may exert on the α -treatments can lead to a reduction in the usage of the α -treatments, not only relative to the β -treatments, but in absolute terms.

5 Experimentation Toward Regulatory Approval

In many problems of interest, a player produces research output to persuade another player to take a decision. For example, a pharmaceutical company may experiment with different drugs or vaccines with the intent of inducing a regulatory authority (e.g., the FDA) to approve its products. Our results can be used to shed light on such problems.

Consider a firm that needs regulatory approval to sell its products. The products differ in their profitability to the firm, but also in their safety. We capture such heterogeneity by assuming that each product belongs in a category $\xi \in \Xi$, where Ξ is a finite set. Each ξ -product can either be safe ($\mu^\xi = 1$) or not ($\mu^\xi = 0$), and this event is unknown both to the firm and to the regulator at the outset. The products' safety is independent across products, conditional on their category. Let $p^\xi(\emptyset) = \Pr(\mu^\xi = 1)$ denote the prior probability that a category- ξ product is safe. Each ξ -product, when sold to the market, brings the firm a flow payoff equal to $(1 - \delta)v^\xi > 0$. There is no value to the firm in selling more than one product per period (for example, because the products are seen as substitutes by the consumers).

At each period, the firm can either experiment with one of the products in its current CS, expand the CS by searching for new products, or sell one of the approved products. Each experiment on a ξ -product generates a binary outcome $\vartheta \in \{G, B\}$, with $\vartheta = G$ denoting a “good” outcome and $\vartheta = B$ a “bad” one. If the product is safe, a good outcome

$\vartheta = G$ is generated with probability $q_1^\xi = \Pr(\vartheta = G | \mu^\xi = 1) \in (0, 1]$. If, instead, the product is unsafe, a bad outcome $\vartheta = B$ is generated with probability $q_0^\xi = \Pr(\vartheta = B | \mu^\xi = 0)$, with $q_1^\xi \geq 1 - q_0^\xi$. Given a history $\theta = (\vartheta_1, \dots)$ of experimentation outcomes, denote by $p^\xi(\theta)$ the posterior probability that a ξ -product is safe. As in [Henry and Ottaviani \(2019\)](#), the history θ is public. Experimentation with a ξ -product entails a cost $\lambda^\xi(\theta)$ to the firm, the magnitude of which may depend on the history of past outcomes. Expansion of the CS entails a constant cost $c \geq 0$ and brings one new product whose category is drawn from Ξ according to a distribution $\rho \in \Delta(\Xi)$, with ρ^ξ denoting the probability that the newly discovered product is of category ξ , with $\sum_{\xi \in \Xi} \rho^\xi = 1$.³³

We model the approval process in reduced form as follows. For any category ξ , there exists a threshold $\Psi^\xi \in (0, 1]$ such that a ξ -product is approved once the posterior probability $p^\xi(\theta)$ that the product is safe exceeds the threshold Ψ^ξ . To avoid trivialities, assume that $p^\xi(\emptyset) < \Psi^\xi$, so that each product must be tested at least once to be approved.

The firm's goal is to maximize the expected discounted payoff from selling its (approved) products, net of all experimentation and search costs. The solution to this problem is also given by the index policy of [Theorem 1](#). In each period prior to receiving approval for one of its products, the firm experiments with one of the products in its current CS with the highest index if such an index is higher than the index for search, and expands its CS otherwise.

Given $\omega^P = (\xi, \theta)$, the index of a product in state ω^P that has not been approved yet is equal to

$$\mathcal{I}^P(\omega^P) = \frac{(1 - \delta) \mathbb{E} \left[- \sum_{s=0}^{\min\{\tau^*, \phi\}} \delta^s \lambda^\xi(\theta) + \delta^\phi \mathbf{1}_{\{\phi < \tau^*\}} v^\xi | \omega^P \right]}{1 - \mathbb{E}[\delta^{\tau^*} | \omega^P]}, \quad (12)$$

where the expectations in [\(12\)](#) are under the process obtained by selecting the product under consideration in each period, ϕ is the first time at which $p^\xi(\theta)$ exceeds the threshold Ψ^ξ ($\phi = \infty$ if this event never occurs), and τ^* is either the first period at which the posterior belief that the product is safe is below the value $p^\xi(\theta)$ at the time the index is computed, when such an event occurs before the product is approved, i.e., before period ϕ , or is equal to $\tau^* = \infty$ otherwise. The index for search is given by

$$\mathcal{I}^S = \frac{(1 - \delta) \left(-c + \delta \sum_{\xi \in \Xi} \rho^\xi \mathbb{E} \chi^* \left[- \sum_{s=0}^{\min\{\tau^{\xi*}, \phi^\xi\}} \delta^s \lambda^\xi(\theta) + \delta^{\phi^\xi} \mathbf{1}_{\{\phi^\xi < \tau^{\xi*}\}} v^\xi | \xi, \emptyset \right] \right)}{1 - \sum_{\xi \in \Xi} \rho^\xi \mathbb{E} \chi^* [\delta^{\tau^{\xi*}} | \xi, \emptyset]}, \quad (13)$$

³³That $\sum_{\xi \in \Xi} \rho^\xi = 1$ is without loss of generality. The case where search brings no product with positive probability can always be captured by letting one of the categories replicate the continuation payoff associated with the arrival of no new product.

where ϕ^ξ is the first time at which $p^\xi(\theta)$ exceeds the approval threshold, whereas $\tau^{\xi*}$ is either the first time at which the value of the index of the new ξ -product brought to the CS drops weakly below \mathcal{I}^S , when this event occurs before ϕ^ξ , or is equal to $\tau^{\xi*} = \infty$ otherwise.

Finally, note that, because experimenting with a product that has been approved already is dominated by selling the approved product, the index of a ξ -product that received regulatory approval is constant and equal to $(1 - \delta)v^\xi$. The optimal policy being an index policy then also implies that, as soon as one of the firm's products is approved, the firm brings to an end its experimentation process and sells the approved product in each of the subsequent periods.³⁴

5.1 Changes in approval standards

Suppose now that the regulator adopts a policy specifying a lower approval threshold for one of the products' categories, possibly in response to political pressure to approve more products from that category. What are the effects of such a change? For simplicity, suppose that, as in the clinical-trials application, there are two product categories, α and β , and two products in the firm's initial CS, one of each category. Further assume that $q_1^\xi = q_0^\xi = q^\xi$ meaning that good and bad signals are equally informative, $\xi = \alpha, \beta$, so that the posterior belief $p^\xi(\theta)$ that a ξ -product is safe depends on the history θ only through the difference between good and bad outcomes. We then have the following result:

Proposition 2. *A relaxation of the approval standards for the category- α products from Ψ^α to $\Psi^\alpha - \varepsilon$, $\varepsilon > 0$, may trigger a reduction in the ex-ante probability that an α -product is approved.*

Proof. See the online Supplementary Material.³⁵ □

The result hinges on the endogeneity of the CS. If the latter was exogenous, the relaxation would unambiguously favor the α -category. When, instead, the CS is endogenous, the relaxation also increases the index of search (provided of course that search brings with positive probability the α -products, i.e., that $\rho^\alpha > 0$). The intuition is similar to the one discussed for clinical trials. The index for search may increase more than the index of the α -products after the evaluation of the latter products yielded a few negative results.

³⁴This follows from the fact that the index of a ξ -product that has been approved already is higher than the index of any ξ -product that has not been approved yet.

³⁵The arguments are related to those establishing Proposition 1 above. Because the derivations are more lengthy, though, the proof is relegated to the online Supplementary Material.

When this is the case, search may crowd out further evaluations of such products, thus contributing to a reduction in the ex-ante probability that α -products are approved.

The result cautions against simplistic policy recommendations when firms' product sets are endogenous (for example because of R&D). The implications of the above result for the selection of the optimal regulatory standards is however left for future work.

6 Pandora's problem with an endogenous set of boxes

Consider the following variant of [Weitzman's \(1979\)](#) "Pandora's problem" in which the set of boxes is endogenously expanded over time. 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 value v is drawn and λ^ξ is the cost of inspecting (i.e., of opening) the box. As in [Weitzman \(1979\)](#)'s original setting, each box's value v is revealed upon its 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 value, or (c) stop and either recall the value of one of the previously opened boxes, or take the outside option, with either one of the last two choices ending the decision problem. For simplicity, assume that each search $m \in \mathbb{N}$ brings exactly one box, whose category ξ is drawn from Ξ according to a time-homogeneous distribution $\rho \in \Delta(\Xi)$ independently across searches, with ρ^ξ denoting the probability that search brings a ξ -box, and $\sum_{\xi \in \Xi} \rho^\xi = 1$.³⁶ The boxes' values are drawn independently across boxes, conditional on their categories. The cost of expanding the set of boxes depends on the number of past searches, with $c(m)$ denoting the cost of the m -th search, where $c(\cdot)$ is a positive and increasing function. The DM discounts the future according to δ .

The solution to this problem is also an index policy with the structure of [Theorem 1](#). To see this, it suffices to follow the familiar trick of considering a relaxed problem in which each open box can continue to be selected, with each selection reversible (i.e., the DM can switch to other boxes) and yielding a flow payoff equal to $(1 - \delta)v$, where v is the box's value. Because the index of each open box whose value is v is constant and equal to $(1 - \delta)v$ after each selection, once, in the relaxed problem, the DM selects an open box, she continues selecting the same box forever. The optimal index policy for the relaxed problem is thus feasible (and hence optimal) also in the primitive problem.

³⁶For simplicity, we assume here that the set Ξ is finite. All the results extend to the case where Ξ is infinite. Likewise, that the distribution from which the box categories are drawn is invariant to the number of past searches is not essential. The results below extend to the case where such a distribution depends on m provided that the indexes for search decline (weakly) with the number of past searches.

For any $\omega^P = (\xi, \emptyset)$, the index of a ξ -box that has not been opened yet is equal to³⁷

$$\mathcal{I}^P(\omega^P) = \frac{-\lambda^\xi + \delta \int_{\frac{\mathcal{I}^P(\omega^P)}{1-\delta}}^{\infty} v dF^\xi(v)}{1 + \frac{\delta}{1-\delta} \left(1 - F^\xi\left(\frac{\mathcal{I}^P(\omega^P)}{1-\delta}\right)\right)}. \quad (14)$$

The index $\mathcal{I}^P(\omega^P)$ in (14) corresponds to what Weitzman (1979) refers to as a box’s “reservation price”.³⁸

Next, for any $l \in \mathbb{R}$, let $\Xi(l) \equiv \{\xi \in \Xi : \mathcal{I}^P(\xi, \emptyset) > l\}$ denote the set of boxes whose reservation price exceeds l . Because all the relevant information about the evolution of the state of the search technology is summarized in the number of times search was conducted in the past, hereafter we abuse notation and let $\mathcal{I}^S(m)$ denote the index for the m -th search, with the latter given by³⁹

$$\mathcal{I}^S(m) = \frac{-c(m) + \delta \sum_{\xi \in \Xi(\mathcal{I}^S(m))} \rho^\xi \left(-\lambda^\xi + \delta \int_{\frac{\mathcal{I}^S(m)}{1-\delta}}^{\infty} v dF^\xi(u)\right)}{1 + \sum_{\xi \in \Xi(\mathcal{I}^S(m))} \rho^\xi \left[\delta + \frac{\delta^2}{1-\delta} \left(1 - F^\xi\left(\frac{\mathcal{I}^S(m)}{1-\delta}\right)\right)\right]}. \quad (15)$$

6.1 Consumer search and eventual purchase

As an illustration of the potential applications of the above model, consider a consumer searching online for the best possible alternative. Despite the importance that sponsored search plays in modern business activities, the models that have been used to study online search typically assume that consumers click ads in the order they are displayed, and that click-through-rates are either invariant to positions or depend on the latter in an exogenous fashion. Such assumptions do not appear to square well with empirical findings.⁴⁰

³⁷To see why the index takes this form, note that, because the value of the box v is fully revealed once opened, the index of an open box coincides with its flow value $(1-\delta)v$. Because the optimal stopping time τ^* in the definition of the index $\mathcal{I}^P(\omega^P)$ in (1) is the first time at which the value of the index drops below its value $\mathcal{I}^P(\omega^P)$ at the time the index is computed, we then have that $\tau^* = 1$ if $(1-\delta)v \leq \mathcal{I}^P(\omega^P)$ and $\tau^* = \infty$ otherwise.

³⁸Weitzman (1979) defines the reservation price $\hat{\mathcal{I}}^P(\omega^P)$ for $\omega^P = (\xi, \emptyset)$ as the solution to $\lambda^\xi = \delta \int_{\hat{\mathcal{I}}^P(\omega^P)}^{\infty} (v - \hat{\mathcal{I}}^P(\omega^P)) dF^\xi(v) - (1-\delta)\hat{\mathcal{I}}^P(\omega^P)$, which yields

$\hat{\mathcal{I}}^P(\omega^P) = \left[-\lambda^\xi + \delta \int_{\hat{\mathcal{I}}^P(\omega^P)}^{\infty} v dF^\xi(v)\right] / \left[1 - \delta + \delta \left(1 - F^\xi\left(\hat{\mathcal{I}}^P(\omega^P)\right)\right)\right]$. The indices in (14) are thus equal to the reservation prices in Weitzman (1979) multiplied by $(1-\delta)$, that is, $\mathcal{I}^P(\omega^P) = (1-\delta)\hat{\mathcal{I}}^P(\omega^P)$.

³⁹By Corollary 6, because the search technology is deteriorating, the optimal stopping-time τ^* in (2) is equal to (a) $\tau^* = \infty$ if the box identified at the m -th search has a reservation price $\mathcal{I}^P(\omega^P) > \mathcal{I}^S(m)$ and its realized flow payoff $v(1-\delta) > \mathcal{I}^S(m)$, (b) $\tau^* = 1$ if $\mathcal{I}^P(\omega^P) \leq \mathcal{I}^S(m)$, and (c) $\tau^* = 2$ if $\mathcal{I}^P(\omega^P) > \mathcal{I}^S(m)$ and $v(1-\delta) \leq \mathcal{I}^S(m)$.

⁴⁰For example, using data from Microsoft Live, Jeziorski and Segal (2015) find that almost half of the users who click on a link do not click in sequential order of positions, and that click-through-rates depend non-trivially on the identity of competing ads.

The Pandora’s problem with an endogenous set of boxes described above can be used to endogenize the relationship between the ad positions and the corresponding click-through-rates. To see this, note that a consumer who reads an ad displayed by a search engine brings the advertised product into her CS. By clicking on the product’s link, the consumer is then directed to the vendor’s website where she learns the value for the vendor’s product and then decides whether or not to finalize the purchase (think of such a value as net of the product’s price). The purchase of a product brings to an end the consumer’s search process. Hence, in this environment, reading an ad corresponds to expanding the CS, clicking on an ad corresponds to opening a box, and purchasing a product from one of the visited vendors corresponds to selecting an open box.⁴¹

Formally, suppose that each category $\xi \in \Xi$ corresponds to a different firm and that each position $m \in \mathbb{N}$ is occupied by the ad of one and only one firm, but possibly with the same firm displaying ads at multiple positions. Reading the ad displayed at the m -th position reveals to the consumer the identity $\xi(m) \in \Xi$ of the firm whose ad occupies the m -th position. From the perspective of the consumer, $\xi(m)$ is drawn from Ξ according to some distribution $\rho \in \Delta(\Xi)$ reflecting the consumer’s beliefs over how the search engine assigns positions to different firms, with ρ^ξ denoting the probability that the consumer assigns to position m being occupied by firm ξ . The consumer believes that each $\xi(m)$ is drawn from ρ , independently across m .⁴²

The discovery of the identity of the firm occupying the m -th position in turn informs the consumer that if she were to click on the ad in the m -th position she would be directed to firm $\xi(m)$ ’s website, where she would incur a cost $\lambda^{\xi(m)}$ to learn her value v for firm $\xi(m)$ ’s product, with v drawn from an absolutely continuous distribution $F^{\xi(m)}$.⁴³ The

⁴¹This formulation implicitly assumes that the consumer does not know the content of an ad before reading it. This seems reasonable given that the selection of the ads is made by the search engine and is not perfectly predictable by the consumer. It also assumes that consumers do not click on ads they have not read. This also seems reasonable given that clicking on an ad without reading it amounts to being directed to vendors’ webpages at random, which is not profitable for the consumer. The formulation above also assumes that consumers read the ads in the order they are displayed by the search engine (but may click on the links of the ads they have read in the order they choose). This also seems consistent with normal practices. Finally, the above formulation assumes that, once a purchase is made, search ends. This also seems plausible if the search is for substitutable goods.

⁴²This assumption simplifies the exposition but is not essential. The results below extend to other search technologies that are deteriorating. For example, the consumer may expect lower positions to be occupied by firms providing, on average, lower-value products.

⁴³Heterogeneity across firms in both F^ξ and λ^ξ may reflect the consumer’s prior experiences, the firms’ reputations, as well as the consumer’s expectations about the easiness to navigate through each firm’s website. The assumption that each F^ξ is absolutely continuous is made only to avoid the need to keep track of possible indifferences in the consumer’s optimal behavior which affect the formulas but not the qualitative results.

consumer expects her values for the different products to be drawn independently across positions, given the firms' identities.⁴⁴

Let $c(m)$ denote the cost of reading the ad displayed at the m -th position (because the consumer reads the ads in the order in which they are displayed, $c(m)$ is also the cost of the m -th expansion of the consumer's CS). Finally, let position $m = 0$ correspond to the consumer's outside option, yielding a payoff normalized to zero.

We then have that the index for the decision to read the ad in the m -th position is equal to $\mathcal{I}^S(m)$, with $\mathcal{I}^S(m)$ as defined in (15), whereas the index for the decision to click on the ad displayed in the m -th position is equal to $\mathcal{I}_m \equiv \mathcal{I}^P(\xi(m), \emptyset)$, with $\mathcal{I}^P(\xi(m), \emptyset)$ as defined in (14).

One can then use the model to endogenize the probability with which the consumer reads the ads, clicks on them, and finalizes her purchases. In a similar setting but with an exogenous CS, Choi, Dai and Kim (2018) (and, independently, Armstrong, 2017) derive a static condition characterizing eventual purchase decisions based on a comparison of "effective values." Proposition 3 below extends the characterization of the eventual purchase probabilities to an endogenous CS. Let v_m denote the value to the consumer for the product sold by the firm advertising in the m -th position.

For all $m \geq 1$, let $w_m \equiv \min\{\mathcal{I}_m, v_m(1 - \delta)\}$ denote the "effective value" of the product advertised in position m (for brevity, product m), when the product is already in the consumer's CS, as defined in Choi, Dai and Kim (2018) and then let $d_m \equiv \min\{w_m, \mathcal{I}^S(m)\}$ denote the product "discovery value," when the product must be brought to the consumer's CS before it can be explored. Let product $m = 0$ correspond to the consumer's outside option, with $w_0 = d_0 = 0$.

Proposition 3. *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 $l \neq m$).*

Proof. See the Appendix. □

As in Choi, Dai and Kim (2018), consumers' purchasing decisions are determined by a static comparison of effective values, as in canonical discrete-choice models. Contrary to Choi, Dai and Kim (2018), however, such effective values account for the order by which the various products are brought to the consumer's CS, which here depends on the order in which the products' ads are displayed by the search engine. In particular,

⁴⁴This also means that, in case the consumer encounters ads from the same firm ξ at different positions, with each ad corresponding to a different product from the same firm, the consumer expects her value for each of firm ξ 's products to be drawn from F^ξ independently across products.

Proposition 3 implies that, provided that the reading cost $c(\cdot)$ is non-decreasing, 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 typically assumed, but not micro-founded, in existing search models.

The result in Proposition 3 follows from the combination of the fact that the optimal policy is an index rule with the fact that the search indexes $\mathcal{I}^S(m)$ decline with m . Heuristically, if a consumer reads the ad displayed at the m -th position, because she must have read also all the ads displayed at positions $l < m$ (meaning that all such products have been brought to the consumer's CS), it must be that the reservation prices \mathcal{I}_l of all products $l < m$, as well as the discovered values $v_l(1-\delta)$ of those products $l < m$ that have been inspected, are no greater than $\mathcal{I}^S(m)$. Furthermore, because the search technology is non-improving, $\mathcal{I}^S(m+1) \leq \mathcal{I}^S(m)$. Hence, if after reading the m 's ad, $\mathcal{I}_m \geq \mathcal{I}^S(m)$, the consumer necessarily clicks on m 's ad and then learns m 's value v_m . Once v_m is learned, if $(1-\delta)v_m \geq \mathcal{I}_m$, the consumer then stops the search and purchases product m . The formal proof in the Appendix shows how the above monotonicity properties can be used to arrive at the result in the proposition.

The above results can also be used to endogenize the positions' *click-through rates* (formally, the fraction of ads assigned by a search engine to each position that are eventually clicked upon, among those that are brought to the consumer's CS). In the setting described above, a product is brought to the consumer's CS after its ad has been displayed to and read by the consumer.⁴⁵ We thus have that, for each position m , the click-through-rate is defined as⁴⁶

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

Click-through-rates play an important role for the pricing of ad space on platforms, as studied at length in the IO and marketing literatures. We then have the following result:⁴⁷

⁴⁵It is convenient to think of each product as displayed on a different page, with the platform tracking whether a consumer moves from one page to the next.

⁴⁶Note that the probability in the definition of $CTR(m)$ is computed ex-ante by integrating over the different products that are advertised at the different positions. It is not the probability that a specific product advertised at a given position is clicked upon. In other words, the click-through-rates are position-specific and not product-specific.

⁴⁷For 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 in the online supplement discusses how alternative ways of breaking the indifferences must be accounted for if one were to compute bounds for such probabilities.

Proposition 4. *The CTR for each position $m \geq 1$ is given by*

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

Proof. See the online Supplementary Material. □

For product 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 ever reading the ad displayed in the m 's position. Once product m is read, 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 advertised further down the list, for otherwise the consumer selects one of the other products before clicking on m .⁴⁸

The result in Proposition 4 also suggests that, while the ex-ante demands are naturally decreasing in the positions, the click-through rates need not be monotone in m . To see this, note that $Pr(\mathcal{I}_m \geq \max_{l < m}\{w_l\})$ is decreasing in m , which contributes to CRTs declining in m . However, for product m to be read, it must be that $\max_{l < m}\{w_l\} \leq \mathcal{I}^S(m)$. Because $\mathcal{I}^S(m)$ is decreasing in m , $Pr(\max_{l < m}\{w_l\} \leq \mathcal{I}^S(m))$ is also decreasing in m , thus contributing to the possibility that CTR are non-monotone in positions.⁴⁹

The results above also have implications for the effects of additional ad space on firms' profits. To see this, consider the following situation. The consumer's initial CS contains three products, one from each firm $\xi \in \Xi = \{A, B, C\}$. By searching online, the consumer is presented with a fourth product drawn from Ξ according to $\rho \in \Delta(\Xi)$. As above, each product yields the consumer a net value v drawn from a distribution F^ξ , independently across products.⁵⁰ The cost to the consumer of learning her value for each ξ 's products is λ^ξ .⁵¹ The consumer has unit demand and each firm's profit is the same for each of its

⁴⁸Note that the assumption 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 $j = m + 1, \dots, l$, $\mathcal{I}^S(j) > \mathcal{I}_m$, 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 discover value exceeds \mathcal{I}_m , the consumer will purchase one of these products before clicking on m .

⁴⁹That $Pr(\mathcal{I}_m \geq \max_{l > m}\{d_l\})$ is increasing in m also contributes to the possibility of CRTs increasing in m .

⁵⁰If the extra product the consumer is presented when searching is from firm ξ , the value the consumer derives from such a product is also drawn from F^ξ , independently from the value she derives from the three products already in her CS.

⁵¹Again, if the extra product the consumer is presented when searching is from firm ξ , because the value she derives from such product is independent from the value she derives from the other product from firm ξ already in her CS, the total cost the consumer incurs to learn her value for each of firm ξ 's products is thus $2\lambda^\xi$.

products. We then have the following result:

Proposition 5. *Consider the environment described above. An increase in the probability ρ^ξ that search brings an additional firm- ξ product may reduce firm ξ 's ex-ante expected profits.*

Proof. See the online Supplementary Material. □

Once again, the result hinges on the endogeneity of the consumer's CS. An 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. As we show in the online supplementary material, this may in turn reduce the probability that one of firm ξ 's product is selected and hence its profits.

7 Extensions

We have shown that the endogeneity of the CS brings novel exploration dynamics and comparative statics fundamentally different from those under an exogenous CS. The results accommodate for a few extensions that may be relevant for applications.

No discounting. All proofs assume that $\delta < 1$. The results, however, extend to $\delta = 1$ (i.e., no discounting). As noted in [Olszewski and Weber \(2015\)](#), bandit problems in which $\delta = 1$ can be thought of as problems with undiscounted “target processes” whereby arms reaching a certain (target) state stop delivering payoffs. A well known result for such problems is that the finiteness they impose allows to take the limit as $\delta \rightarrow 1$ (see, e.g., [Dumitriu, Tetali, and Winkler, 2003](#)).

Irreversible Choice. In many decision problems, in addition to learning about existing options and searching for new ones, the DM can irreversibly commit to one of the alternatives, bringing to an end the exploration process. As mentioned above, in general, such problems do not admit an index solution. In the online Supplementary Material, we derive a sufficient condition under which the optimality of an index rule extends to such problems. We assume the DM must explore each alternative of category ξ at least $M_\xi \geq 0$ times before she can irreversibly commit to it (for example, a consumer must visit a vendor's webpage at least once to finalize a transaction with that vendor, as in the consumer search problem of Section 6). The condition guarantees that, once an alternative reaches a state in which the DM can irreversibly commit to it, its “retirement value” (that is, the value of irreversibly committing to it) either drops below the value of the outside option, or improves over time. This property is related to a similar condition

in [Glazebrook \(1979\)](#), who establishes the optimality of an index policy in a class of bandit problems with stoppable processes. Our proof, however, is different and accommodates for the possibility that the set of alternatives evolves endogenously over time.

Relative length of expansion. In order to allow for frictions in the search for new alternatives, we assume that, whenever the DM searches, she cannot explore any of the alternatives in the CS, with search occupying the same amount of time as the exploration of any of the alternatives in the CS. All the results extend to a setting in which both the time that each search occupies and the time that each exploration takes vary stochastically with the state.⁵² Furthermore, because the time that each exploration takes can be arbitrary, by rescaling the payoffs and adjusting the discount factor appropriately, one can make the length of time during which the exploration of the existing alternatives is paused because of search arbitrarily small. The results therefore also apply to problems in which search and learning occur “almost” in parallel.

Multiple expansion possibilities. In certain problems of interest, the decision to search also involves an intensive margin, as when the DM chooses “how much” to invest in search. As we show in the online Supplementary Material, in general, such problems do not admit an index solution because of the correlation in the search outcomes. Instead, the analysis readily extends to an environment in which there are multiple search possibilities with independent outcomes, by allowing for multiple “search arms”.

References

- Armstrong, M.** (2017). “Ordered consumer search,” *Journal of the European Economic Association*, 15(5), 989-1024.
- Armstrong, M. and J. Vickers** (2015). “Which demand systems can be generated by discrete choice?” *Journal of Economic Theory*, 158, 293-307.
- Austen-Smith, D. and C. Martinelli** (2018). “Optimal exploration.” Working paper.
- Bergemann, Dirk and J. Välimäki** (2008). “Bandit problems”. In: *The New Palgrave Dictionary of Economics*. Ed. by Steven N. Durlauf and Lawrence E. Blume. Basingstoke: Palgrave Macmillan.
- Berry, D. A.** (2006). “Bayesian clinical trials,” *Nature Reviews Drug discovery*, 5(1), 27-36.
- Caplin, A., Dean, M. and D. Martin** (2011). “Search and satisficing.” *American Economic Review*, 101(7), 2899-2922.

⁵²More generally, all of the results can be extended to a semi-Markov environment, where time is not slotted.

- Caplin, A., Dean, M. and J. Leahy** (2019). “Rational inattention, optimal consideration sets, and stochastic choice,” *The Review of Economic Studies*, 86(3), 1061-1094.
- Che, Y. K. and K. Mierendorff** (2019). “Optimal dynamic allocation of attention.” *American Economic Review*, 109(8), 2993-3029.
- Choi, M., Dai, A. Y. and K. Kim** (2018). “Consumer search and price competition.” *Econometrica*, 86(4), 1257-1281.
- Choi, M. and L. Smith** (2016). “Optimal sequential search among alternatives.” mimeo University of Wisconsin.
- Dickstein, M. J.** (2014). “Efficient provision of experience goods: Evidence from antidepressant choice.” mimeo Stanford University.
- Doval, L.** (2018). “Whether or not to open Pandora’s box.” *Journal of Economic Theory* 175, 127–158.
- Dumitriu, I., Tetali, P. and P. Winkler** (2003). “On playing golf with two balls,” *SIAM Journal on Discrete Mathematics*, 16(4), 604-615.
- Edelman, B., Ostrovsky, M. and M. Schwarz** (2007). “Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords,” *American Economic Review*, 97(1), 242-259.
- Eliasz, K. and R. Spiegler** (2011). “Consideration sets and competitive marketing,” *The Review of Economic Studies*, 78(1), 235-262.
- Fershtman, D. and A. Pavan** (2020). “Soft Affirmative Action and Minority Recruitment,” mimeo, Northwestern and Tel-Aviv University.
- Garfagnini, U. and B. Strulovici** (2016). “Social experimentation with interdependent and expanding technologies,” *Review of Economic Studies*, 83(4), 1579-1613.
- Gittins, J. C.** (1979). “Bandit processes and dynamic allocation indices,” *Journal of the Royal Statistical Society. Series B (Methodological)* 41 (2), 148–177.
- Gittins, J. and D. Jones** (1974). “A dynamic allocation index for the sequential design of experiments. In J. Gani (Ed.),” *Progress in Statistics*, pp. 241-266. Amsterdam, NL: North- Holland.
- Glazebrook, K. D.** (1979). “Stoppable families of alternative bandit processes.” *Journal of Applied Probability* 16.4: 843-854.
- Gossner, O., Steiner, J. and C. Stewart** (2019). “Attention, please!,” mimeo.
- Greminger, R.** (2020). “Optimal search and discovery,” mimeo, Tilburg University.
- Hauser, J. R. and B. Wernerfelt** (1990). “An evaluation cost model of consideration sets,” *Journal of Consumer Research*, 16(4), 393-408.
- Henry, E., and M. Ottaviani** (2019). “Research and the approval process: the organi-

- zation of persuasion,” *American Economic Review*, 109(3), 911-55.
- Jeziorski, P. and I. Segal** (2015). “What makes them click: Empirical analysis of consumer demand for search advertising,” *American Economic Journal: Microeconomics*, 7(3), 24-53.
- Katehakis, M., and C. Derman** (1986). “Computing optimal sequential allocation rules in clinical trials.” *Lecture Notes-Monograph Series*, 8, 29-39.
- Ke, T. T., Z.-J. M. Shen and J. M. Villas-Boas** (2016). “Search for information on multiple products,” *Management Science* 62 (12), 3576-3603.
- Ke, T. T. and J. M. Villas-Boas** (2019). “Optimal learning before choice,” *Journal of Economic Theory* 180, 383-437.
- Keller, G. and A. Oldale** (2003). “Branching bandits: a sequential search process with correlated pay-offs,” *Journal of Economic Theory*, 113(2), 302-315.
- Liang, A., Mu, X., and V. Syrgkanis** (2019). “Dynamically aggregating diverse information,” mimeo.
- Mandelbaum, A.** (1986). “Discrete multi-armed bandits and multi-parameter processes,” *Probability Theory and Related Fields* 71(1), 129–147.
- Manzini, P. and M. Mariotti** (2014). “Stochastic choice and consideration sets,” *Econometrica*, 82(3), 1153-1176.
- Masatlioglu, Y., Nakajima, D. and E. Y. Ozbay** (2012). “Revealed attention,” *American Economic Review*, 102(5), 2183-2205.
- Olszewski, W. and R. Weber** (2015). “A more general Pandora rule?,” *Journal of Economic Theory* 160, 429–437.
- Roberts, J. H. and J. M. Lattin** (1991). “Development and testing of a model of consideration set composition,” *Journal of Marketing Research*, 28(4), 429-440.
- Schneider, J. and C. Wolf** (2019). “Switching forth and back: Problem solving under time pressure,” mimeo.
- Simon, H. A.** (1955). “A behavioral model of rational choice,” *The Quarterly Journal of Economics*, 69(1), 99-118.
- Villar, S. S., Bowden, J., and J. Wason** (2015). “Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges.” *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, 30(2), 199.
- Weber, R.**, (1992). “On the Gittins Index for Multiarmed Bandits,” *The Annals of Applied Probability*, Vol. 2(4), 1024-1033.
- Weiss, G.** (1988). “Branching bandit processes,” *Probability in the Engineering and Informational Sciences*, 2(3), 269-278.

Weitzman, M. (1979). “Optimal search for the best alternative,” *Econometrica* 47 (3), 641–654.

8 Appendix

Preliminaries: General states and processes. The analysis applies to a broader class of problems than the one described in Section 2, in which the evolution of the different alternatives may originate in shocks other than the arrival of information, and where the endogeneity of the shocks may reflect for example a preference for variety, or habit formation. To facilitate the analysis, it will be useful to introduce a more general framework with more general processes and states (which need not reflect learning).

Suppose that exploring alternative $i \in C_t$ at period t generates a stochastic flow payoff $u_{it} \in \mathbb{R}$, the distribution of which is a function of the alternative’s state. Each alternative has a *category* ξ , an element of an arbitrary topological space Ξ , which determines the stochastic process governing the evolution of its state. The process corresponding to each alternative of category ξ is Markov and time-homogeneous (i.e., invariant to calendar time). Slightly abusing notation, denote by $\omega^P = (\xi, \theta) \in \Omega^P = \Xi \times \Theta$ the alternative’s current *state*, where θ is an element of an arbitrary set Θ . While ξ is fixed and captures the alternative’s “category,” or process, θ evolves over time. Denote by σ the sigma-field associated with Ω^P , and by $H_{\omega^P} \in \Delta(\Omega^P)$ the distribution over next period’s state when the alternative’s current state is ω^P . Without loss of generality, we assume that the flow payoffs the DM receives over time are governed by the same process describing the evolution of the alternative’s state (that is, they can be represented by a deterministic function of the alternative’s state). The first time the DM explores an alternative of category ξ , the latter is in state (ξ, \emptyset) .

The process governing the search technology – i.e., the cost incurred as well as the number and categories of new alternatives introduced through search – is also Markov time-homogeneous, and is the same as the one described in Section 2. The state of the decision problem is also defined as in Section 2. A policy χ for the above decision problem is a rule governing the decisions in each period – whether to explore an alternative in the CS or search for new ones – based on the available information. Given a sequence of feasible decisions $(d_t)_{t \geq 0}$, the state process $(\mathcal{S}_t)_{t \geq 0}$ generates a natural filtration $\{\mathcal{F}_t\}_{t \geq 0}$. A policy χ is then a \mathcal{F}_t -measurable sequence of feasible decisions. Denote by $U_t \equiv \sum_{j=0}^{\infty} x_{jt} u_{jt} - c_t y_t$ the realized period- t net payoff. A policy χ is *optimal* if it maximizes the expected discounted sum of the net payoffs $\mathbb{E}^\chi [\sum_{t=0}^{\infty} \delta^t U_t | \mathcal{S}_0]$, where $\delta \in (0, 1)$ denotes the discount factor. To guarantee that the process of the expected net payoffs is

well behaved, we assume that for any state \mathcal{S} and policy χ , $\delta^T \mathbb{E}^\chi [\sum_{s=T}^\infty \delta^s U_s | \mathcal{S}] \rightarrow 0$ as $T \rightarrow \infty$.⁵³

Proof of Theorem 1. The proof is in three steps. Step 1 below first establishes the result in part (ii) and then uses the recursive representation of the index of search in (3) to show that, when the DM follows an index policy, her expected (per-period) payoff satisfies the representation in (4), thus establishing part (iii). Steps 2 and 3 then use the representation in (4) to show that the DM's payoff under the proposed index rule satisfies the Bellman equation for the dynamic program under consideration, thus proving the optimality of the index policy.

Step 1. We start by proving part (ii). Let $\hat{\tau}$ be the optimal stopping time in the definition of $\mathcal{I}^S(\omega^S)$. Note that, at $\hat{\tau}$, the index of each alternative in the CS (including search) must be weakly smaller than $\mathcal{I}^S(\omega^S)$. Otherwise, by selecting one of the alternatives (including search) in the CS at $\hat{\tau}$ for which the index is larger than $\mathcal{I}^S(\omega^S)$ and stopping optimally from that moment onward, the DM would attain an average payoff

$$\frac{\mathbb{E}^\pi [\sum_{s=0}^{\hat{\tau}-1} \delta^s U_s | \omega^S]}{\mathbb{E}^\pi [\sum_{s=0}^{\hat{\tau}-1} \delta^s | \omega^S]}$$

strictly greater than $\mathcal{I}^S(\omega^S)$, contradicting the definition of $\mathcal{I}^S(\omega^S)$.⁵⁴ This implies that $\hat{\tau}$ is weakly greater than τ^* . Moreover, since at τ^* the index of each alternative in the CS (including search) is weakly below $\mathcal{I}^S(\omega^S)$, the average payoff between τ^* and $\hat{\tau}$ is at most $\mathcal{I}^S(\omega^S)$. Hence, the average payoff under the optimal selection rule in the definition of $\mathcal{I}^S(\omega^S)$ from 0 to τ^* must be exactly equal to $\mathcal{I}^S(\omega^S)$. This implies that the optimal stopping time in the definition of $\mathcal{I}^S(\omega^S)$ can be taken to be τ^* . Because the index policy χ^* selects in each period between 0 and τ^* the alternative for which the average payoff is the largest (including search), we have that the optimal selection rule in the definition of $\mathcal{I}^S(\omega^S)$ must coincide with the index policy χ^* . That $\mathcal{I}^S(\omega^S)$ satisfies the representation in part (ii) then follows from the arguments above.

Next, consider part (iii). We construct the following stochastic process based on the values of the indices, and the expansion of the CS through search, *under the index policy*. Starting with the initial state $\mathcal{S}_0 = (\mathcal{S}_0^P, \omega_0^S)$, let $v^0 = \max\{\mathcal{I}^*(\mathcal{S}_0^P), \mathcal{I}^S(\omega_0^S)\}$. Let $t(v^0)$ be the first time at which, when the DM follows the policy χ^* , all indices are *strictly* below v^0 , with $t(v^0) = \infty$ if this event never occurs. Note that $t(v^0)$ dif-

⁵³This property is immediately satisfied if payoffs and costs are uniformly bounded.

⁵⁴Since infinity is allowed as a value of the stopping time, the supremum in the definitions of \mathcal{I}^P and \mathcal{I}^S is attained, that is, an optimal stopping time exists (see Mandelbaum, 1986).

fers from $\kappa(v^0)$, as $\kappa(v^0) = 0$ is the first time at which all indices are *weakly* below v^0 . Next let $v^1 \equiv \max\{\mathcal{I}^*(\mathcal{S}_{t(v^0)}^P), \mathcal{I}^S(\omega_{t(v^0)}^S)\}$ be the value of the largest index at $t(v^0)$, where $\mathcal{S}_{t(v^0)} = (\mathcal{S}_{t(v^0)}^P, \omega_{t(v^0)}^S)$ is the state of the system in period $t(v^0)$. Note that, by construction, $t(v^0) = \kappa(v^1)$. Furthermore, when $t(v^0) < \infty$, if $v^0 > \mathcal{I}^S(\omega_0^S)$ then $\omega_{t(v^0)}^S = \omega_0^S$. We can proceed in this manner to obtain a strictly decreasing sequence of values $(v^i)_{i \geq 0}$, with corresponding stochastic times $(\kappa(v^i))_{i \geq 0}$. Note that the values v^i are all non-negative, as the DM's outside option is normalized to zero. Next, for any $i = 0, 1, 2, \dots$, let $\eta^i \equiv \sum_{s=\kappa(v^i)}^{\kappa(v^{i+1})-1} \delta^{s-\kappa(v^i)} U_s$ denote the discounted sum of the net payoffs between periods $\kappa(v^i)$ and $\kappa(v^{i+1}) - 1$, when the DM follows the index policy and let $(\eta^i)_{i \geq 0}$ define the corresponding sequence of discounted accumulated net payoffs, with $\eta^i = 0$ if $\kappa(v^i) = \infty$.

Denote by $\mathcal{V}(\mathcal{S}_0)$ the expected (per-period) net payoff under the index policy χ^* , given the initial state of the problem \mathcal{S}_0 . That is, $\mathcal{V}(\mathcal{S}_0) = (1 - \delta)\mathbb{E}^{\chi^*} [\sum_{t=0}^{\infty} \delta^t U_t | \mathcal{S}_0]$. By definition of the processes $(\kappa(v^i))_{i \geq 0}$ and $(\eta^i)_{i \geq 0}$, $\mathcal{V}(\mathcal{S}_0) = (1 - \delta)\mathbb{E}^{\chi^*} [\sum_{i=0}^{\infty} \delta^{\kappa(v^i)} \eta^i | \mathcal{S}_0]$. Next, using the definition of the indices (1) and (2), observe that

$$v^i = \frac{(1 - \delta)\mathbb{E}^{\chi^*} [\eta^i | \mathcal{S}_{\kappa(v^i)}]}{\mathbb{E}^{\chi^*} [1 - \delta^{\kappa(v^{i+1}) - \kappa(v^i)} | \mathcal{S}_{\kappa(v^i)}]}. \quad (16)$$

To see why (16) holds, recall that, at period $\kappa(v^i)$, given the state of the system $\mathcal{S}_{\kappa(v^i)}$, the optimal stopping time in the definition of the index v^i is the first time at which the index of the physical alternative corresponding to v^i (if v^i corresponds to a physical alternative), or the index of search and of all alternatives introduced through future searches (in case v^i corresponds to the index of the search technology), drop below v^i .⁵⁵

Rearranging, multiplying both sides of (16) by $\delta^{\kappa(v^i)}$, and using the fact that $\delta^{\kappa(v^i)}$ is known at $\kappa(v^i)$, we have that

$$(1 - \delta)\mathbb{E}^{\chi^*} [\delta^{\kappa(v^i)} \eta^i | \mathcal{S}_{\kappa(v^i)}] = v^i \mathbb{E}^{\chi^*} [\delta^{\kappa(v^i)} - \delta^{\kappa(v^{i+1})} | \mathcal{S}_{\kappa(v^i)}].$$

Taking expectations of both sides of the previous equality given the initial state \mathcal{S}_0 , and

⁵⁵Note that if, at period $\kappa(v^i)$, there are multiple options ("physical" alternatives and search) with index v^i , the average sum $\mathbb{E}^{\chi^*} [\eta^i | \mathcal{S}_{\kappa(v^i)}]$ of the discounted net payoffs across all alternatives with index v^i until the indices of all such options drop below v^i (in case search is included, also those of the new arriving alternatives), per unit of average discounted time, $\mathbb{E}^{\chi^*} [1 - \delta^{\kappa(v^{i+1}) - \kappa(v^i)} | \mathcal{S}_{\kappa(v^i)}] / (1 - \delta)$, is the same as the average sum of the discounted net payoffs of each individual option with index v^i normalized by the average discounted time until the index of that alternative falls below v^i . This follows from the independence of the processes. Hence, Condition (16) holds irrespectively of whether, at $\kappa(v^i)$, there is a single or multiple options with index v^i .

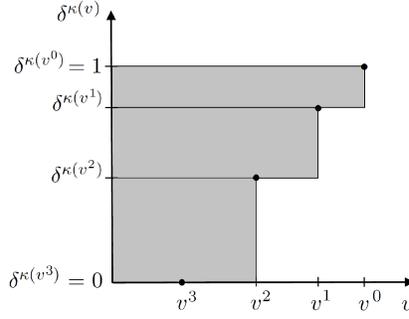


Figure 1: An illustration of the function $\delta^{\kappa(v)}$ and the region $\sum_{i=0}^{\infty} v^i (\delta^{\kappa(v^i)} - \delta^{\kappa(v^{i+1})}) = \int_0^{\infty} v d\delta^{\kappa(v)}$, for a particular path with $\kappa(v^3) = \infty$.

using the law of iterated expectations, we have that

$$(1 - \delta)\mathbb{E}^{\chi^*} [\delta^{\kappa(v^i)} \eta^i | \mathcal{S}_0] = \mathbb{E}^{\chi^*} [v^i (\delta^{\kappa(v^i)} - \delta^{\kappa(v^{i+1})}) | \mathcal{S}_0].$$

It follows that

$$\mathcal{V}(\mathcal{S}_0) = \mathbb{E}^{\chi^*} \left[\sum_{i=0}^{\infty} v^i (\delta^{\kappa(v^i)} - \delta^{\kappa(v^{i+1})}) | \mathcal{S}_0 \right]. \quad (17)$$

Next, note that $\delta^{\kappa(v^i)} = 0$ whenever $\kappa(v^i) = \infty$, and that, for any $i = 0, 1, \dots$, $\kappa(v) = \kappa(v^{i+1})$ for all $v^{i+1} < v < v^i$. It follows that (17) is equivalent to

$$\mathcal{V}(\mathcal{S}_0) = \mathbb{E}^{\chi^*} \left[\int_0^{\infty} v d\delta^{\kappa(v)} | \mathcal{S}_0 \right] = \mathbb{E}^{\chi^*} \left[\int_0^{\infty} (1 - \delta^{\kappa(v)}) dv | \mathcal{S}_0 \right] = \int_0^{\infty} (1 - \mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | \mathcal{S}_0]) dv. \quad (18)$$

The construction of the integral function (18) is illustrated in Figure 1.

Step 2. We use the representation of the DM's payoff under the index rule in (4) 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. This will permit us to establish in Step 3 the optimality of χ^* through dynamic programming.

We first introduce some notation. Define $\mathcal{S}_t^1 \vee \mathcal{S}_t^2 \equiv (\mathcal{S}_t^1(\omega) + \mathcal{S}_t^2(\omega) : \omega \in \Omega)$ and $\mathcal{S}_t^1 \setminus \mathcal{S}_t^2 \equiv (\max\{\mathcal{S}_t^1(\omega) - \mathcal{S}_t^2(\omega), 0\} : \omega \in \Omega)$. 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 $\mathcal{S}_t(\hat{\omega}^S) = 1$ and such that $\mathcal{S}_t(\omega^S) = 0$ 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 state by $e_t(\omega)$, where $\omega \in \Omega$ is the state of the search technology in case (i) and of the single physical alternative in case (ii).⁵⁶

Lemma 1. *For any $v \in \mathbb{R}$ and states \mathcal{S}^1 and \mathcal{S}^2 , $\kappa(v|\mathcal{S}^1 \vee \mathcal{S}^2) = \kappa(v|\mathcal{S}^1) + \kappa(v|\mathcal{S}^2)$.*

Proof of Lemma 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 θ of the various alternatives evolve independently of one another and of the evolution 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 evolution of 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 indices 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 $\mathbb{E}[u|\omega^P]$ the immediate expected payoff from exploring an alternative in state ω^P (the expectation is taken under the distribution H_{ω^P}) and by $\tilde{\omega}^P$ the new state of that alternative triggered by its exploration. Let

$$V^P(\omega^P|\mathcal{S}_0) \equiv (1 - \delta)\mathbb{E}^{\chi^*}[u|\omega^P] + \delta\mathbb{E}^{\chi^*}[\mathcal{V}(\mathcal{S}_0 \setminus e(\omega^P) \vee e(\tilde{\omega}^P))|\omega^P] \quad (19)$$

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

$$V^S(\omega^S|\mathcal{S}_0) \equiv -(1 - \delta)\mathbb{E}^{\chi^*}[c|\omega^S] + \delta\mathbb{E}^{\chi^*}[\mathcal{V}(\mathcal{S}_0 \setminus e(\omega^S) \vee e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S))|\omega^S] \quad (20)$$

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 $\mathbb{E}^{\chi^*}[c|\omega^S]$ is the

⁵⁶Throughout the analysis below, we maintain the assumption that an outside option with value equal to zero is available to the DM. That is, although we do not explicitly treat the outside option as a separate alternative, we implicitly assume that each state \mathcal{S} includes this option.

immediate expected cost from searching, $\tilde{\omega}^S$ is the new state of search after the first search is carried out, and $W^P(\tilde{\omega}^S)$ is the state of the new alternatives brought to the CS by search, with c and $W^P(\tilde{\omega}^S)$ jointly drawn from the distribution H_{ω^S} .⁵⁷

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 $\kappa(v|\mathcal{S}_0 \vee e(\omega_M^A)) = \kappa(v|\mathcal{S}_0)$, and if $v < M$, then $\kappa(v|\mathcal{S}_0 \vee e(\omega_M^A)) = \infty$. Hence, the representation in (4), adapted to the fictitious environment that includes the auxiliary option, implies that

$$\begin{aligned} \mathcal{V}(\mathcal{S}_0 \vee e(\omega_M^A)) &= \int_0^\infty \left(1 - \mathbb{E}^{\chi^*} \left[\delta^{\kappa(v)}|\mathcal{S}_0 \vee e(\omega_M^A)\right]\right) dv = M + \int_M^\infty \left(1 - \mathbb{E}^{\chi^*} \left[\delta^{\kappa(v)}|\mathcal{S}_0\right]\right) dv \\ &= \mathcal{V}(\mathcal{S}_0) + \int_0^M \mathbb{E}^{\chi^*} \left[\delta^{\kappa(v)}|\mathcal{S}_0\right] dv. \end{aligned} \quad (21)$$

The definition of χ^* , along with Conditions (19) and (20), then imply the following:

Lemma 2. *For any (ω^S, ω^P, M) ,*

$$\mathcal{V}(e(\omega^S) \vee e(\omega_M^A)) = \begin{cases} V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A)) & \text{if } M \leq \mathcal{I}^S(\omega^S) \\ M > V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A)) & \text{if } M > \mathcal{I}^S(\omega^S) \end{cases} \quad (22)$$

$$\mathcal{V}(e(\omega^P) \vee e(\omega_M^A)) = \begin{cases} V^P(\omega^P|e(\omega^P) \vee e(\omega_M^A)) & \text{if } M \leq \mathcal{I}^P(\omega^P) \\ M > V^P(\omega^P|e(\omega^P) \vee e(\omega_M^A)) & \text{if } M > \mathcal{I}^P(\omega^P). \end{cases} \quad (23)$$

Proof of Lemma 2. First note that the index corresponding to the auxiliary option is equal to M . Hence, if $M \leq \mathcal{I}^S(\omega^S)$, χ^* prescribes to start with search. If, instead, $M > \mathcal{I}^S(\omega^S)$, χ^* prescribes to select the auxiliary option forever, yielding 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 $V^S(\omega^S|e(\omega^S) \vee e(\omega_M^A)) = \mathbb{E}^{\chi^*} \left[(1 - \delta) \sum_{s=0}^{\hat{\tau}-1} \delta^s U_s^{\hat{\pi}} + \delta^{\hat{\tau}} M | \omega^S \right]$, for some stopping and selection rules $\hat{\tau}, \hat{\pi}$. 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

⁵⁷Note that $W^P(\tilde{\omega}^S)$ is a deterministic function of the new state $\tilde{\omega}^S = ((c_0, E_0), (c_1, E_1))$ of search. To see this, recall that, for any m , the function E_m counts how many alternatives have been added to the CS, for each possible state ω^P , as a result of the m -th search.

subsequent periods. By definition of $\mathcal{I}^S(\omega^S)$,

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

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

$$V^S(\omega^S | e(\omega^S) \vee e(\omega_M^A)) = \mathbb{E}^{\hat{\pi}} \left[(1 - \delta) \sum_{s=0}^{\hat{\tau}-1} \delta^s U_s + \delta^{\hat{\tau}} M | \omega^S \right] < M \mathbb{E}^{\hat{\pi}} \left[(1 - \delta) \sum_{s=0}^{\hat{\tau}-1} \delta^s + \delta^{\hat{\tau}} | \omega^S \right] = M.$$

Similar arguments establish Condition (23). \square

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 physical 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.⁵⁸

Lemma 3. *Let \mathcal{S}_0 be the initial state of the decision problem, with $\omega^S \in \Omega^S$ denoting the state of the search technology, as specified in \mathcal{S}_0 . We have that⁵⁹*

$$\begin{aligned} D^S(\omega^S | \mathcal{S}_0) &= \int_0^{\mathcal{I}^*(\mathcal{S}_0^P)} D^S(\omega^S | e(\omega^S) \vee e(\omega_v^A)) d\mathbb{E}^{\chi^*} \left[\delta^{\kappa(v)} | \mathcal{S}_0 \setminus e(\omega^S) \right] \\ &\quad + \mathbb{E}^{\chi^*} \left[\delta^{\kappa(0)} | \mathcal{S}_0 \setminus e(\omega^S) \right] D^S(\omega^S | e(\omega^S) \vee e(\omega_0^A)). \end{aligned} \quad (24)$$

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

$$\begin{aligned} D^P(\omega^P | \mathcal{S}_0) &= \int_0^{\max\{\mathcal{I}^*(\mathcal{S}_0^P \setminus e(\omega^P)), \mathcal{I}^S(\omega^S)\}} D^P(\omega^P | e(\omega^P) \vee e(\omega_v^A)) d\mathbb{E}^{\chi^*} \left[\delta^{\kappa(v)} | \mathcal{S}_0 \setminus e(\omega^P) \right] \\ &\quad + \mathbb{E}^{\chi^*} \left[\delta^{\kappa(0)} | \mathcal{S}_0 \setminus e(\omega^P) \right] D^P(\omega^P | e(\omega^P) \vee e(\omega_0^A)). \end{aligned} \quad (25)$$

Proof of Lemma 3. Using Condition (21), we have that, given the state $\mathcal{S}_0 \vee e(\omega_M^A)$ of

⁵⁸In 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 .

⁵⁹Recall that $\mathcal{I}^*(\mathcal{S}_0^P)$ is the largest index of the alternatives in the CS under the state \mathcal{S}_0 .

the decision problem, and $\omega^S \in \Omega^S$,

$$D^S(\omega^S | \mathcal{S}_0 \vee e(\omega_M^A)) = \mathcal{V}(\mathcal{S}_0) + \int_0^M \mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | \mathcal{S}_0] dv + (1 - \delta) \mathbb{E}^{\chi^*} [c | \omega^S] \\ - \delta \mathbb{E}^{\chi^*} \left[\mathcal{V}(\mathcal{S}_0 \setminus e(\omega^S) \vee e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S)) + \int_0^M \mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | \mathcal{S}_0 \setminus e(\omega^S) \vee e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S)] dv | \omega^S \right], \quad (26)$$

where the equality follows from combining (20) with (21). Similarly,

$$D^S(\omega^S | e(\omega^S) \vee e(\omega_M^A)) = \mathcal{V}(e(\omega^S)) + \int_0^M \mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | e(\omega^S)] dv + (1 - \delta) \mathbb{E}^{\chi^*} [c | \omega^S] \\ - \delta \mathbb{E}^{\chi^*} \left[\mathcal{V}(e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S)) + \int_0^M \mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S)] dv | \omega^S \right]. \quad (27)$$

Differentiating (26) and (27) with respect to M , using the independence across alternatives and search and Lemma 1, we have that

$$\frac{\partial}{\partial M} D^S(\omega^S | \mathcal{S}_0 \vee e(\omega_M^A)) = \mathbb{E}^{\chi^*} [\delta^{\kappa(M)} | \mathcal{S}_0 \setminus e(\omega^S)] \frac{\partial}{\partial M} D^S(\omega^S | e(\omega^S) \vee e(\omega_M^A)). \quad (28)$$

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 indices 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 $\omega^P \in \{\hat{\omega}^P \in \Omega^P : \mathcal{S}_0(\hat{\omega}^P) > 0\}$,

$$\frac{\partial}{\partial M} D^P(\omega^P | \mathcal{S}_0 \vee e(\omega_M^A)) = \mathbb{E}^{\chi^*} [\delta^{\kappa(M)} | \mathcal{S}_0 \setminus e(\omega^P)] \frac{\partial}{\partial M} D^P(\omega^P | e(\omega^P) \vee e(\omega_M^A)). \quad (29)$$

Let $M^* \equiv \max\{\mathcal{I}^*(\mathcal{S}_0^P), \mathcal{I}^S(\omega^S)\}$. Integrating (28) over the interval $(0, M^*)$ of possible values for the auxiliary option and rearranging, we have that

$$D^S(\omega^S | \mathcal{S}_0 \vee e(\omega_0^A)) = D^S(\omega^S | \mathcal{S}_0 \vee e(\omega_{M^*}^A)) - \int_0^{M^*} \mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | \mathcal{S}_0 \setminus e(\omega^S)] \frac{\partial}{\partial v} D^S(\omega^S | e(\omega^S) \vee e(\omega_v^A)) dv \\ = D^S(\omega^S | \mathcal{S}_0 \vee e(\omega_{M^*}^A)) - D^S(\omega^S | e(\omega^S) \vee e(\omega_{M^*}^A)) \\ + \mathbb{E}^{\chi^*} [\delta^{\kappa(0)} | \mathcal{S}_0 \setminus e(\omega^S)] D^S(\omega^S | e(\omega^S) \vee e(\omega_0^A)) \\ + \int_0^{M^*} D^S(\omega^S | e(\omega^S) \vee e(\omega_v^A)) d\mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | \mathcal{S}_0 \setminus e(\omega^S)],$$

where the second equality follows from integration by parts and from the fact that $\mathbb{E}^{\chi^*} [\delta^{\kappa(M^*)} | \mathcal{S}_0 \setminus e(\omega^S)] = 1$. 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 $D^S(\omega^S | \mathcal{S}_0 \vee$

$e(\omega_{M^*}^A)) = D^S(\omega^S | e(\omega^S) \vee e(\omega_{M^*}^A))$.⁶⁰ Therefore, we have that

$$\begin{aligned} D^S(\omega^S | \mathcal{S}_0) &= \int_0^{M^*} D^S(\omega^S | e(\omega^S) \vee e(\omega_v^A)) d\mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | \mathcal{S}_0 \setminus e(\omega^S)] \\ &\quad + \mathbb{E}^{\chi^*} [\delta^{\kappa(0)} | \mathcal{S}_0 \setminus e(\omega^S)] D^S(\omega^S | e(\omega^S) \vee e(\omega_0^A)). \end{aligned} \quad (30)$$

Similar arguments imply that

$$\begin{aligned} D^P(\omega^P | \mathcal{S}_0) &= \int_0^{M^*} D^P(\omega^P | e(\omega^P) \vee e(\omega_v^A)) d\mathbb{E}^{\chi^*} [\delta^{\kappa(v)} | \mathcal{S}_0 \setminus e(\omega^P)] \\ &\quad + \mathbb{E}^{\chi^*} [\delta^{\kappa(0)} | \mathcal{S}_0 \setminus e(\omega^P)] D^P(\omega^P | e(\omega^P) \vee e(\omega_0^A)). \end{aligned} \quad (31)$$

To complete the proof of Lemma 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 (24) in the lemma follows directly from (30). Thus consider Condition (25). First observe that, for any state $\omega^P \in \Omega^P$ such that $M^* > \max\{\mathcal{I}^*(\mathcal{S}_0^P \setminus e(\omega^P)), \mathcal{I}^S(\omega^S)\}$, 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 (31) is equal to zero over the interval $[0, \mathcal{I}^P(\omega^P)]$ and hence also over the interval $[0, \max\{\mathcal{I}^*(\mathcal{S}_0^P \setminus e(\omega^P)), \mathcal{I}^S(\omega^S)\}]$. We thus have that, in this case, Condition (25) clearly holds. Next observe that, for any state $\omega^P \in \Omega^P$ such that $M^* = \max\{\mathcal{I}^*(\mathcal{S}_0^P \setminus e(\omega^P)), \mathcal{I}^S(\omega^S)\}$, Condition (25) follows directly from (31).

Case (2): given \mathcal{S}_0 , χ^* specifies starting with search (i.e., $M^* = \mathcal{I}^S(\omega^S)$). Then, for any $\omega^P \in \Omega^P$, $\max\{\mathcal{I}^*(\mathcal{S}_0^P \setminus e(\omega^P)), \mathcal{I}^S(\omega^S)\} = M^*$, in which case Condition (25) in the lemma follows directly from (31). That Condition (24) 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 (30) is equal to zero over the entire interval $[0, \max\{\mathcal{I}^*(\mathcal{S}_0^P \setminus e(\omega^P)), \mathcal{I}^S(\omega^S)\}]$. \square

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

Lemma 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 $\mathcal{I}^S(\omega^S) \geq \mathcal{I}^*(\mathcal{S}_0^P)$;

⁶⁰This 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 $\mathbb{E}^{\chi^*} [\mathcal{V}(\mathcal{S}_0 \setminus e(\omega^S) \vee e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S) \vee e(\omega_{M^*}^A)) | \omega^S] = \mathbb{E}^{\chi^*} [\mathcal{V}(e(\tilde{\omega}^S) \vee W^P(\tilde{\omega}^S) \vee e(\omega_{M^*}^A)) | \omega^S]$. 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^* .

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 4. *Part 1.* First, use (22) 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 if $v \leq \mathcal{I}^S(\omega^S)$. Therefore, from (24), $D^S(\omega^S|\mathcal{S}_0) \geq 0$ – 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 (23) 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 (25), $D^P(\omega^P|\mathcal{S}_0) \geq 0$ with the inequality holding as equality if and only if $\mathcal{I}^P(\omega^P) \geq \max\{\mathcal{I}^*(\mathcal{S}_0^P \setminus e(\omega^P)), \mathcal{I}^S(\omega^S)\}$. The result in part 2 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 $\delta^T \mathbb{E}^\chi [\sum_{s=T}^\infty \delta^s U_s | \mathcal{S}] \rightarrow 0$ as $T \rightarrow \infty$ guarantees $\mathcal{V}(\mathcal{S}_0) = \mathcal{V}^*(\mathcal{S}_0)$, and hence the optimality of χ^* . \square

This completes the proof of the theorem. \blacksquare

Proof of Proposition 3. Since product 0 corresponds to the outside option, one of the products 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, 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 $\mathcal{I}^S(m) \geq d_m > d_l = \min\{\mathcal{I}_l, (1 - \delta)v_l, \mathcal{I}^S(l)\}$, 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

$$\min\{\mathcal{I}_l, (1 - \delta)v_l\} < d_m \leq \min\{\mathcal{I}_m, (1 - \delta)v_m\}. \quad (32)$$

Furthermore, because the search technology is non-improving, $\mathcal{I}^S(l + 1) \geq \dots \geq \mathcal{I}^S(m - 1) \geq \mathcal{I}^S(m)$. 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 (32) 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. ■