Searching in the Dark and Learning Where to Look

Suraj Malladi*

January 6, 2023

Abstract

We study a model of ordered search with learning. Our focus is on settings where an agent knows little about how observable attributes of items map to quality but wants to avoid missing good discoveries. Knowing only that items with similar attributes have similar qualities, she sequentially decides where to look next or stops searching to take her best discovery to date. We characterize worst-case optimal search procedures. Search effort is 'inverse-v-shaped' in search complexity. If there is a sweet-spot in attribute space, the search path is 'funnel-shaped'. Optimal search is dynamically consistent and computationally tractable. And the model readily generalizes to limited consideration sets and multidimensional search.

*Thanks to Martino Banchio, Daniel Chen, Tommaso Denti, David Easley, Fancesca Molinari, Andrzej Skrzypacz, Joerg Stoye, Ruqing Xu, seminar participants at SITE: Dynamic Games, Northwestern and Arizona State University for their feedback and suggestions. I am particularly grateful to Ilya Morozov for so many helpful conversations. Cornell University. Email: surajm@cornell.edu.

1 Introduction

People search for various things: designs for a new product, ideas for a paper, a new tennis racquet. They examine some items in detail, and in doing so get a sense for the quality of similar, unexplored alternatives. They then assess whether to settle for their best discoveries so far, and if not, which of many heterogeneous alternatives to consider next.

The basic question of *where* forward-looking searchers would explore given what they had learned is scarcely addressed in economic theory. Classic models treat search as a pure stopping problem (McCall, 1970; Rothschild, 1974) or assume that search outcomes are independent, shutting down scope for learning (Weitzman, 1979). Recent models study settings where search order is exogenously fixed (Urgun and Yariv, 2021) or searchers are either short-lived or myopic (Callander, 2011; Garfagnini and Strulovici, 2016; Hodgson and Lewis, 2020). A broad literature on bandits and optimization studies various heuristics rather than how rational agents optimally search and stop.

This paper studies ordered search with learning, particularly how forward-looking searchers explore *unfamiliar territories*. These are settings where searchers know little about how observable attributes of items map to quality but have incentives to avoid missing good discoveries, should they exist. We model searchers as maximizing their worst-case eventual payoffs so as to be robust to the shape of this unknown mapping. The model yields a tractable and intuitive characterization of dynamically optimal search. The following examples motivate the model:

Firm R&D A startup prototypes different designs of a new technology, as does its strong competitor. Small tweaks to the earlier designs can only affect quality by so much, but larger changes can produce a wider range of outcomes. Upon bringing its technology to market, the startup would make little profit should its competitor find a much better design. The startup extensively explores the space of possible designs to prevent this from happening.

Dose-ranging study A pharmaceutical company must choose the dosage of a compound in a new drug and again is concerned with a competitor. Similar dosage levels have similar effects and there is thought to be a sweet spot, or, 'therapeutic window': too high a dose is toxic and too low a dose is ineffective. But what is too high or low is to be determined through costly experimentation.

Consumer search An amateur photographer shops for his first camera, knowing little about how he would value various features. He can learn his expected surplus for a camera by taking the time to read its description and reviews. He would dislike finding out later that he missed out on a much better purchase and searches to avoid this scenario.

The base model considers a searcher who searches (or explores) items in some compact, one-dimensional attribute space. The searcher does not know the true mapping from attributes to quality, i.e., the quality index. She knows only that similar items have similar qualities (i.e., the quality index is Lipschitz continuous), and perhaps that there is a "sweet spot" in attribute space (i.e., the quality index is Lipschitz continuous and quasiconcave). The Lipschitz constant is her perception of search complexity.[1]

The searcher sequentially explores different items to learn their true quality, and in the process, narrows down the set of rationalizable quality indices. After each search, she decides whether to continue exploring or settle with the best item she had discovered so far.

Her payoff increases in the quality of the best item she discovers, decreases in the quality of the best attainable item, and decreases with costly search. At each turn, she acts to maximize her worst-case payoff upon eventually concluding search. She takes this worst-case over the rationalizable quality indices at that point in time. In effect, she fears going on a wild goose chase only to miss an unexplored high quality item. But the more she learns, the less scope there is for this unhappy outcome: if a nearby option were so good, then what she had found could not have been too far off.

We characterize optimal search procedures. Consider an alternative framework, akin to the classical model of simultaneous search (Stigler, 1961), where the searcher must irrevocably decide the sequence of searches she will make beforehand and pick the best of the discovered items. We show that all optimal sequential search procedures are 'greedy' optimal simultaneous search procedures: in each period, the searcher solves for an optimal simultaneous search procedure, and then take only the first step prescribed by that procedure. From this we derive comparative statics and an explicit optimal search algorithm.

One basic comparative static is that the search intensity is 'inverse-v-shaped' in search complexity. When search complexity is low, unexplored items must be similar in quality to those that have been discovered, so the searcher concludes search. The searcher also concludes when complexity is high: making good discoveries is like finding a needle in a haystack so the effort is not worth the cost. Under a decreasing returns to search assumption, search only takes place for intermediate levels of search complexity. And under constant search costs, the searcher plans to do more searches when complexity is higher, conditional on doing any search at all. This gives rise to a tipping point in complexity, where the searcher goes from exerting maximal effort to giving up on search entirely.

We also study the impact of past search outcomes on future search. The impact of discovering high or low quality items is particularly striking in the case where the searcher believes that the quality index is quasiconcave. Such discoveries alert the searcher to the location of the 'sweet spot' and rule out other regions of the search space. This produces funnel-like search dynamics empirically observed by Bronnenberg et al. (2016) and Blake

[1]For example, pharmaceutical firms may perceive high complexity when working with lithium. Small changes in lithium levels tend to have large effects on the efficacy or toxicity of a drug.

et al. (2016) in the context of online-shopping. The searcher initially explores broadly in attribute space but hones in on a particular point over time.

The results for the baseline model have close analogues in extensions to limited consideration sets (i.e., the searcher does not know all searchable items *a priori* but may learn of their existence over time), multidimensional attribute spaces (e.g., a shopper explores cameras by sensor size *and* resolution), and multidimensional learning (i.e., the shopper learns her values for sensor size and resolution separately and not just the camera as a whole).

Finally, we give polynomial-time algorithms for optimal search when search costs are constant. These algorithms tie out Weitzman's intuition about ordered search with learning:

> It appears plausible that other things being equal it would be better to open a box whose reward is highly correlated with other rewards because this adds a positive informational externality. But translating such an effect into a simple search rule seems difficult except in the most elementary cases. (Weitzman, 1979)

When exploring unfamiliar territory, the searcher picks her targets to minimize a measure of distance between the set of explored and 'relevant' unexplored items. And roughly speaking, proximity in search space corresponds to correlation in a Bayesian model.

1.1 Related literature

An broad literature on search with learning studies settings where search order is either irrelevant (i.e., items are homogeneous) or exogenous. Rothschild (1974) first considers an agent who draws independent samples from an unknown distribution and learns its shape while deciding when to stop. Bikhchandani and Sharma (1996) among others generalize these results to allow for recall and other distributions. Schlag and Zapechelnyuk (2021) find stopping policies that achieve a substantial fraction of achievable surplus regardless of the true distribution. In addition to optimal stopping, Urgun and Yariv (2021) study the speed with which searchers explore the path of a Brownian motion from left to right.[2]

Far less is known about optimal search with learning when agents freely choose where to look. Weitzman (1979) studies ordered search without learning: rewards to different items are independent so learning about one item teaches the agent nothing about another. While Weitzman's model with correlations is generally intractable, Adam (2001) solves an important special case where all items have independent and unknown distributions, but items of the same observable type are identically distributed. In a much richer setting, Callander (2011) characterizes where on a realized path of a Brownian motion a sequence

[2] Wong (2021) studies experimentation in a similar model with uncertain flow payoffs.

of short-lived social learners would choose to search.[3] Our primary theoretical contribution to analyze a model of ordered search with learning by a forward-looking agent (as in Adam (2001)) where the qualities of items are flexibly intertwined (as in Callander (2011)).

We follow a growing body of work in economics that considers maxmin or minmax regret objectives.[4] The modeling approach is also related to the applied mathematics literature on adversarial multi-arm bandits and Lipschitz function optimization (see reviews by Slivkins (2019) and Hansen et al. (1992), resp.).

In contrast to bandits, searchers do not face uncertain flow utility and the ensuing exploration-exploitation trade-off. And in contrast to optimization, rational search involves optimal stopping. A more methodological distinction is that we build a model suitable for analyzing optimal search and robust comparative statics under general preferences as opposed to heuristics achieving certain long-run linear regret bounds.

2 The model

2.1 Preliminaries

There is a set of items represented as a compact set $S \subset \mathbb{R}$. Let $Q \subset [0,1]^S$ be the set of potential *quality indices*—mappings from the search space to a measure of quality. There is some true quality index $q \in Q$, so each item $x \in S$ has a quality $q(x) \in [0,1]$.

There is a searcher who knows S and Q but not the identity of the true quality index. She can learn the quality of technologies in S through costly search. This way, she narrows down the set of candidate true quality indices in Q.

In each period, $t = 1, 2, 3 \ldots$, the searcher takes one of two kinds of actions. She either searches a new item $x_t \in S$ to learn its quality, $q(x_t)$. Or she concludes her search, $x_t = \emptyset$, and adopts the highest quality item that she had searched with so far.

Formally, let $h_t = \{(x_i, z_i)\}_{i=0}^{t-1}$ be the time t partial history when the searcher has not yet concluded search, with $z_i = q(x_i)$. Let X_{h_t} be the set of items which were searched at this history. Let $z^*_{h_t} = \max_{i=0,\ldots,t-1} z_i$. If $x_i \in X_{h_t}$ and $z_i = z^*_{h_t}$, then x_i is *an optimal item at* h_t. $X^*_{h_t} \subset X_{h_t}$ denotes the set of optimal items at h_t. Let $Q_{h_t} \subset Q$ be the set of quality indices that are *consistent* with what the searcher had observed so far at history h_t. That is, Q_{h_t} is the set of indices q' satisfying $q'(x_i) = z_i$ for all $i = 0, \ldots, t-1$. Let H denote the set of all partial histories h where Q_h is nonempty.

[3]Notably, Garfagnini and Strulovici (2016) study a bandit problem in a social learning framework similar to Callander (2011) but with forward-looking agents who live two periods.

[4]See reviews by Carroll (2019) (contracting), Manski (2011) (treatment choice), and Banerjee et al. (2017) (experimental design).

Let Q_L^{MP} denote the set of all Lipschitz continuous mappings $S \to [0,1]$ with Lipschitz constant $L > 0$. Let $Q_L^{QC} \subset Q_L^{MP}$ denote the set of all quasiconcave Lipschitz continuous mappings $S \to [0,1]$.[5] We assume the following throughout:

Assumption **1.** *Either* $Q = Q_L^{MP}$ *or* $Q = Q_L^{QC}$.

In words, the searcher little about the true quality index. She knows the limits to quality, as $0 < q < 1$. If $Q = Q_L^{MP}$, she knows that proximate items in S cannot be too different in quality. And if $Q = Q_L^{QC}$, she additionally knows that there is an ideal range (or "sweet spot") in S, and that items farther from this range are of lower quality. Under Assumption 1, an item's location in S can be thought of as an aggregate index of its observable attributes—items with similar attributes have similar qualities.

2.2 Payoffs

The searcher's benefit to adopting item x is,

$$U(q(x), \max_{y \in S} q(y)),$$

where q is the true quality index, and U is differentiable almost everywhere and continuous. U is non-decreasing in the quality of the adopted item and non-increasing in the quality of the best item in S. Moreover, $U_1 + U_2 \geq 0$ almost everywhere, so for any item choice, translating the true quality index q upward always benefits to the searcher.[6]

One interpretation of the payoff function is that a the searcher is a firm whose future profits depend not only on the quality of its chosen technology but also on the technology chosen by a strong potential competitor. In the worst case, the competitor discovers the best innovation possible, hurting the profits that the first firm stands to make.[7] Alternatively, the searcher could be an online shopper who experiences disutility from later learning that she missed out on buying a much better product.

[5] We include the subscript here to emphasize that L is known by the searcher to be an upper bound on the rate of change of q; but we drop the subscript elsewhere, taking this to be understood.

[6] This also ensures that if the searcher discovers the best quality item, she is weakly better off when that item is of higher quality.

[7] The model can also capture a searcher with an outside option that gives a payoff of u_0 independent of the true quality index. Note that such a searcher's benefit to concluding search at history h_t would be

$$U'(z^*_{h_t}, \max_{y \in S} q(y)) \equiv \begin{cases} u_0, \text{ if } \min_{q' \in Q_{h_t}} U(z^*_{h_t}, \max_{y \in S} q'(y)) < u_0. \\ U(z^*_{h_t}, \max_{y \in S} q(y)), \text{ otherwise.} \end{cases}$$

Since U is nondecreasing in the first argument and non-increasing in the second argument, the same is true for U'; and $U'_2 \leq U'_1$, as well.

The searcher's cost of searching item x in period t is given by

$$C(x, h_t),$$

where $C : S \times H \to \mathbb{R}_{++}$ is continuous and bounded away from 0. Next, C may depend on the sequence of items explored so far but not their qualities.[8]

When the quality index is q, the searcher's total payoff after concluding search at history $h_t \in H$ is given by:

$$p(h_t, q) = U(z^*_{h_t}, \max_{x \in S} q(x)) - \sum_{i=1}^{t} C(x_i, h_t(i)).$$

2.3 The searcher's problem

At each history, and for each plan of how to search thereafter, the searcher worries about the quality index that is consistent with her earlier searches but would minimize her payoff upon concluding search. She seeks a strategy that is *robust* to such scenarios, i.e., she searches as if the true quality index were adversarially selected.

We recast the searcher's problem as a dynamic zero-sum game. Each period proceeds in three stages. At history h_t, in the first stage, the searcher explores a new item, x_t, or concludes search, $\emptyset$. In the second stage, an imaginary adversary observes the searcher's action and picks a quality index $q' \in Q_{h_t}$ that is consistent with previous searches at this history. In the third stage, either the quality of this item is revealed; or if she had concluded search in the first stage, the searcher realizes her payoff $p(h_t, q')$.

Let $\sigma : H \to \Delta S \cup \{\emptyset\}$ and $\sigma^A : H \times \{S \cup \{\emptyset\}\} \to \Delta Q$ denote strategies of the searcher and its imagined adversary, respectively. A crucial detail is that the adversary responds to the realization of the searcher's choice after every partial history, so mixed strategies do not help the searcher. The adversary is constrained to choosing strategies which satisfy $\sigma^A(h_t, \cdot) \in \Delta Q_{h_t}$ for every $h_t \in H$; such strategies are said to be *feasible*.

Let $h^{(\sigma,\sigma^A)}$ denote the history after which the searcher concludes search when both agents follow their respective strategies. The payoff to the searcher is $p(h^{(\sigma,\sigma^A)}, \sigma^A(h^{(\sigma,\sigma^A)}))$ and the payoff to the adversary is $-p(h^{(\sigma,\sigma^A)}, \sigma^A(h^{(\sigma,\sigma^A)}))$. If the searcher never concludes search under (σ, σ^A), then let $p(h^{(\sigma,\sigma^A)}, \cdot) = -\infty$.

Given these preferences, if the searcher concludes search at a history h_t, she anticipates

[8] For example, in some settings, a searcher may find it less costly to experiment with technologies nearer to those she had explored earlier.

the adversary to choose a quality index in

$$\begin{aligned}&\arg\min_{q'\in Q_{h_t}}\left(U(z^*_{h_t},\max_{y\in S}q'(y))\right)\\ =&\arg\max_{q'\in Q_{h_t}}\left(\max_{y\in S}q'(y)\right).\end{aligned}$$

If the searcher continues searching at h_t, she anticipates that the adversary would take actions to minimize her eventual payoff inclusive of search costs. It is not immediately clear what those actions would be, because they would affect her search path and vice-versa.[9]

A strategy $\sigma : H \to \Delta S \cup \{\emptyset\}$ for the searcher is an *optimal sequential search procedure at* $h_t \in H$ if there exists a feasible strategy σ^A for the adversary such that (σ, σ^A) is a subgame-perfect equilibrium of the subgame starting at h_t. In subgame-perfect equilibria, the searcher chooses robustly optimal strategies even at off-path histories. We can therefore study how optimal search unfolds for any true $q \in Q$.

3 Simultaneous search procedures

We introduce a subclass of the searcher's strategy space and define an auxiliary game where the searcher is restricted to strategies in this class. This is useful for the analysis in Section 4.

A simultaneous search procedure is a strategy that does not depend on the qualities of items explored. For example, a searcher may decide to search items z, x, y and w in that order and regardless of what she learns about the qualities of these items along the way. Intuitively, if the searcher decides to tie her own hands in this manner, it is immaterial whether she learns along the way or learns the qualities of explored items 'simultaneously' at the end of her search when it is time to pick the best one.

Formally, two partial histories $h'_t, h''_t \in H$ *differ only by quality* if the same sequence of items are explored in both histories, i.e., $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1}$ and $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1}$. A strategy σ_s is a *simultaneous search procedure at* h_t if $\sigma_s(h'_\tau) = \sigma_s(h''_\tau)$ for any two histories $h'_\tau, h''_\tau \in H$ following h_t that differ only by quality.[10]

Let Γ_{h_t} denote all simultaneous search procedures at h_t. Consider a modification of the sub-game starting at history h_t where the searcher must commit to a strategy in Γ_{h_t} after which an adversary chooses a feasible quality index following every history and choice by the searcher (just as in the sequential search game). This is the *simultaneous search game at* h_t. A strategy σ_s is an *optimal simultaneous search procedure at* h_t if $\sigma_s \in \Gamma_{h_t}$ and if there exists

[9] For example, the searcher may worry about making discoveries that would trick her into going on a 'wild goose chase' with a lot of costly search but no results to show for it.

[10] Payoffs to simultaneous search procedures may differ if the order of search is permuted because search costs can depend on this order.

some strategy σ^A for the adversary such that (σ_s, σ^A) is a subgame perfect equilibrium of the simultaneous search game at h_t.[11]

Returning to the original sequential search game, we say a strategy σ *follows an optimal simultaneous search procedure* if at every history $h_t \in H$, there exists an optimal simultaneous search procedure at h_t, σ_s^*, such that $\sigma(h_t) = \sigma_s^*(h_t)$. In other words, this strategy solves for an optimal simultaneous search procedure at every history and takes the first action prescribed by that procedure. Similarly, for any $h_t \in H$, a strategy σ *follows an optimal simultaneous search procedure from h_t onward* if the same condition holds at h_t and for any history preceded by h_t.

Note that a strategy that follows an optimal simultaneous search procedure typically is not itself a simultaneous search procedure. Such a strategy uses the information about the latest discovery to recompute an optimal simultaneous search strategy every period, whereas the searcher's path of play under a simultaneous search procedure is unaffected by the discoveries made.

4 Optimal sequential search procedures

Section 4.1 characterizes optimal sequential search procedures as those that follow optimal simultaneous search procedures. Section 4.2 then characterizes all optimal simultaneous search procedures. Section 4.3 sketches the proofs and discusses the role of the assumptions. Section 4.4 describes an algorithm that returns optimal search procedures when search costs are constant.

4.1 Optimal sequential search

The first result establishes that following optimal simultaneous search procedures is optimal when $Q = Q^{MP}$ or $Q = Q^{QC}$.

Theorem **1.** *Let $h_t \in H$. A search strategy σ is an optimal sequential search procedure at h_t if and only if it follows an optimal simultaneous search procedure from h_t onward.*[12]

The proof of Theorem 1 is in Appendix F.1, while a sketch is given in Section 4.3.

Theorem 1 says that optimal sequential search is effectively 'greedy' simultaneous search. The searcher solves for an optimal simultaneous search procedure at every period, but only executes the first step of that plan. If the realized quality from that experiment is what she

[11] All the subgame-perfect equilibria of the simultaneous search game are also Nash equilibria, so this refinement is immaterial for the definition of optimal simultaneous search procedures.

[12] Continuity of C and U ensure that optimal sequential search procedures exist.

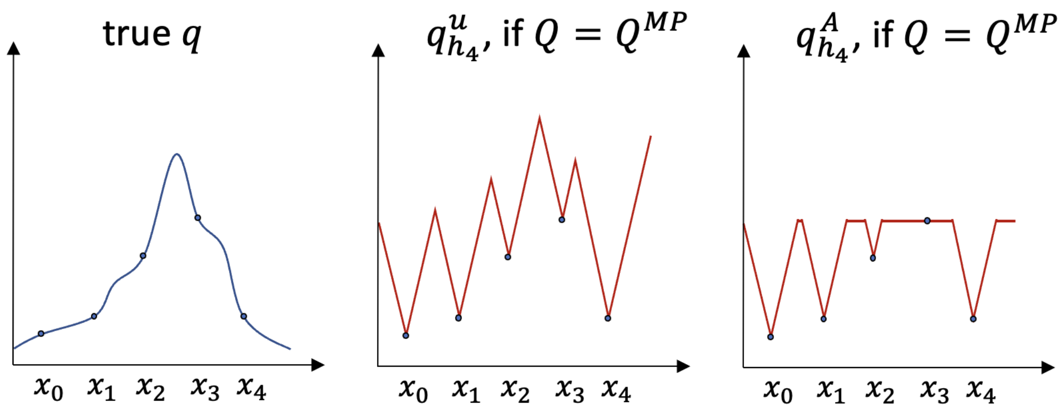

Figure 1: The figure on the left displays the true quality index. At history h_4, the searcher has observed the qualities at x_0, x_1, x_2, x_3 and x_4 in that order. The figure in the middle plots $q^u_{h_4}$, the upper envelope of feasible quality indices when $Q = Q^{MP}$, which is what the adversary chooses under σ^A_d if the searcher were to conclude search at this history. On the right is the adversary's chosen quality index under σ^A_d if the searcher were to continue search at this history. When the adversary plays σ^A_d, the 'no news' strategy, the searcher expects to discover items that are no better or no worse (if the feasibility constraint does not bind) than the best item she has discovered so far.

expected, she continues to carry out the previously calculated optimal simultaneous search procedure. If quality is different than what she anticipated, she formulates a new plan.

A key property of optimal sequential search procedures is that they satisfy certain notions of dynamic consistency. Even if the searcher could commit to a fully contingent search plan, committing to follow an optimal simultaneous search procedure remains an optimal strategy. Of course, this claim hinges on how the corresponding commitment game is modeled, in particular, on the timing of the adversary's moves (adaptive versus non-adaptive) and the space of the searcher's strategies (deterministic versus randomized). Appendix A shows that under natural formulations of the commitment game such that optimal sequential search procedures are indeed dynamically consistent.

4.2 Optimal simultaneous search

Theorem 1 reduces the problem of finding optimal sequential search procedures to finding optimal simultaneous search procedure at every history. Naturally, we next characterize optimal simultaneous search procedures.

Let $q^u_{h_t}$ denote the upper envelope of quality indices in Q_{h_t} for any $h_t \in H$, and let

$q^A_{h_t} = \min\{q^u_{h_t}, z^*_{h_t}\}$. Let $\sigma^A_d(h_t, x) = q^A_{h_t}$ for all $x \neq \emptyset$, and $\sigma^A_d(h_t, \emptyset) = q^u_{h_t}$. We refer to σ^A_d as the 'no news' strategy for the adversary, and Figure 1 explains why.

PROPOSITION 1. *Let $\sigma \in \Gamma_{h_t}$. Then σ is an optimal simultaneous search procedure at h_t if and only if σ is a best-response to σ^A_d in the simultaneous search game at $h_t \in H$.*

Proposition 1 only characterizes the searcher's behavior in simultaneous search games, not that of the fictitious adversary. There may be equilibria of the sequential search game where the adversary plays strategies other than σ^A_d. But by Proposition 1, for every such equilibrium, there is an equilibrium where the searcher's behaviour is unchanged but the adversary plays σ^A_d instead.

4.3 Discussion

We sketch the proofs of the results and highlight the role of the assumptions.

4.3.1 Proof ideas for Proposition 1

In the simultaneous search game, the adversary's 'no news' strategy serves two purposes when the searcher continues searching. First, the adversary does not let the quality of any new discovery exceed $z^*_{h_t}$ because that would improve the quality of the searcher's chosen item more than it increases the quality of the potential best alternative. Such a trade-off would go in the searcher's favor by the assumption that $U_1 + U_2 \geq 0$. Second, the adversary picks the quality of new discoveries to be as high as possible conditional on it not exceeding $z^*_{h_t}$. This improves the quality of the best unexplored alternative under some feasible quality index without improving the quality of the searcher's eventually chosen item.

This reasoning suggests that 'no news' is a weakly dominant strategy for the adversary in the sense that no other strategy delivers a strictly higher payoff for any strategy of the searcher. There are generally other weakly dominant strategies for the adversary. But by analogous arguments to the minimax theorem, swapping out the adversary's action profile with the 'no news' strategy yields another equilibrium. We therefore conclude that a simultaneous search procedure is optimal if and only if it best responds to the 'no news' strategy of the adversary.

4.3.2 Proof ideas for Theorem 1

Unlike in the simultaneous search game, the adversary's actions can affect the searcher's path of play in the sequential search game. These dynamics create an additional incentive for the adversary to keep the searcher engaged in wasteful and unproductive search. Restated from the searcher's perspective, she fears being led onto a 'wild goose chase' when she can gather and react to information dynamically.

Intuitively, it helps the adversary to keep the searcher in the dark as to where good or bad discoveries might be found. For example, upon discovering a very low quality item, the searcher learns that items with similar attributes are also low quality (by the Lipschitz continuity of the true quality index). She would thereafter redirect her attention to a more promising area of the search space or conclude search early. Clearly, the adversary would try to avoid tipping off the searcher in this way.

The 'no news' strategy shines in this respect as well. Discovering items with qualities higher or lower than $z^*_{h_t}$ would tip the searcher off as to which areas to explore or avoid. Discoveries that reveal neither good nor bad news minimize information leakage to the searcher.

In sum, the 'no news' strategy minimizes the quality of the searcher's best discoveries, maximizes the quality of the best unexplored alternative conditional on this, and minimizes information leakage to the searcher. Therefore, 'no news' is optimal for the adversary in the both the simultaneous and sequential search games. This implies that the set of searcher best responses at every history are also identical across both games, establishing Theorem 1.[13]

4.3.3 $Q = Q^{MP}$ versus $Q = Q^{QC}$

The path of optimal sequential search can play out very differently when $Q = Q^{MP}$ versus when $Q = Q^{QC}$. This is because at any history, the worst-case quality-index from the searcher's perspective depends on whether or not she believes the quality index to be quasiconcave. Figure 2 illustrates how, starting from the same history, the upper-envelope of feasible quality indices (and therefore σ^A_d and the searcher's best response) vary with Q.

4.3.4 The role of searcher's knowledge

The assumption that $Q = Q^{MP}$ or $Q = Q^{QC}$ is crucial for establishing a link between optimal simultaneous and sequential search. Greedy simultaneous search may be dominated by other sequential strategies when the searcher has other information about the shape of potential quality indices.

For example, let $Q = \{q_r, q_b, q_g\}$ as pictured in Figure 3. Suppose that search costs are constant: $C(x, h_t) = \epsilon$ for some $\epsilon > 0$, for all $x \in [0, 1]$ and $h_t \in H$. Suppose U is strictly

[13] This intuition helps us guess the form of the searcher's optimal sequential search procedure, and the proof in Appendix F verifies it. We first show that, in equilibrium, following optimal simultaneous search procedures (in the sequential search game) and committing to optimal simultaneous search procedures (in the simultaneous search game) both produce the same path of play. We then use this result and Proposition 1 to argue there are no profitable one-shot deviations for either player.

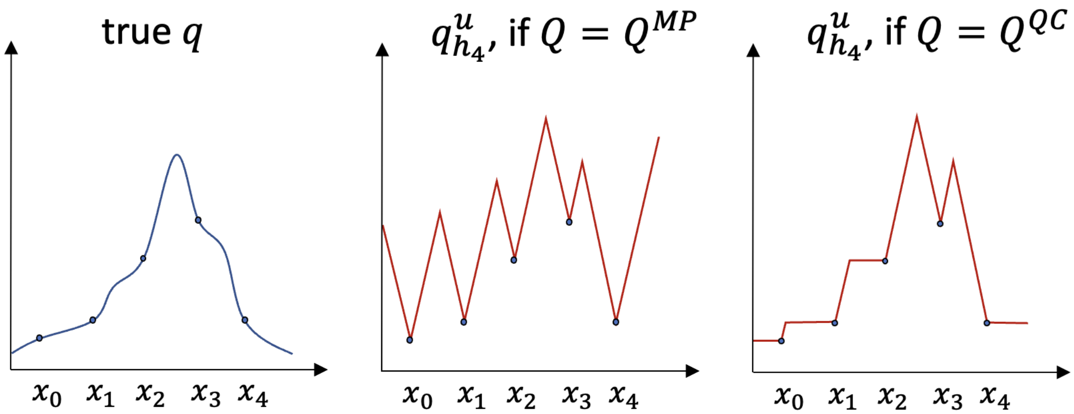

Figure 2: The left and middle figure are the same as in Figure 1. The figure on the right plots $q^u_{h_4}$ when $Q = Q^{QC}$ (recall that the upper envelope of quasiconcave functions need not be quasiconcave). When $Q = Q^{MP}$, the searcher may be interested in searching to the right of x_4 to rule out the possibility of finding something of high quality there. But when $Q = Q^{QC}$, the searcher is guaranteed not to make good discoveries in this region.

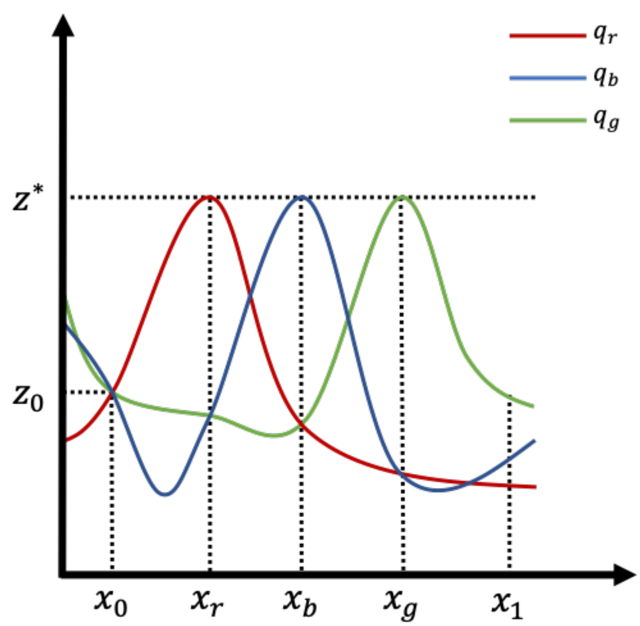

Figure 3: The quality indices q_r, q_b, and q_g attain their maxima at x_r, x_b, and x_g respectively.

increasing in the quality of the adopted item. The searcher starts by observing the quality of item x_0. Since $q_r(x_0) = q_b(x_0) = q_g(x_0) = z_0$, the searcher cannot narrow down the set of possible quality indices at history $h_0 = \{(x_0, z_0)\}$.

For ϵ sufficiently small, the optimal simultaneous procedure is to experiment with x_r, x_b and x_g in any order. One of these technologies is guaranteed to be of quality z^*. Experimenting with any other technologies is clearly wasteful. And leaving one or more of x_r, x_b or x_g unexplored lowers the searcher's worst-case payoff. For example, she worries that if she only searches x_b and x_g, q_r may turn out to be the true quality index.

On the other hand, no optimal sequential search procedure starts with x_r, x_b or x_g.

For instance, if the searcher first explores x_r, the true quality index would turn out not to be q_r in the worst case. Because $q_b(x_r) = q_g(x_r)$, the searcher would entertain the possibility that either q_b or q_g could be the true quality index. If she were to next explore x_b, the true quality index would turn out not to be q_b in the worst case. She would then know the true quality index is q_g and conclude search upon exploring x_g. So in the worst case, following an optimal simultaneous search procedure would take three searches to guarantee finding a maximal quality item.

However, if the searcher were to first explore x_1, she would immediately identifies the true quality index, as $q_r(x_1) \neq q_b(x_1) \neq q_g(x_1)$. In the next step, she would choose the technology that maximizes that index. This strategy guarantees z^* with two costly searches, rather than three. Therefore, the searcher is better off not following an optimal simultaneous search strategy at h_0.

In this example, the potentially high quality items are different that the more informative items; optimal simultaneous search procedures only explore the former while optimal sequential search procedures initially explore the latter. In exploring unfamiliar territories where the searcher does not have precise information about the possible shapes of the true quality index, such a distinction does not arise.[14]

4.4 An algorithm for optimal sequential search

We describe an algorithm for computing optimal sequential search procedures when search costs are constant ($C(\cdot,\cdot) = c > 0$) and the search space S is an interval in $\mathbb{R}$, leaving a more formal description and the case of finite search spaces for Appendix B.

[14] A related example where the link between simultaneous and sequential search may fail is when Q is the set of all polynomials of degree n. The issue here is that the searcher can back out the the true quality index with any $n + 1$ distinct searches.

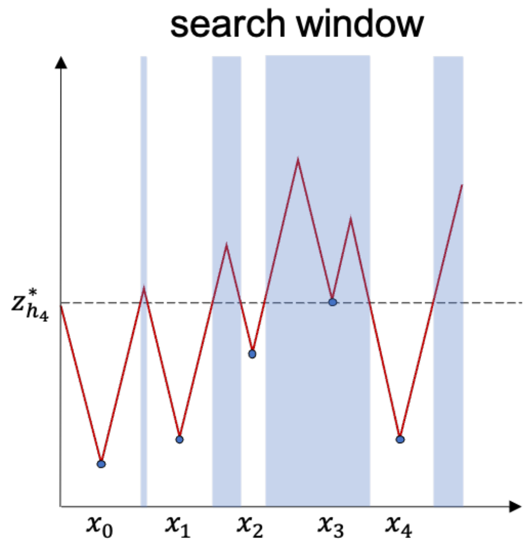

Figure 4: Continuing the example from Figure 2, the figure plots $q^u_{h_4}$ when $Q = Q^{MP}$ and highlights the search window.

4.4.1 Preliminaries

For any $h_t \in H$, the *search window*, $W_{h_t} \subset S$, is the set of technologies x for which $q^A_{h_t}(x) = z^*_{h_t}$ (or equivalently, $q^u_{h_t}(x) \geq z^*_{h_t}$); see Figure 4. Outside of the search window, the largest realizable quality is strictly below $z^*_{h_t}$ for any $q \in Q_{h_t}$. When costs are constant (or more generally, history-independent), there is no value in searching outside the search window. Such a search would return a lower payoff than $z^*_{h_t}$ without affecting the costs of future searches.

By Proposition 1, the searcher expects to discover a technology of quality $z^*_{h_t}$ for every search. This implies that the searcher expects the search window to remain unchanged over the course of search. Now the search window is a disjoint union of intervals, as pictured in Figure 4. Let Y_{h_t} be the union of (a) the endpoints of the disjoint intervals that comprise the search window and (b) the items in $X^*_{h_t}$ (e.g., x_3 in Figure 4).

At the midpoint of two items $y, y' \in Y_{h_t}$ that are in the same interval of the search window, $q^u_{h_t}$ has a local maximum. Clearly, the size of that local maximum is increasing in the distance between y and y'. Let h_{t+1} denote the history where $y'' \in (y, y')$ is explored next and has quality $z^*_{h_t}$, as the searcher expects (by Proposition 1). Then $q^u_{h_{t+1}}$ will have peaks in (y, y'') and (y'', y'), both smaller than the peak $q^u_{h_t}$ has in (y, y').

4.4.2 The algorithm

At the time of concluding search, the searcher's goal is to have the height of the tallest peak of q^u be as small as possible. The height of this peak is the quality of the best unexplored alternative in the worst case. If the searcher were to conclude search k periods from h_t, she can plan to pick her k search points so that (1) Y_{h_t} along with these k points partition W_{h_t}, and (2) the longest interval in this partition is as short as possible. Doing so minimizes the global maximum of q^u after k searches when the adversary plays the 'no news' strategy.

This logic pins down the location of searches conditional on k. The role of U and c is only to determine anticipated search intensity, k, i.e., the number of items the searcher plans to explore before concluding. Excessive search eventually pushes payoffs below the that of concluding search immediately. The number of searches $\overline{k}_t$ for which the cost, $\overline{k}_t \cdot c$, exceeds the best-case marginal benefit of optimally conducting $\overline{k}_t$ searches is as upper bound on k.[15]

The algorithm proceeds as follows. For each k such that $0 \leq k \leq \overline{k}_t$, the searcher determines the location of the k items that most evenly partition W_{h_t}. Next, she computes the payoff to conducting these k searches and concluding, should the adversary play the 'no news' strategy. She chooses a k^* corresponding to the highest total payoff among all $0 \leq k \leq \overline{k}_t$. She selects any one of the optimal k^* items and learns the quality of that item. She repeats these steps until $k^* = 0$, at which point she concludes search.[16]

4.4.3 The algorithm's time complexity

Initialized at a history h_t where the quality of n items are known, the time complexity of the algorithm is $O(\overline{k}_t^3 + \overline{k}_t^2 \cdot n)$. The algorithm therefore runs in polynomial time in the maximum conceivable number of searches and the number of items already explored.

The tractability of optimal sequential search strengthens the case for the model as a positive theory of how agents search. Indeed, one questions the predictions of models where optimization entails solving problems that are intractable even for computers.[17] Another implication of tractability is that the model may have use for estimating preferences of forward-looking agents from ordered-search data.

[15] The best case scenario after conducting any search is that the highest feasible quality is achieved. So the marginal benefit of search in the best case scenario is $U(\max_x q^u_{h_t}(x), \max_x q^u_{h_t}(x)) - U(z^*_{h_t}, \max_x q^u_{h_t}(x))$.

[16] Suppose the searcher has already explored l items since h_t, finding herself at a history h_{t+l}. She only considers search intensities k such $0 \leq k \leq \overline{k}_{t+l} \leq \overline{k}_t - l$ when computing her $l+1$st search. Therefore, the algorithm concludes in at most $\overline{k}_t$ steps.

[17] For example, Salant (2011) and Camara (2021) explore the behavioral consequences of the premise that agents cannot solve such computationally complex problems.

4.4.4 Discussion

In each period, the algorithm plans a set of searches so that upon the unexplored items in the search window will be close as possible (in the sense of Hausdorff distance) to the explored items upon concluding search. This procedure appears to reflect Weitzman's intuition that ordered search in the presence of correlated rewards should favor exploring items most correlated with other unexplored options. In the present prior-free model with continuous quality indices, proximity is roughly the analogue of correlation.

However, the underlying logic is distinct. The searcher here does not try to learn the location of good discoveries. Indeed, she searches as if she does not anticipate good discoveries. But by searching close to other unexplored items, she ensures that the items that she misses are not so different in quality than than those that she has searched.

Another feature of search with constant costs (and history-independent costs, more generally) is that the searcher is indifferent to the order in which she explores the items she plans on searching. Outside of this case, the searcher may prefer certain search orders; for example, searching farther away from the most recently searched item may be more costly.[18]

5 Search complexity and search intensity

Our results can be used to characterize how search complexity, measured by L, affects search intensity, the number of additional searches the searcher plans to conduct.

Search complexity captures how hard it is to discover a relatively good outcome. A pharmaceutical company may experiment with a compound whose efficacy and safety is typically very sensitive to dosage. Similarly, a software firm may be competing to make a product where slight differences in design significantly affect user experience and market share.

The first result is that search intensity is non-monotonic in search complexity. In particular, the searcher does not explore at all if search is sufficiently complex or simple.

Proposition **2.** *Suppose $|S| = \infty$. Let $h \in H$. There exist $\underline{L}, \overline{L} \in \mathbb{R}_{++}$ such that if search complexity $L < \underline{L}$ or $L > \overline{L}$, every optimal sequential search procedure at h concludes search immediately.*[19]

When L is sufficiently small, the searcher expects the quality index to be relatively flat and close to quality of preciously explored items. The value of search is therefore low.

[18] Taking the limit as such 'travel costs' go to zero is one way of selecting an equilibrium while retaining the tractability of the constant cost case. Appendix C discusses equilibrium refinements.

[19] We implicitly assume that L is such that h is a feasible history, i.e., Q_h is nonempty for this L.

If, on the other hand, L is sufficiently large, then the searcher is discouraged for a different reason: she may spend a lot of resources searching and still come nowhere near the peak outcome. Making a worthwhile discovery is like finding a needle in a haystack. This intuition hinges on the bounds to achievable quality (i.e., $0 \leq q \leq 1$) and having sufficiently many unexplored items to open. When there are many $x \in S$ which could potentially be of maximal quality, a single search is uninformative and too many items have to be opened to avoid missing out on good outcomes.

Under additional assumptions, Proposition 2 has a converse. We say there are *decreasing returns to search* if $U_{22} < 0$. In this case, the searcher experiences larger losses from missing better discoveries. Because additional search reduces the scope for missing better discoveries, the value of search decreases over time. Next, in light of Proposition 1, we say $h_t \in H$ is *on-path* if the quality of explored items is identical, i.e., $h_t = \{(x_i, z)\}_{i=0}^{t}$.

Proposition 3. *Suppose $|S| = \infty$. Let $h \in H$ be on-path and suppose there are decreasing returns to search. There exist $\underline{L}, \overline{L} \in \mathbb{R}_{++}$ such that if search complexity $L < \underline{L}$ or $L > \overline{L}$, every optimal sequential search procedure at h concludes search immediately. Conversely, if $L \in (\underline{L}, \overline{L})$, no optimal sequential search procedure concludes search immediately.*

Proposition 3 implies in particular that if the searcher starts out knowing only the quality of a single item, she continues search if and only if search complexity falls in an intermediate region. Figure 5 gives a graphical intuition for the claim.

Proposition 3 partitions $\mathbb{R}_{++}$ into regions of search complexity for which the searcher either concludes or continues searching under every equilibrium. We next pursue the question: how extensively does the searcher plan on searching if she indeed decides to continue?[20]

Let $\mathcal{I}^h(L)$ be the set of $n \in \mathbb{N}$ such that the searcher conducts n more searches in equilibrium at history $h \in H$ in some optimal sequential search procedure when search complexity is L. Define the *search intensity* correspondence as $\mathcal{I}^h : \mathbb{R}_+ \rightrightarrows \mathbb{N}$. Consider the case of *constant costs* where $C(x, h) = c > 0$ for all $x \in S$ and $h \in H$.

Proposition 4. *Let $h \in H$ be on-path and suppose there are decreasing returns to search and constant costs. Every selection from the search intensity correspondence $\mathcal{I}^h$ is non-decreasing on the region where search does not immediately conclude under some equilibrium.*

An intuition for Proposition 4 is that conditional on some fixed number of searches being profitable at complexity levels $L' < L''$, additional search is weakly more profitable when complexity is higher if any search is profitable at all (by the decreasing returns to search

[20] There is a difference between anticipated and realized search effort in a sequential search setting. While the searcher plans for the worst-case, she may shorten on lengthen her search if she makes unanticipated discoveries along the way. Callander (2011) describes this property as the *law of unintended consequences*.

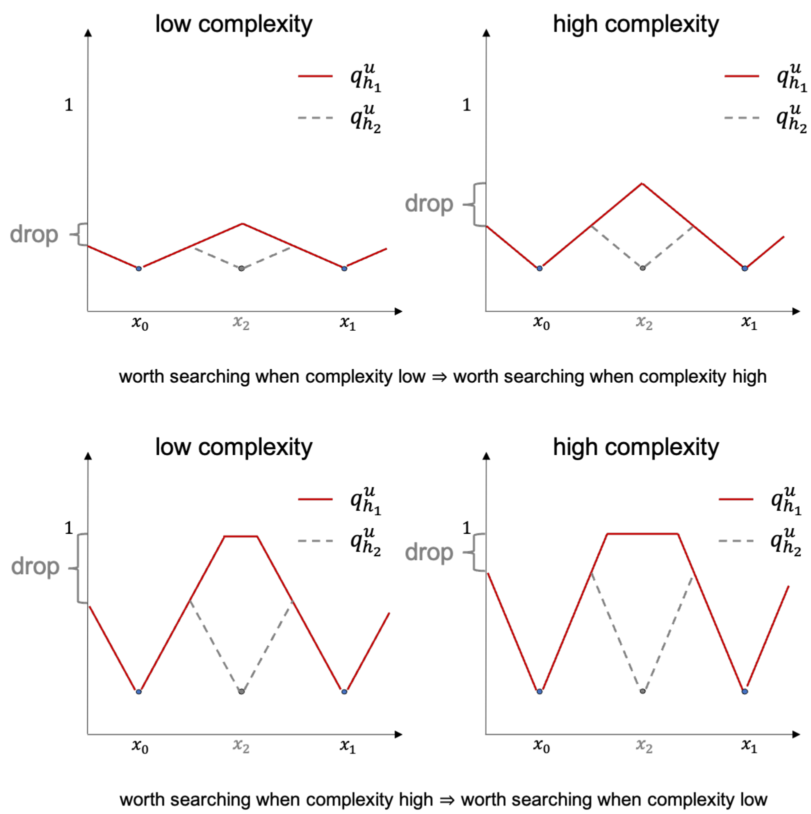

Figure 5: The proof of Proposition 3 breaks into two cases. Let $q^u_{h,L}$ be the upper envelope of feasible quality indices at history h when complexity is L. Let L^* be the threshold search complexity where $\max q^u_{h,L}(x) < 1$ if $L < L^*$ (top) and $\max q^u_{h,L}(x) = 1$ if $L > L^*$ (bottom). The dominance argument shown in each case implies that the set of complexities where search does not immediately conclude is an interval about L^*.

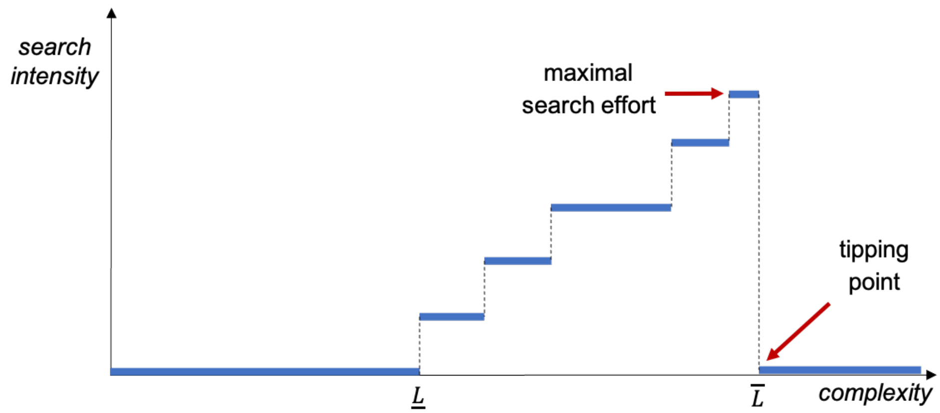

Figure 6: An example of the intensity correspondence when the assumptions of Proposition 3 and Proposition 4 are met.

assumption). In the case of constant search costs, the greater willingness to spend on search also translates into more search.

Taken together, Proposition 3 and Proposition 4 imply that search intensity is 'inverse-v-shaped' in complexity (see Figure 6). When complexity is low, there is no search. As complexity increases, search intensity initially ramps up. But beyond a tipping point where maximal search occurs, search intensity collapses back to zero.

The model predicts that small differences in the perceived difficulty of exploring an unfamiliar territory may determine whether the searcher exerts maximal effort or gives up on exploration entirely. This suggests, for example, that there may be large impacts to subsidizing search effort when agents are paralyzed by the perceived difficulty making good discoveries.

6 News and search dynamics

By Proposition 1, the searcher explores as if expecting to make neither good nor bad discoveries (within the search window) along the equilibrium path. The interpretation is not literal: the adversary is only a useful metaphor to characterize the behavior of a searcher who maximizes her payoff guarantee. Indeed in our motivating examples, there is no force driving quality indices to be adversarially chosen. It is therefore of interest to explore how the searcher's behavior changes when 'off-path' good or bad news events occur.

Fix the true quality index q and a search strategy. Consider a partial history $h_{t+1} \in H$ where the searcher does not conclude search, and let $h_{t+1}(t)$ be its time t sub-history. The searcher learns *good news* at h_{t+1} if $z_{t+1} > z^*_{h_{t+1}(t)}$. Similarly, the searcher learns *bad news* at h_{t+1} if $z_{t+1} < z^*_{h_{t+1}(t)}$.

6.1 Search step size

The next proposition describes the impact of sufficiently bad news on step size in an optimal search procedure. It captures the intuitive idea that bad news should drive the searcher to avoid the surrounding area and search elsewhere.

We say costs are *history-independent* if $C(x, \{(x_i, z_i)\}_{i=0}^t) = f(x)$ for some continuous $f : S \to \mathbb{R}_+$, bounded below by some $c > 0$

PROPOSITION **5.** *Suppose costs are history-independent. Let* σ *be an optimal sequential search procedure and consider any history* $h_t \equiv \{(x_i, z_i)\}_{i=0}^t \in H$ *at which* $\sigma(h_t) \neq \emptyset$ *and* $\min_{q \in Q_{h_t}} q(\sigma(h_t)) > 0$*. Let* x *be the closest technology to* $\sigma(h_t)$ *among those that had been discovered so far. Then there exists* $z' < z^*_{h_t}$ *such that:*

1. $Q_{h'_{t+1}}$ *is non-empty, where* $h'_{t+1} \equiv \{(x_0, z_0), \ldots, (x_t, z_t), (\sigma(h_t), z')\}$.
2. *If* $z_{t+1} \leq z'$*, then either* $\sigma(h_{t+1}) = \emptyset$ *or* $|\sigma(h_{t+1}) - \sigma(h_t)| > |\sigma(h_t) - x|$.

The second part of Proposition 5 says that after sufficiently bad news, an optimal search procedure either concludes or jumps beyond the nearest technology to another part of the search space. The first part of Proposition 5 says that it is possible to hear such sufficiently bad news at any history satisfying the stated assumptions.

With bad news, the best technology discovered so far remains unchanged, but not so with good news. The searcher's reaction to good news varies from case to case, and there are many cases to consider.

For example, the searcher may care less about being near the top as the quality of her best discovery rises. In that case, she may reduce search intensity in response to good news and make larger jumps in search space.

Alternatively, the searcher may care more about being near the top as the quality of her best discovery rises. For example, if a firm's search for an efficient flying car technology produces poor outcomes, the firm might come to believe the market for such cars will never be large. But if the firm chances upon a good design, it learns that there could be a large market and the winner of the innovation race would capture a disproportionate share. Such a searcher may ramp up search intensity in reaction to good news and start by testing designs close to her best discovery.

Even these simple intuitions do not easily generalize. In addition to the submodularity or supermodularity of U, the searcher's response to good news depends on the history of discoveries, the cost function and the equilibrium in question.

6.2 Funnel shaped search

The impact of good or bad news on search dynamics is more striking when the quality index is known to be quasiconvex. The searcher can use *any* good or bad news to significantly narrow down the space over which future searches will occur.

Starting at some history $h_t \in H$, a searcher *never searches in some set* $T \subset S$ *under strategy* σ if for any $h \in H$ such that $h(t) = h_t$, either $\sigma(h) \in S/T$ or $\sigma(h) = \emptyset$.

PROPOSITION **6.** *Let* $Q = Q^{QC}$ *and suppose costs are history-independent. Suppose the searcher uses an optimal sequential search procedure* σ*, and let* $h_{t+1} = \{(x_i, z_i)\}_{i=0}^{t} \in H$ *be a history at which* $\sigma(h_{t+1}) \equiv x_{t+1} \in S$*.*

1. *Suppose the searcher learns good news at* x_{t+1}*, i.e.,* $z_{t+1} > z^*_{h_t}$*.*
   (a) *If* $x_{t+1} > x_t$*, the searcher will never search left of* $x_t + \frac{1}{L}(z_{t+1} - z^*_{h_t})$*.*
   (b) *If* $x_{t+1} < x_t$*, the searcher will never search right of* $x_t - \frac{1}{L}(z_{t+1} - z^*_{h_t})$*.*
2. *Suppose the searcher learns bad news at* x_{t+1}*, i.e.,* $z_{t+1} < z^*_{h_t}$*. Let* $x^*_{h_t} \in X^*_{h_t}$*.*
   (a) *If* $x_{t+1} > x^*_{h_t}$*, the searcher will never search right of* $\sigma(h_t) - \frac{1}{L}(z^*_{h_t} - z_{t+1})$*.*
   (b) *If* $x_{t+1} < x^*_{h_t}$*, the searcher will never search left of* $\sigma(h_t) + \frac{1}{L}(z^*_{h_t} - z_{t+1})$*.*

Good and bad news events cause search to unfold in a 'funnel shape': the walls close in on the space of products or technologies over which search continues. This pattern of first searching broadly in attribute space and then narrowing in on a particular region was observed, for example, by Bronnenberg et al. (2016) in ordered-search data from online shoppers. In our model, the existence of a 'sweet-spot' generates such dynamics.

These dynamics also resemble the 'triangulation phase' that Callander (2011) finds in a model where, effectively, a myopic agent tries to find the zeroes of a Brownian motion path. As Hörner and Skrzypacz (2017) note, understanding how dynamics change in "Callander's model for patient agents is an important open problem". The quasiconcave case is an analogous prior-free model of a patient agent, and Proposition 6 illustrates the robustness of the triangulation pattern in Callander (2011). However, patience changes where searchers explore and when they stop relative to myopic searchers—the subject of Appendix D.

7 Extensions

The base model of search considers a single dimensional search space in which the searcher knows all available items. In this section, we relax both assumptions by allowing for limited consideration sets and multidimensional search.

7.1 Limited consideration sets

Suppose the searcher perceives S to be the set of items that exist. However, at a history h_{t+1}, the searcher can only search items that lie in a *consideration set* $S_{h_t} \subset S$. Moreover, the searcher's cannot shrink over time. That is, if $h_{t+1}(t)$ is a time t sub-history of h_{t+1}, then $S_{h_{t+1}(t)} \subset S_{h_{t+1}}$.

The notion of a consideration set captures a searcher who, over time, discovers the existence of more items that can be searched (and does not forget previously discovered items). An example is a shopper who learns about the existence of new products while searching as a platform recommends related items.

We generalize the model to allow for limited consideration by supposing the searcher employs a strategy that is robust to the way her consideration set evolves. That is, the searcher behaves as if she faces an adversary who can choose not only a feasible quality index, but also her consideration set at every history.

More formally, a feasible partial history is now a sequence of tuples $h_t = \{(x_i, z_i, S_{i+1})\}_{i=0}^{t-1}$ with the property that $S_i \subset S_{i+1} \subset S$ and $x_i \in S_i$ for all i. Each period proceeds in three stages. First, at h_t, the searcher either concludes search or searches an unexplored item $x_t \in S_t$. Next the adversary chooses a feasible quality index $q' \in Q_{h_t}$ and a consideration set $S_{t+1} \supset S_t$. Finally, if the searcher had concluded search, she anticipates a payoff of $p(h_t, q')$. Otherwise, $q'(x_t)$ and the new consideration set S_{t+1} are revealed to the searcher.

The agent's payoffs and the adversary's payoffs are the same as in the baseline model. In particular, upon concluding search at a history h_t, the agent still worries about the quality of the best undiscovered item in S, rather than in her limited consideration set. S_t is simply a constraint on what the searcher is able to explore, but there may be items in S/S_{t+1} that are unreachable to her but still act as reference points.

It is straightforward to see that Theorem 1 continues to hold for this generalized model with consideration sets: optimal sequential search procedures are those that follow an optimal simultaneous search procedure at every history. In particular, the searcher anticipates that she will not learn the existence of any new items.

To see why, suppose at some history, the adversary deviates to showing the searcher new items. If none of these items are part of the searcher's optimal simultaneous search strategy, then this one-shot deviation does not help. If they are part of the searcher's optimal search strategy, they the adversary is weakly worse off, since the searcher chose not implement the (still feasible) strategy that she would have before learning the existence of the new items.[21] We omit a full proof as it follows similar steps to the proof of Theorem 1.

[21] It is also straightforward to argue that double deviations, where the adversary gives good or bad news and introduces new items into the consideration set would only help the searcher.

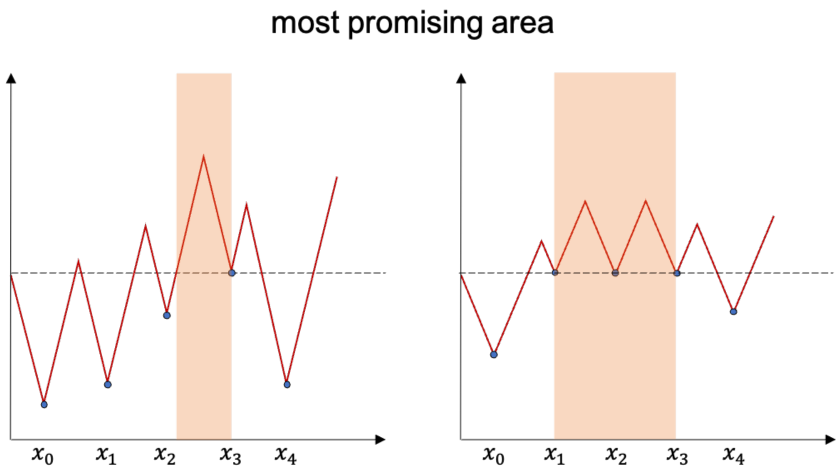

Figure 7: For different histories and Lipschitz constants, both figures plot upper-envelopes of feasible quality indices and highlight the most promising area. This is the set of unexplored items, which if explored, would reduce the value of one of tallest peaks in equilibrium.

We can use the observation that the consideration set does not expand along the equilibrium path to give a necessary condition for search to continue.

Let $A(h) \equiv \arg\max_{x \in S} q_h^u$ for any $h \in H$. Let h_x denote the history immediately after h where item $x \in S$ is searched and found to have quality z_h^*. Define *the most promising area* to be the subset P of unexplored items in S such that if $x \in P$ is explored at h, then either $\max_{x \in S} q_{h_x}^u < \max_{x \in S} q_h^u$ or $A(h_x)$ is a strict subset of $A(h)$; see Figure 7.

The following proposition says that if the searcher's consideration set does not include items in the most promising area, then she will conclude search.

Proposition 7. *Let $h_t = \{(x_i, z_i, S_{i+1})\}_{i=0}^{t-1}$ and let P be the the most promising area of S at h_t. If $S_t \cap P = \phi$, the searcher concludes search at h_t.*

The proof is simple. If the searcher cannot explore the region where the best alternative lies, then she anticipates never being able to explore in that region, so search does not help her in the worst case.

Proposition 7 has implications for how online platforms might make product recommendations for first-time buyers searching unfamiliar territory. Recommendations affect search by altering a shopper's consideration sets. A platform selling identically priced goods would want to recommend products that induce more search to avoid the shopper leaving before making a purchase (see Footnote 7). Proposition 7 suggests that the platform should not initially display a narrow set of products: repeated searches in the same area would not reduce the scope for missing good discoveries. On the other hand, displaying a diverse offering assures the searcher that she can form worthwhile search plans.

7.2 Multidimensional search

Technologies and products typically have multiple observable attributes. For example, online marketplaces may make the resolution, size and price of TVs readily observable before shoppers have to click on the image of the product to learn more.

In some settings, searchers may be able to discern how attributes contribute to the quality of an item (e.g., learning one's tastes for thread-count and material used in bedding). In other settings, they might only be able to learn the quality of an item as a whole (e.g., noticing only that a certain pillow is or is not comfortable).

This suggests two natural ways of extending the single-dimensional model to multidimensional search. Appendix E considers a different model where the searcher learns her preferences for each attribute. Here, we consider a more straightforward extension where the searcher only learns her value of an item as a whole.

The search space, S, is a compact subset of $\mathbb{R}^k$ in the *multidimensional search model*, which is otherwise identical to the base model.

THEOREM **1′.** *Suppose* $Q = Q^{MP}$ *and* S *is a compact subset of* $\mathbb{R}^n$*. Then a search strategy* σ *is an optimal sequential search procedure if and only if it follows an optimal simultaneous search procedure at every history.*

This result needs no separate proof, as the proof of Theorem 1 in Appendix F never makes use of the one-dimensionality of S.

As a consequence, comparative statics with respect to search complexity (Proposition 2, Proposition 3, and Proposition 4) and news (Proposition 5) generalize to the multidimensional case as well, when $Q = Q^{MP}$. Appendix B describes an optimal search algorithm for the case of constant costs and a finite search space.

On the other hand, there is no simple analogue to *triangulation* when the search space is multidimensional. Moreover, the characterization of optimal search when $Q = Q^{QC}$ does not readily generalize; Appendix F.2 highlights some difficulties that arise.

8 Conclusion

We study ordered-search with learning in unfamiliar territories. The model tries to capture the limited information that searchers work with when trying to innovate or learn their tastes for new products. In such settings, we characterize optimal forward-looking and even multidimensional search.

We close by speculating about two potential applications.

First, a problem for online platforms is to order products so as to help shoppers discover what they want. Blake et al. (2016), for example, find that "[where] a search engine like the one at eBay is tuned to encourage immediate purchase, the site might be better served if it thought holistically about this search funnel and helped consumers learn about the attributes of products in a way that ultimately led them step by step down the process instead of assuming that they are at the end of it". Addressing how to make product offerings easier to navigate requires a theory of how shoppers navigate. Our model might help address such design questions.

Second, a literature in economics and marketing estimates preferences and search costs from consumer search data. In a review article, Honka et al. (2019) note a need for theories that better match empirical search patterns and can rationalize search sequences. The resemblance to the funnel shaped search dynamics observed in Bronnenberg et al. (2016) and the computational tractability of our model suggest that it may be useful for structural estimation.

References

Adam, Klaus, "Learning while searching for the best alternative," *Journal of Economic Theory*, 2001, *101* (1), 252–280.

Banerjee, Abhijit V, Sylvain Chassang, and Erik Snowberg, "Decision theoretic approaches to experiment design and external validity," in "Handbook of Economic Field Experiments," Vol. 1, Elsevier, 2017, pp. 141–174.

Bikhchandani, Sushil and Sunil Sharma, "Optimal search with learning," *Journal of Economic Dynamics and Control*, 1996, *20* (1-3), 333–359.

Blake, Thomas, Chris Nosko, and Steven Tadelis, "Returns to consumer search: Evidence from ebay," in "Proceedings of the 2016 ACM Conference on Economics and Computation" 2016, pp. 531–545.

Bronnenberg, Bart J, Jun B Kim, and Carl F Mela, "Zooming in on choice: How do consumers search for cameras online?," *Marketing science*, 2016, *35* (5), 693–712.

Callander, Steven, "Searching and learning by trial and error," *American Economic Review*, 2011, *101* (6), 2277–2308.

Camara, Modibo K, "Computationally Tractable Choice," Technical Report 2021.

Carroll, Gabriel, "Robustness in mechanism design and contracting," *Annual Review of Economics*, 2019, *11*, 139–166.

Doval, Laura, "Whether or not to open Pandora's box," *Journal of Economic Theory*, 2018, *175*, 127–158.

Garfagnini, Umberto and Bruno Strulovici, "Social experimentation with interdependent and expanding technologies," *Review of Economic Studies*, 2016, *83* (4), 1579–1613.

Hansen, Pierre, Brigitte Jaumard, and Shi-Hui Lu, "Global optimization of univariate Lipschitz functions: I. Survey and properties," *Mathematical programming*, 1992, *55* (1), 251–272.

Hodgson, Charles and Gregory Lewis, "You can lead a horse to water: Spatial learning and path dependence in consumer search," 2020.

Honka, Elisabeth, Ali Hortaçsu, and Matthijs Wildenbeest, "Empirical search and consideration sets," in "Handbook of the Economics of Marketing," Vol. 1, Elsevier, 2019, pp. 193–257.

Hörner, Johannes and Andrzej Skrzypacz, "Learning, experimentation and information design," 2017.

Kreps, David M, "Decision problems with expected utility critera, i: upper and lower convergent utility," *Mathematics of Operations Research*, 1977, *2* (1), 45–53.

Manski, Charles F, "Choosing treatment policies under ambiguity," *Annu. Rev. Econ.*, 2011, *3* (1), 25–49.

McCall, John Joseph, "Economics of information and job search," *The Quarterly Journal of Economics*, 1970, pp. 113–126.

Rothschild, Michael, "Searching for the Lowest Price When the Distribution of Prices Is Unknown," *Journal of Political Economy*, 1974, *82* (4), 689–711.

Salant, Yuval, "Procedural analysis of choice rules with applications to bounded rationality," *American Economic Review*, 2011, *101* (2), 724–48.

Schlag, Karl H and Andriy Zapechelnyuk, "Robust sequential search," *Theoretical Economics*, 2021, *16* (4), 1431–1470.

Slivkins, Aleksandrs, "Introduction to multi-armed bandits," *arXiv preprint arXiv:1904.07272*, 2019.

Stigler, George J, "The economics of information," *Journal of political economy*, 1961, *69* (3), 213–225.

Urgun, Can and Leeat Yariv, "Retrospective search: Exploration and ambition on uncharted terrain," Technical Report, National Bureau of Economic Research 2021.

Weitzman, Martin L, "Optimal search for the best alternative," *Econometrica: Journal of the Econometric Society*, 1979, pp. 641–654.

Wong, Yu Fu, "Spatial experimentation," *Available at SSRN 3611751*, 2021.

A Dynamic Consistency

A.1 Adaptive adversary

Consider an alternative game where (1) the searcher can commit a fully history-contingent strategy, and (2) the adversary chooses a consistent quality index after observing the searcher's action at every history, just as in the sequential search game. We call this *the commitment game with an adaptive adversary.*

Proposition **8.** *Any optimal sequential search procedure in the sequential game is an optimal strategy for the searcher in the commitment game with an adaptive adversary.*

The proof is in Appendix F.5. The idea is that the adversary could always play the same 'no news' strategy it does in the sequential search game (see Proposition 1). The adversary should not receive a higher payoff from playing 'no news' than she would in the simultaneous search game, where the searcher is weaker (i.e., where she can only commit to simultaneous search strategies). But that means a searcher optimal strategy in the commitment game must best respond to 'no news'. Since the adversary cannot be better off than it is in the simultaneous search game, its payoff to 'no news' is the highest it can hope for. But then, the searcher is no worse off committing to an optimal simultaneous search procedure. By Theorem 1, she is is also no worse off committing to an optimal sequential search procedure.

A.2 Static adversary

Consider an alternative game where (1) the searcher can commit a fully history-contingent pure strategy, and (2) the adversary subsequently moves once to choose a consistent quality index. We call this the *pure strategy commitment game with a static adversary.*

Proposition **9.** *Any deterministic optimal sequential search procedure in the sequential game is an optimal strategy for the searcher in the pure strategy commitment game with a static adversary.*

The proof idea is that static adversary effectively as strong as the adaptive one in Proposition 8. It can compute the path of play under q^A to predict which items would be explored and make only the unexplored items high quality.

The restriction of the searcher to pure strategies is the natural commitment analogue of the sequential search game, where the adversary observes the searcher's action at every history. An interpretation of that timing is that the searcher cannot commit to randomizing where to look next. She may roll the dice to determine which item to pick. But upon seeing the realized roll, she starts to worry that the item she is about to select is the wrong choice.

Symmetrically, the searcher in the pure strategy commitment game with a static adversary can commit to a contingent plan of search but cannot commit to randomization.

On the other hand, a commitment game where the searcher can choose randomized strategies corresponds to an alternative sequential search model in which the searcher and adversary simultaneously choose actions at every history.

B Optimal sequential search algorithms

We describe algorithms returning optimal sequential search procedures when search costs are constant and fixed at $c > 0$. Appendix B.1 describes an algorithm for search over a finite set set of items, allowing for the search space to be multidimensional. Appendix B.2 more formally describes the algorithm discussed in Section 4.4 for search over a single-dimensional, connected search space.

B.1 Finite (multidimensional) search spaces

Suppose there are $n \in \mathbb{N}$ objects that can be searched, i.e., $|S| = n$. Let h_t be some history where $t < n$ objects have been discovered already. Let W_t be the search window at this history. Let $\overline{k}_t > 0$ be the smallest natural number such that if the searcher searches more than $\overline{k}_t$ times after history h_t, her payoff is guaranteed to be less than her payoff to ending search immediately, regardless of the search outcome. Note that $\overline{k}_t$ is independent of n.

The following algorithm returns an action at h_t that is part of an optimal sequential search procedure.

1. For each x in S and y in W_t, let $D(x, y) \equiv \max\{d(s, y)|s \in Conv(\{x, y\}),\ q^u_{h_t}(s) \geq z^*_{h_t}\}$.

2. For each $k \in \{0, 1, \ldots, \overline{k}_t\}$, compute B_i: the subset of W_t of size k that solves:

$$\min_{B \subset W_t, |B|=k} \max_{y \in S/B \cup X_{h_t}} D(y, B \cup X_{h_t}),$$

   where $B_0 = \emptyset$, and where for each y in W_t and $S' \subset S$, $D(y, S') = \min_{x \in S'} D(x, y)$.

3. Let B^* be the set $B_k \in \{B_1, \ldots, B_{\overline{k}_t}\}$ that maximizes the searcher's payoff, were he to follow a simultaneous search strategy that searches all items in B_k in some order and the adversary plays the 'no news' strategy.

4. The searcher picks some item in B^* to search if it is non-empty, and concludes search otherwise.

Intuitively, the algorithm picks items so that each undiscovered item is as close as possible to some discovered item. That way, by the bounds imposed by Lipschitz continuity, no undiscovered item can be that much better than a discovered item.

More precisely, in the third step, the searcher finds subsets of items (of different cardinalities) that, if searched, would be closest to the remaining undiscovered items. In the fourth step, by the multidimensional analogue of Theorem 1, she picks the subset that maximizes her payoff along the equilibrium path. Finally, the searcher picks one item in this subset and repeats the procedure until search concludes.

The first step and third step adjust for the right notion of distance. In the case where there is no item or only one item had been discovered so far, $D(y, S')$ is the usual Hausdorff distance between item y and a set S'. If item x had been discovered and has quality below $z^*_{h_t}$, then we measure the distance between y and the closest point between y and x (not necessarily in S) at which quality could conceivably exceed $z^*_{h_t}$.

Complexity analysis The first two step requires $O(n^2)$ computations (computing the modified notion of distance for, at most, every pair of items). The third and fourth steps collectively require $O(\overline{k}_t)$ computations. The second step requires $O(\binom{n}{1} + \ldots \binom{n}{\overline{k}_t}) \times O(n^2) = O(n^{\overline{k}+2})$ computations. She repeats these operations at most $\overline{k}_t$ times (with fewer computations each time, as $\overline{k}_{t+k} < \overline{k}_t$), giving an overall time complexity of $O(\overline{k}_t \cdot n^{\overline{k}_t+2})$.

The algorithm is therefore polynomial in n, e.g., the number items, and exponential in $\overline{k}_t$. The searchers utility parameters and search costs, U and c, determine $\overline{k}_t$. If c is large, or if U_2 is close to zero, then $\overline{k}_t$ is small. We are typically not interested in the limit as $\overline{k}_t$ grows large. On the other hand, we may be interested in solving problems where the number of items grows large, e.g., an online market place adds to the selection of cameras a shopper may choose from. This is a computationally tractable limit.

B.2 Connected (one-dimensional) search spaces

Without loss of generality, we take $S = [0, 1]$. Fix some history $h_t \in H$.

The boundary points of W_{h_t} and the technologies in $X^*_{h_t}$ partition W_{h_t} into sub-intervals. Let $mesh(W_{h_t}, X_{h_t})$ be the length of the longest of these sub-intervals. If $0 \in W_{h_t}/X_{h_t}$, let l_{h_t} denote the length of the sub-interval in W_{h_t} containing item 0 ; otherwise, let $l_{h_t} = 0$. Likewise, if $1 \in W_{h_t}/X_{h_t}$, let r_{h_t} denote the length of the sub-interval in W_{h_t} containing item 1; otherwise, let $r_{h_t} = 0$. Define the *weighted mesh at* h_t, denoted by $\omega(W_{h_t}, X^*_{h_t})$, as $\max\{mesh(W_{h_t}, X^*_{h_t}),\ 2l_{h_t},\ 2r_{h_t}\}$.

For any $k \in \mathbb{N}$, let $T(k, h_t)$ and denote a solution to the following problem:

$$\min_{T \in S/X_{h_t};|T|=k} \omega(W_{h_t}, T \cup X^*_{h_t}).$$

Label the elements of $T(k, h_t)$ as $x_{t+1}, \ldots, x_{t+k}$ in any order. Let

$$h_{t+k} = \{(x_0, z_0), \ldots, (x_t, z_t), (x_{t+1}, q^A_{h_t}(x_{t+1})), \ldots, (x_{t+k}, q^A_{h_t}(x_{t+k}))\},$$

and let $v(k, h_t) \equiv p(h_{t+k}, q^u_{h_{t+k}})$. Recall the definition of $\overline{k}_t$ from Section 4.4 (in particular, Footnote 15):

1. Let k^* be a solution to $\max_{k=0,\ldots,\overline{k}_t} v(k, h_t)$.

2. Let $\sigma(h_t) \in T(k^*, h_t)$.

Repeating these steps until $k^* = 0$ returns an optimal sequential search procedure, σ.

Complexity analysis Starting at a history where n items were searched, the search window can be found with $O(n)$ computations; this pre-processing does not affect the overall time complexity of the algorithm. Next, at most $\overline{k}_t$ searches occur after this history, regardless of the realization of the true quality index. Therefore, after l searches, the searcher's anticipated search intensity falls anywhere between 0 and $\overline{k}_t - l$, and the algorithm considers each of these levels. For any given anticipated search intensity, k, the placement of the optimal k searches can be computed in $O(n + k)$.[22] Therefore, the worst-case time complexity is $O(\overline{k}_t^3 + \overline{k}_t^2 \cdot n)$.

C Equilibrium refinements for history-independent costs

As the algorithms highlight, there are typically many optimal sequential search procedures when costs history-independent. This need not be the case for some history-dependent search costs, but the constant costs case is more tractable and relevant for some applications. Here we consider two natural equilibrium selection criteria.

C.1 Travel costs

Let $g : \mathbb{R}_+ \to \mathbb{R}_+$ be strictly increasing and strongly convex with $g(0) = 0$. Let $f : S \to \mathbb{R}_+$ be bounded away from 0. We say there are *travel costs* if $C(x, \{(x_i, z_i)\}_{i=0}^t) = f(x) + g(|x - x_t|)$.

Travel costs consist of a fixed component, which depends on the technology or product being searched, and a travel component, which depends on the distance between the current and the previously explored items in S. This can describe, for example, the form of search costs of an online shopper who is navigating some attribute-ordered list of products. Searching an item involves the fixed cost of clicking and reading about it. Jumping to a distant

[22] $\frac{n}{2} + 2$ is the maximum number of disjoint intervals in the search window. Sorting these by interval length takes $O(n)$ computations. Then the searches are placed so that $Y_{h_{t+1}}$ partitions $W_{h_{t+k}}$ as evenly as possible; this takes $O(k)$ operations.

item in price-space involves more scrolling or navigating to a new page. Alternatively, it may be easier to try similar designs when innovating than to attempt something very different.

Introducing travel costs reduces the multiplicity of optimal search procedures, as maintaining the direction of search and exploring nearer items first saves on such costs. We can use this idea while maintaining the tractability of the fixed cost model by considering the case of vanishing travel costs (e.g., reading about a product is much more expensive than scrolling).

Let $\{\alpha_n\}_{n=1}^{\infty} > 0$, and let $C_n(x, h_t) = f(x) + \alpha_n g(n|x - x_t|)$, where $\alpha_n g(n|x - x_t|) \to 0$ as $n \to \infty$. In effect, travel costs go to zero while the cost of the largest jump outstrips the other travel costs.

Suppose S is finite and let $h_t \in H$. At h_t, let T_{h_t} be the a set of items that the searcher would plans to explore before concluding search along some equilibrium path in some optimal sequential procedure when $C(x, h_t) = f(x)$. Optimal sequential search procedures converge to procedures where the searcher starts on a *hamiltonian cycle* through T_{h_t} that minimizes the maximum step size between all future searches.[23] Because hamiltonian cycles minimizing the max step size are generically unique (i.e., after slightly perturbing the locations of items in T_{h_t}), this selection criterion gives a sharp prediction for search dynamics in the fixed cost model.

This selection rule captures settings where travel costs are a secondary concern to the costs of conducting experiments themselves.

C.2 Search distractions

One can also consider a modification of the fixed cost model where search becomes prohibitively expensive with some small probability at every step. This can be interpreted as a distraction leading to a premature conclusion of search. As this probability vanishes, an optimal sequential search procedure with search distractions converges to an optimal sequential search procedure where at each history h_t, the searcher chooses a technology that leads to the biggest drop in $\max_{x \in s} q^u_{h_t}(x)$. Unlike the travel costs refinement, search distractions need not generically select for unique optimal search procedures.

D Myopic search

Here we characterize optimal myopic search procedure. Such strategies take the best action (e.g., either concluding search or searching some new technology) at each history under the

[23] Given a set of n points $T \subset R$ and an initial point, $t_0 \in T$, a *hamiltonian cycle* is a sequence $(t_0, t_1, \ldots, t_n)$ that includes each other point in T exactly once. Since travel costs are convex and grow steeper in n, the largest step size $|t_{i+1} - t_i|$ dominates the aggregate travel costs.

constraint that the searcher must conclude search by the next period at the latest.

Formally, let $M > 0$ be large enough such that if $x \in S$ is at least M, an optimal search procedure would conclude search immediately at any $h_t \in H$. For any $h_t \in H$, let $\sigma_{h_t}^{SL}$ denote a optimal sequential search procedure if costs where given by C^{SL}, where $C^{SL}(x, h_t) = C(x, h_t)$ for any new technology $x \in S$, and $C^{SL}(\cdot, h_{t+i}) = M$ for all $h_{t+i} \in H$ such that $h_{t+i}(t) = h_t$. A *myopic strategy* σ^M is a strategy that follows $\sigma_{h_t}^{SL}$ at every history $h_t \in H$.

D.1 Propensity to Search

A simple observation is that myopic searchers have a lower propensity to continue search than forward looking searchers. More formally, at any $h_t \in H$, if there exists an optimal sequential search procedure σ such that $\sigma(h_t) = \emptyset$, then there is a myopic strategy σ^M such that $\sigma^M(h_t) = \emptyset$. This follows by the optimality of σ and the definition of σ^M: if no myopic strategy concludes at h_t, the payoff to continuing search at h_t for one period and concluding search subsequently is strictly greater than concluding immediately, so the forward-looking searcher would not conclude either.

D.2 Search Location

The next proposition gives a general characterization of where myopic search occurs.

Proposition 10. *Let $Q = Q^{MP}$ or $Q = Q^{QC}$, and let $h_t \in H$. For any myopic strategy, σ^M, either $\sigma^M(h_t) = \emptyset$, or $\sigma^M(h_t)$ is in most promising area of S at h_t.*

Proof of Proposition 10. Search outside of the search window is wasteful. Searching a new technology inside the search window but outside of the most promising area results in a quality less than or equal to $z_{h_t}^*$ in the worst case. However, at this history, h_{t+1}, $\max_{x \in S} q_{h_t}^u(x) = \max_{x \in S} q_{h_t}^u(x)$. Therefore, concluding search at h_t would be an improvement for the searcher over reaching h_{t+1} and concluding search. Therefore, a myopic strategy that continues search at h_t searches inside the most promising area. □

Intuitively, any future search in the most promising area results in one of two possible outcomes: either the searcher learns that the highest possible quality is smaller than expected or she finds a better technology than previously discovered. Outside of the most promising area, finding technologies that are worse than what was previously discovered does not change the searcher's perception of what the highest possible quality could be. Searches in this region are of no value to a myopic searcher who does not expect to find something of higher quality than $z_{h_t}^*$.

Unlike myopic strategies, optimal sequential search procedures do not always search in the most promising area of S at every history.

For example, suppose the cost of experimenting with a set of promising technologies is prohibitively high for a new startup with limited resources. However, experimenting first with a more basic but less promising technology makes exploring the more promising technologies accessible. A forward-looking startup may start with the basic technology to unlock lower costs of future experimentation. A myopic startup that wishes to sell for the highest valuation after one period may only find it worthwhile to experiment with the most promising technologies directly.

D.3 Myopic search with constant fixed costs

In the case of constant fixed costs, one natural myopic strategy is simple to describe. Let $h_t \in H$. Let $x \in \arg\max q^u_{h_t}$, and let h_{t+1} be the history after h_t where x is searched and found to have quality $z^*_{h_t}$.[24] Define

$$\sigma^M(h_t) \equiv \begin{cases} \emptyset, \text{ if } p(h_t, q^u_{h_t}) > p(h_{t+1}, q^u_{h_{t+1}}) \\ x, \text{ otherwise} \end{cases}$$

It is clear that σ^M is a myopic strategy. In the event search does not conclude immediately, σ^M explores the technology that leads to the one-period greatest decrease in the quality of the best available alternative, on the equilibrium path (Proposition 1).

The technology explored by σ^M is farther away from Y_{h_t}, in terms of Hausdorff distance, than any other technology in the search window. At every history, a myopic searcher who uses σ^M ventures into a more unexplored part of the search space than any optimal sequential strategy.

D.4 Optimistic search and the constant fixed costs

As an aside, consider an optimistic searcher, i.e., one who believes, at every history $h_t \in H$, the best outcome in Q_{h_t} will obtain given her actions. At any $h_t \in H$ at which she does not conclude search, she explores an $x \in \arg\max q^u_{h_t}$, just as in σ^M (it is sub-optimal for her to search elsewhere). Therefore, she always explores a more unexplored part of the search space than the searcher in the base model.[25] It is also straightforward to see that at any history, it

[24] It is easy to verify that $Q_{h_{t+1}} \neq \emptyset$

[25] Outside of constant fixed costs, location choices for myopic and optimistic search need not coincide. Indeed, even in a more general fixed costs model, optimistic searchers may search outside of the most promising area entirely, if search in this region is significantly more expensive than search in other regions.

is optimal for the baseline searcher (and therefore, the myopic searcher) to conclude search if the optimistic searcher does so.

E Attribute learning model

Here, we define a multidimensional *attribute learning model*, where upon searching an item, the searcher learns the quality of each of its attributes.

Suppose items have k attributes, e.g., the size, resolution and brand of TVs. Let S_i be the one-dimensional sets of values that the ith attribute can take. The search space is $S \equiv S_1 \times \ldots \times S_k$.

An *attribute quality index* is a mapping $\kappa : S \to [0,1]^k$, where $\kappa \in (Q^{MP})^l \times (Q^{QC})^{k-l}$ for some $l \in \{0, 1, \ldots, k\}$.[26] Each time the searcher explores an item, she learns her value for each dimension of the object, e.g., how much she likes a large versus small TV, how much she values more resolution, etc.

The quality index aggregates the qualities of the separate dimensions: for any $x \in S$, $q(x) = f \circ \kappa(x)$, where $f : [0,1]^k \to [0,1]$ is increasing in all its arguments. We say *attributes are substitutes* if f is submodular. While κ is unknown to the searcher, f is known: the searcher knows the weight she gives to dimension of a item's quality but not the attribute qualities themselves.

S is a *rectangular search space* if for every $x, y \in S$, there exists some $z \in S$ such that $\kappa(z) = \kappa(x) \vee \kappa(y)$. In words, if the searcher likes some attributes of one item and some of another, she can find an item in search space that has all the better of all attributes.

To simplify the analysis, we assume that the searcher can conclude search by taking a previously unexplored option without paying an additional search cost. For example, if she considers a 65 inch 4k TV and a 75 inch 5k TV and learns that she likes the smaller size but a higher resolution, she can purchase a 65 inch 5k TV without incurring additional search costs.[27]

The attribute learning model is otherwise identical to the baseline model. We have the following analogue of Theorem 1.

Corollary 1. *Consider an attribute learning model where S is a rectangular search space and attributes are substitutes. Then a search strategy σ is an optimal sequential search procedure if and only if it follows an optimal simultaneous search procedure at every history.*

[26] Lipschitz constants may differ across attributes.

[27] In the base model, even if the searcher could take an option that she had not previously explored, she would never do so when the lower envelope of Q_h is greater than 0: the quality of an unexplored item is below z^* in the worst case. By comparison, allowing the searcher to take an unexplored option can significantly change the solution in Weitzman's model; see Doval (2018).

See Appendix F.6 for the proof.

Corollary 1 implies that previous comparative statics results and algorithms can be readily adapted to the attribute learning model. For example, triangulation in search space occurs along those dimensions of κ which the searcher perceives as being quasiconcave.

When $S = [0, 1]^k$, the same algorithm described in Section 4.4 can be applied to choose search location for each attribute. A similar procedure works when S is finite as well.

The assumptions that the search space is rectangular or that attributes are substitutes rule out some interesting applications. For example, it may be possible to purchase a powerful but bulky computer or a slower but more portable model. However, there may not exist a very fast and perfectly portable computer on the market. Similarly, a searcher may like a bubblegum pink convertible but may find the same color to be distasteful for a station wagon. Studying optimal search in the presence of attribute complementarities is an interesting problem for future work.

F Proofs

F.1 Proof of Theorem 1 in the $Q = Q^{MP}$ case

Let $h_t = \{(x_i, z_i)\}_{i=0}^{t-1} \in H$. For each $x_i \in X_{h_t}$ and $y \in S$, let $f_{h_t,x_i}(y) = L||y - x_i|| + z_i$. Let $f_{h_t} = \min_{x \in X_{h_t}} f_{h_t,x}$. Recall that $q^u_{h_t}$ is the upper envelope of the quality indices in Q_{h_t} for any history $h_t \in H$.

Lemma **1.** *Suppose that $Q = Q^{MP}$. Then $q^u_{h_t} \in Q_{h_t}$ for any $h_t \in H$. Moreover, $q^u_{h_t} = \min\{f_{h_t}, 1\}$.*

Proof of Lemma 1. We proceed by proving three claims.

Claim 1: $\min\{f_{h_t}, 1\}$ is L-Lipschitz.

Let $x, y \in S$. Then there exists some $x_i, x_j \in X_{h_t}$ such that $|f_{h_t}(x) - f_{h_t}(y)| = |f_{h_t,x_i}(x) - f_{h_t,x_j}(y)|$. Suppose without loss of generality that $f_{h_t,x_i}(x) \geq f_{h_t,x_j}(y)$. Then

$$\begin{aligned}
|f_{h_t}(x) - f_{h_t}(y)| &= f_{h_t,x_i}(x) - f_{h_t,x_j}(y) \\
&\leq f_{h_t,x_j}(x) - f_{h_t,x_j}(y) \\
&= L||x - x_j|| + z_j - L||y - x_j|| - z_j \\
&= L||x - x_j|| - L||y - x_j|| \\
&\leq L||x - y||,
\end{aligned}$$

where the first inequality follows from the definition of f_{h_t}. Therefore, f_{h_t}, and therefore $\min\{f_{h_t}, 1\}$, is L-Lipschitz.

Claim 2: Every quality index in Q_{h_t} is bounded above pointwise by $\min\{f_{h_t}, 1\}$.

Let $h_t = \{(x_i, z_i)\}_{i=0}^{t-1} \in H$, and let $q' \in Q_{h_t}$. For any $x_i \in T$ and $y \in S$, $|q'(y) - q'(x_i)| \leq L||y - x_i||$; Moreover, note that $|f_{h_t,x_i}(y) - f_{h_t,x_i}(x_i)| = f_{h_t,x_i}(y) - q'(x_i) = L||y - x_i||$. Together, this implies that $q'(y) \leq f_{h_t,x_i}(y)$ for all $x_i \in T$. Therefore, $q'(y) \leq \min\{f_{h_t}(y), 1\}$ for all $q' \in Q_{h_t}$.

Claim 3: $\min\{f_{h_t}, 1\}$ is consistent at h_t.

For any $x_i \in X_{h_t}$ and $q' \in Q_{h_t}$, $z_i = q'(x_i) \leq f_{h_t}(x_i) \leq f_{h_t,x_i}(x_i) = z_i$, where the first inequality follows from *Claim 2*, and the second is by the definition of f_{h_t}. Therefore $f_{h_t}(x_i) = \min\{f_{h_t}(x_i), 1\} = z_i$ for all $x_i \in X_{h_t}$, so $\min\{f_{h_t}, 1\}$ is consistent.

The first and third claims imply that $\min\{f_{h_t}, 1\} \in Q_{h_t}$. So by the second claim, f_{h_t} is the upper envelope of the quality indices in Q_{h_t}, i.e., $q^u_{h_t} = \min\{f_{h_t}, 1\}$. □

Let $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1}$ and $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1}$ be two partial histories that differ only by quality. Suppose $z'_i \geq z''_i$ for all i. Then we say h'_t *dominates* h''_t in *quality*.

Lemma **2.** *Suppose that $Q = Q^{MP}$. If h'_t dominates h''_t in quality, then*

$$\max_{q' \in Q_{h'_t}, x \in S} q'(x) \geq \max_{q' \in Q_{h''_t}, x \in S} q'(x).$$

Proof of Lemma 2. The statement of the lemma is equivalent to showing that:

$$\max_{x \in S} q^u_{h'_t} \geq \max_{x \in S} q^u_{h''_t}.$$

Let $T \equiv X_{h'_t} = X_{h''_t}$ be the set of searched technologies in $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1}$ and $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1}$. By Lemma 1, it suffices to show that $f_{h'_t,x_i}(y) = L||y - x_i|| + z'_i \geq L||y - x_i|| + z''_i = f_{h''_t,x_i}$ for every $x_i \in T$ and $y \in S$. This follows immediately from the assumption that $z'_i \geq z''_i$ for all i. □

Lemma **3.** *Suppose that $Q = Q^{MP}$. Fix some history $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1} \in H$ and let x be an optimal technology at h'_t. Let $T = X^*_{h'_t} \setminus \{x\}$. For each $x_i \in T$, let $\epsilon_i \in \mathbb{R}$. Consider an alternate history $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1} \in H$ where $z''_i = z'_i + \epsilon_i$ for $x_i \in T$ and $z''_i = z'_i$ otherwise. Suppose moreover that $Q_{h'_t}$ and $Q_{h''_t}$ are nonempty . The searcher's payoff to concluding search in h''_t is greater than or equal to the searcher's payoff to concluding search at h'_t.*

Proof of Lemma 3. If T is empty, the statement is trivially true. Suppose T is non-empty.

Case 1: If $\epsilon_i \leq 0$ for each $x_i \in T$, then note that the quality of the searcher's chosen technology is the same in both h''_t and h'_t, but by Lemma 2, the adversarially selected quality of the best technology in S is weakly higher in h'_t. Therefore, the searcher's benefit to concluding search is weakly greater at h''_t than at h'_t.

Case 2: Suppose there exists some $x_i \in T$ such that $\epsilon_i > 0$. Let ϵ denote the largest ϵ_i among all i such that $x_i \in T$. Consider a third partial history $h_t''' = \{(x_i, z_i' + \epsilon)\}_{i=0}^{t-1}$, i.e., where the quality of *all* technologies searched in h_t' is higher than in h_t' by ϵ. It is immediate by definition that $q^u_{h_t'''} \leq q^u_{h_t'} + \epsilon$, with equality if the upper bound of 1 does not bind, so

$$\max_{q' \in Q_{h_t'''}, x \in S} q'(x) \leq \epsilon + \max_{q' \in Q_{h_t'}, x \in S} q'(x).$$

But then the searcher's benefit to concluding search is weakly higher at h_t''' than at h_t', because by assumption, $U(a+\epsilon, b+\epsilon) \geq U(a, b)$ for any $a, b \in \mathbb{R}_+$.

The quality of the searcher's chosen technology is the same in both h_t'' and h_t'''. But by Lemma 2 again, the searcher's benefit to concluding search is weakly greater at h_t'' than at h_t''', and therefore weakly greater at h_t'' than at h_t'.

Recall C is independent of the qualities of searched technologies. Therefore, the searcher's expected costs of experimentation would be the same upon concluding search after partial h_t' or h_t''. Therefore, the searcher's payoff is weakly greater if she concludes search at h_t'' than at h_t'. □

Let $q^A_{h_t} = \min\{q^u_{h_t}, z^*_{h_t}\}$. For all $h_t \in H$, let $\sigma^A_d(h_t, x) = q^A_{h_t}$ for all $x \neq \emptyset$, and $\sigma^A_d(h_t, \emptyset) = q^u_{h_t}$.

LEMMA **4.** *Suppose that $Q = Q^{MP}$. Then σ^A_d is a feasible strategy for the adversary.*

Proof of Lemma 4. By Lemma 1, it suffices to prove that $q^A_{h_t} \in Q_{h_t}$. Clearly, $q^A_{h_t}$ is L-Lipschitz since it is the minimum of an L-Lipschitz function and a constant. By definition, it is also consistent. □

LEMMA **5.** *Suppose that $Q = Q^{MP}$ and let $\sigma \in \Gamma_{h_t}$. Then σ is an optimal simultaneous search procedure at h_t if and only if σ is a best-response to σ^A_d in the simultaneous search game at $h_t \in H$.*

Proof of Lemma 5. Let T denote the set of technologies explored under σ after h_t. Because $\sigma \in \Gamma_{h_t}$, T is independent of the adversary's strategy. Then by Lemma 3, σ^A_d weakly dominates any other strategy for the adversary.

Therefore, if σ is a best response to σ^A_d, then (σ, σ^A_d) is a Nash equilibrium of the simultaneous search game at h_t, i.e., σ is an optimal simultaneous search procedure.

For the other direction, suppose for contradiction that σ is not a best response to σ^A_d but that (σ, σ^A) is a Nash equilibrium of the simultaneous search game at h_t. Let σ' be a best response to σ^A_d, so that (σ', σ^A_d) is another equilibrium. Let $\preceq$ represent the searcher's preferences over outcomes.

First, $(\sigma', \sigma^A) \preceq (\sigma, \sigma^A)$, by the assumption that the latter is a Nash equilibrium.

Next, $(\sigma, \sigma^A) \preceq (\sigma, \sigma_d^A)$: the former is an equilibrium of the zero-sum game, so adversary is weakly worse off if she deviates, and the searcher is weakly better off.

However, it is also true that $(\sigma, \sigma_d^A) \preceq (\sigma, \sigma^A)$ because σ_d^A weakly dominates any other strategy for the adversary, leaving the searcher with a weakly smaller payoff.

Finally, $(\sigma, \sigma_d^A) \prec (\sigma', \sigma_d^A)$, since σ is not a best response to σ_d^A while σ' is a best response.

Putting this all together, $(\sigma', \sigma^A) \preceq (\sigma, \sigma^A) \sim (\sigma, \sigma_d^A) \prec (\sigma', \sigma_d^A)$, i.e., $(\sigma', \sigma^A) \prec (\sigma', \sigma_d^A)$. In other words, when the searcher plays σ', she is strictly better off when the adversary plays σ_d^A over σ^A. This implies the adversary is strictly worse off playing σ_d^A over σ^A when the searcher plays σ', contradicting the fact that σ_d^A weakly dominates any other strategy for the adversary.

Therefore if σ is not a best response to σ_d^A, it is not an optimal simultaneous search procedure. □

LEMMA **6.** *Suppose that $Q = Q^{MP}$ and that σ follows an optimal simultaneous search procedure at every history. Let $h_t \in H$, and let $\sigma' \in \Gamma_{h_t}$ be a simultaneous search procedure at h_t which replicates the searcher's path of play under (σ, σ_d^A), conditional on reaching h_t. Then σ' is an optimal simultaneous search procedure at h_t.*

Proof of Lemma 6. Fixing the adversary's strategy as σ_d^A, Lemma 5 implies that σ is a conserving strategy for the searcher, i.e., the highest payoff possible at h_t equals the highest payoff possible after the searcher takes the action $\sigma(h_t)$.

Next, let h be an infinite history where the searcher never concludes search, and consider the best payoff possible for the searcher at an alternate history that matches h for the first t periods: $\sup_{h' \in H, h'(t) = h(t)} \min_{q \in Q_{h'}} p(h', q)$. Because C is bounded away from zero, the searcher's best achievable payoff decreases without bound as she searches indefinitely, i.e., $\sup_{h' \in H, h'(t) = h(t)} \min_{q \in Q_{h'}} p(h', q) \to -\infty = p(h, \cdot)$ as $t \to \infty$.

Therefore conserving strategies are optimal in this setting (e.g., see Kreps (1977)), i.e., σ is an optimal strategy for the searcher when the adversary plays σ_d^A.

This implies that σ' is a best response to σ_d^A in the simultaneous search game, so by Lemma 5, σ' is an optimal simultaneous search procedure. □

LEMMA **7.** *Suppose that $Q = Q^{MP}$. If σ follows an optimal simultaneous search procedure at every history, then (σ, σ_d^A) is a sub-game perfect equilibrium of the sequential search game.*

Proof of Lemma 7. First we show that the searcher's strategy is unimprovable. Let $h_t \in H$. Let σ' denote some one-shot deviation from σ at h_t. Let σ'' and $\sigma''' \in \Gamma_{h_t}$ be simultaneous search procedures at h_t that replicate the searcher's path of play under (σ, σ_d^A) and (σ', σ_d^A), respectively.

Then σ' is a strict improvement over σ only if σ''' is a strict improvement over σ'' in the simultaneous search game at h_t. But by Lemma 6, σ'' is an optimal simultaneous search procedure and therefore a best response to σ_d^A in Γ_{h_t}. Therefore, σ''' is not a strict improvement over σ'', and so σ' is not a strict improvement over σ.

Next, we prove that σ_d^A is unimprovable. Fix the strategy of the searcher to be σ and suppose the searcher picks technology x_t at history h_t. Suppose that the adversary deviates at this history to $q' \neq q_{h_t}^A$ and returns to following σ^A thereafter. Denote this history as h'_{t+1} There are three cases to consider.

Case 1: $q'(x_t) = q_{h_t}^A(x_t)$. In this case, all future histories would proceed as if there was no deviation at all, so this is not a strict improvement for the adversary.

Case 2: $q'(x_t) > q_{h_t}^A(x_t)$. Suppose for a moment that the searcher behaves as if the quality of x_t was $q_{h_t}^A(x_t)$ and continues to follow at h'_{t+1} what was the optimal simultaneous search procedure at h_t, i.e., σ^*_{s,h_t}. The searcher would be better off upon concluding search, by Lemma 3, than she would have been had the adversary not deviated from σ^A at h_t. If the searcher plays σ and follows $\sigma^*_{s,h'_{t+1}}$ at h'_{t+1}, she is better off still. Therefore the searcher is weakly better off when the adversary makes this one-shot deviation.

Case 3: $q'(x_t) < q_{h_t}^A(x_t)$. As in the previous case, we first consider what would happen if the searcher followed σ^*_{s,h_t} at h'_{t+1} onward. Clearly, the quality of the best searched technology is the same whether the adversary makes this one-shot deviation or not. By Lemma 2, if the searcher follows σ^*_{s,h_t} at h'_{t+1} onward, she is weakly better off when the adversary deviates at h_t. This implies that even when she follows σ, she is weakly better off when the adversary deviates, i.e., the adversary is weakly worse off.

Since the adversary has no strictly profitable one-shot deviation, the strategies (σ, σ_d^A) constitute a sub-game perfect equilibrium. □

LEMMA **8.** *Suppose that* $Q = Q^{MP}$*. If* σ *is an optimal sequential search procedure then* σ *follows an optimal simultaneous search procedure at every history partial* $h_t \in H$*.*

Proof of Lemma 8. Since σ is an optimal sequential procedure, there exists some strategy of the adversary σ^A such that (σ, σ^A) is a sub-game perfect equilibrium. In particular, let $h_t \in H$; then (σ, σ^A) is a Nash equilibrium of the sub-game starting at h_t. Let σ' be some strategy that follows an optimal simultaneous search procedure at every history. Then (σ', σ_d^A) is also a Nash equilibrium of the sub-game starting at h_t by Lemma 7.

Note that because (σ', σ_d^A) is an equilibrium of the zero-sum sub-game at h_t, the adversary is weakly better off at (σ, σ_d^A) as any deviation leaves the searcher weakly worse-off. And over (σ, σ_d^A), the adversary is weakly better off at (σ, σ^A), since σ^A is a best response to σ by assumption. And again because (σ, σ^A) is an equilibrium, the adversary's payoff weakly improves under the strategy profile (σ', σ^A). And finally, the adversary's payoff weakly improves from (σ', σ^A) if she best responds to the searcher instead at (σ', σ_d^A).

This loop of weak inequalities implies that the adversary's payoffs are identical at all of these strategy profiles. In particular, this means that (σ, σ_d^A) is an equilibrium in the sub-game at h_t. Therefore a strategy $\sigma'' \in \Gamma_{h_t}$ that replicates the searcher's path of play under (σ, σ_d^A) is an optimal simultaneous search procedure by Lemma 6.

In other words, σ follows an optimal simultaneous search procedure at h_t. Since the choice of h_t was arbitrary, the result follows. □

Proof of Theorem 1 in the $Q = Q^{MP}$ case. Lemma 7 and Lemma 8 together give the result. □

F.2 Proof of Theorem 1 in the $Q = Q^{QC}$ case

The analogues of Lemma 1 and Lemma 2 no longer holds when $Q = Q^{QC}$.

To see this, consider the following counter-example: Let $S = [0, 4]$ and $L = 1$. Denote by h_3 the partial history where technologies $\{0, 2, 3, 4\}$ have been searched and all have quality equal to 0, i.e., $h_3 = \{(0, 0), (2, 0), (3, 0), (4, 0)\}$.

First, note that the upper envelope of Q_{h_3} is a saw-tooth shaped function and therefore not quasiconcave.

Next, note that the highest possible quality for some technology under some $q \in Q_{h_3}$ is equal to 1. This is uniquely achieved at:

$$q(x) = \begin{cases} x & 0 \leq x < 1 \\ 2 - x & 1 \leq x < 2 \\ 0 & 2 \leq x \leq 4. \end{cases}$$

Now consider the history $h_3' = \{(0, 0), (2, 0), (3, 0.5), (4, 0)\}$, which dominates h_3 in quality. Since every quality index in $Q_{h_3'}$ is quasiconcave, it must now be the case that $q'(1) = 0$ for every $q' \in Q_{h_3}$. The highest possible quality for some technology under some $q' \in Q_{h_3'}$ is equal to 0.75. This is achieved at:

$$q'(x) = \begin{cases} 0 & 0 \leq x < 2 \\ x - 2 & 2 \leq x < 2.75 \\ 3.5 - x & 2.75 \leq x < 3.5 \\ 0 & 3.5 \leq x \leq 4. \end{cases}$$

However, an analogue of Lemma 4 and a weaker form of Lemma 2 hold, which suffice for the proof of Theorem 1 when $Q = Q^{QC}$. To this end, let $q^u_{h_t}$ again denote the upper envelope of quality indices in Q_{h_t} for any $h_t \in H$, and let $q^A_{h_t} = \min\{q^u_{h_t}, z^*_{h_t}\}$.

LEMMA **9.** *Suppose that $Q = Q^{QC}$. Then for all $h_t \in H$, $q^A_{h_t} \in Q_{h_t}$*

Proof. The argument that $q^A_{h_t}$ is L-Lipschitz and consistent is exactly as in the proofs of Lemma 1 and Lemma 4. It only remains to be shown that $q^A_{h_t}$ is quasiconcave.

Let $h_t = \{x_i, z_i\}_{i=1}^{t-1} \in H$. If $q' \in Q_{h_t}$, then by quasiconcavity, q' is non-decreasing on $[0, \min X^*_{h_t})$ and non-increasing on $(\max X^*_{h_t}, 1]$. This implies $q^u_{h_t}$ is non-decreasing on $[0, \min X^*_{h_t})$ and non-increasing on $(\max X^*_{h_t}, 1]$. Therefore $q^A_{h_t}$ is quasiconcave. □

LEMMA **10.** *Suppose that $Q = Q^{QC}$. If $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1}$ dominates $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1}$ in quality, and $z^*_{h'_t} = z^*_{h''_t}$, then*

$$\max_{q' \in Q_{h'_t}, x \in S} q'(x) \geq \max_{q' \in Q_{h''_t}, x \in S} q'(x).$$

Proof of Lemma 10. First we argue that $q^A_{h'_t} \geq q^A_{h''_t}$.

Since h'_t dominates h''_t and $z^*_{h'_t} = z^*_{h''_t}$, it follows that $X^*_{h''_t} \subset X^*_{h'_t}$. Therefore, $q^A_{h'_t} \geq q^A_{h''_t}$ on $[\min X^*_{h'_t}, \max X^*_{h'_t}]$.

Next, it follows from Lemma 9 that $q^A_{h'_t}$ and $q^A_{h''_t}$ are non-decreasing on $[0, \min X^*_{h'_t})$. This implies that $\max\{q^A_{h'_t}, q^A_{h''_t}\}$ is also non-decreasing on this interval. Reasoning analogously, $\max\{q^A_{h'_t}, q^A_{h''_t}\}$ is non-increasing on $(\min X^*_{h'_t}, 0]$. Finally, observe that $\max\{q^A_{h'_t}, q^A_{h''_t}\}$ is consistent with what the searcher had observed so far at h'_t, i.e., $\max\{q^A_{h'_t}, q^A_{h''_t}\} \in Q_{h'_t}$. Therefore, $q^A_{h'_t} = \max\{q^A_{h'_t}, q^A_{h''_t}\} \geq q^A_{h''_t}$.

Now let $q'' \in Q_{h''_t}$. It suffices to show that there exists a $q' \in Q_{h'_t}$ such that $\max_{x \in S} q'(x) = \max_{x \in S} q''(x)$.

Define q' as follows: $q'(x) = q^A_{h'_t}(x)$ if $q''(x) \leq z^*_{h'_t}$ and $q'(x) = q''(x)$ otherwise. By the preceding, $q^A_{h'_t} \geq q^A_{h''_t} \geq \max\{q'', z^*_{h''_t}\}$. Moreover, because $\{x \in [0,1] | q''(x) > q^A_{h'_t}(x)\} \subset [\min X^*_{h'_t}, \max X^*_{h'_t}]$, it is clear that q' is quasiconcave and that $q' \in Q_{h'_t}$.

Since $\max_{x \in S} q'(x) = \max_{x \in S} q''(x)$, the result follows. □

Proof of Theorem 1 in the $Q = Q^{QC}$ case. The proofs of the analogous lemmas to those in Appendix F.1 are identical, with Lemma 10 in place of Lemma 2 whenever the latter is referenced. □

Proof of Proposition 1. See the proof of Lemma 5. The proof when $Q = Q^{QC}$ is identical. □

F.3 Proofs for Section 5

Proof of Proposition 2. First we construct $\underline{L}$. If $L = 0$ and Q_h is nonempty, then clearly there is no value in search, as Q_h is a singleton containing only a constant function. Recall that C is bounded away from 0, so suppose that $c > 0$ is such that $C(x, h') > c$ for all $x \in S$ and $h' \in H$. Let $\epsilon > 0$ be small enough so that $U(z_h^*, z_h^*) - U(z_h^*, z_h^* + \epsilon) < c$. By compactness of S, there exists $\underline{L} > 0$ small enough so that $q_h^u < z_h^* + \epsilon$ when search complexity L is such that $L \leq \underline{L}$ and Q_h is nonempty. For contradiction, suppose search does not conclude immediately when $L \leq \underline{L}$ and Q_h is nonempty: there is some history h' after h on the equilibrium path at which the searcher concludes search. Then by Theorem 1 or **??** and by Proposition 1, the searcher's anticipated payoff is

$$\begin{aligned} p(h', q_{h'}^u) &= U(z_h^*, \max_{x \in S} q_{h'}^u(x)) - \sum_{i=1}^{t} C(x_i, h'(i)) \\ &\leq U(z_h^*, z_h^*) - c \\ &< U(z_h^*, z_h^* + \epsilon) \\ &< p(h, \max_{x \in S} q_h^u(x)), \end{aligned}$$

a contradiction. Therefore concluding search immediately is optimal if $L \leq \underline{L}$.

Next we construct $\overline{L}$. Let $\delta = U(z_h^*, z_h^*) - U(z_h^*, 1)$. Let $n = \lceil \frac{\delta}{c} \rceil$. Let h' be any history on the equilibrium path at which there have been at least n searches, i.e., $t \geq n$. Then

$$\begin{aligned} p(h', q_{h'}^u) &= U(z_h^*, \max_{x \in S} q_{h'}^u(x)) - \sum_{i=1}^{t} C(x_i, h'(i)) \\ &\leq U(z_h^*, z_h^*) - n \cdot c \\ &< U(z_h^*, 1). \end{aligned}$$

Therefore, it suffices to construct $\overline{L}$ such that if $L \geq \overline{L}$, then after any history with n searches where the adversary plays σ_d^A, $\max_{x \in S} q_{h'}^u(x) = 1$. This would imply that the searcher is better off concluding search immediately.

Pick $n+1$ points $x_1', \ldots, x_{n+1}' \in S/X_h$; let T denote the union of these points and X_h and let d be the minimum over distances between two points in T. Let $\overline{L} \geq \frac{1 - z_h^*}{d}$ and suppose $L > \overline{L}$. Consider a history h' after h where some $x_1, \ldots, x_n \in S/X_h$ are searched in some order, with the quality of x_i being $q_h^A(x_i)$. There exists by construction a $q \in Q_{h'}$ such that $q(x_i) = q_h^A(x_i)$ for all i, but $q(x_j') = 1$ for some j. Therefore $\max_{x \in S} q_{h'}^u(x) = 1$ for any history after which there are n searches. ∎

When considering comparative statics with respect to different levels of search complexity, say L' and L'', we subscript variables by L' or L'' to denote which case we are considering.

LEMMA **11.** *Let $0 < L' < L''$ be two different levels of search complexity. Let $h = \{(x_i, z)\}_{i=0}^t \in H$ be on-path, and $\max_{x\in S} q^u_{h,L''}(x) < 1$. Let $y \in S/X_h$, and let $h' \in H$ be the on-path history immediately after at which y is searched. Then,*

$$\max_{x\in S} q^u_{h,L'}(x) - \max_{x\in S} q^u_{h',L'}(x) \leq \max_{x\in S} q^u_{h,L''}(x) - \max_{x\in S} q^u_{h',L''}(x).$$

Moreover, if $\max_{x\in S} q^u_{h,L'}(x) - \max_{x\in S} q^u_{h',L'}(x) > 0$, the preceding inequality is strict.

Proof. Note that because $Q_{h'} \subset Q_h$, $\max_{x\in S} q^u_{h,L''}(x) - \max_{x\in S} q^u_{h',L''}(x) \geq 0$. So, the result holds when $\max_{x\in S} q^u_{h,L'}(x) - \max_{x\in S} q^u_{h',L'}(x) = 0$.

Suppose that $\max_{x\in S} q^u_{h,L'}(x) - \max_{x\in S} q^u_{h',L'}(x) > 0$.

Let $\underline{x}, \overline{x} \in X_h$ be the closest previously searched technologies to the left and right of y (and $\min S$ or $\max S$, respectively, if there are no such technologies). Define $\underline{x_s}, \overline{x_s}' \in X_h$ similarly as the endpoints in $X_h \cup \{\min S, \max S\}$ of the sub-interval containing the second largest peak of q^u_h.

Let $f(L) \equiv \max_{x\in[\underline{x},\overline{x}]} q^u_{h,L}(x)$. Similarly, let $g(L) \equiv \max_{x\in[\underline{x_s},\overline{x_s}]} q^u_{h,L}(x)$.

It is readily verified (for example, by Lemma 1 and an analogous result for the $Q = Q^{QC}$ case) that $f(L) = z + D(\frac{\overline{x},\underline{x}}{2}) \cdot L$, where $D(a, b) \equiv \frac{b-a}{2}$ if $a, b \in X_h$, and $D(a, b) \equiv b - a$ otherwise. Similarly, let $g(L) = z + D(\frac{\overline{x_s},\underline{x_s}}{2}) \cdot L$.

For the remainder of the proof, we consider only the case where $\underline{x}, \overline{x}, \underline{x_s}, \overline{x_s} \in X_h$. We obtain the same conclusion when one or more of $\underline{x}, \overline{x}, \underline{x_s}, \overline{x_s}$ are not in X_h. There are three cases to consider.

Case 1: At history h', $\max_{x\in S} q^u_{h',L'}(x) = g(L')$ and $\max_{x\in S} q^u_{h',L''}(x) = g(L'')$.

Now $f(L) - g(L) = L \cdot \frac{\overline{x}-\underline{x}-\overline{x_s}+\underline{x_s}}{2} > 0$ is linear in L with a positive slope, which implies $f(L) - g(L)$ is strictly increasing in L. Therefore $f(L') - g(L') \leq f(L'') - g(L'')$, which is the desired result.

Case 2: At history h',

$$\max_{x\in S} q^u_{h',L'}(x) = \max_{x\in S\cap[\underline{x},y]} q^u_{h',L'}(x),$$

and

$$\max_{x\in S} q^u_{h',L''}(x) = \max_{x\in S\cap[\underline{x},y]} q^u_{h',L''}(x).$$

Let $\alpha = \frac{y-\underline{x}}{\overline{x}-\underline{x}}$. Then by the property of similar triangles,

$$\max_{x\in S} q^u_{h',L'}(x) = \alpha \max_{x\in S} q^u_{h,L'}(x),$$

and

$$\max_{x\in S} q^u_{h',L''}(x) = \alpha \max_{x\in S} q^u_{h,L''}(x).$$

Since $(1-\alpha) \max_{x\in S} q^u_{h',L'}(x) < (1-\alpha) \max_{x\in S} q^u_{h',L''}(x)$, the result follows.

Case 3: At history h',

$$\max_{x \in S} q^u_{h',L'}(x) = \max_{x \in S \cap [y, \overline{x}]} q^u_{h',L'}(x),$$

and

$$\max_{x \in S} q^u_{h',L''}(x) = \max_{x \in S \cap [y, \overline{x}]} q^u_{h',L''}(x).$$

The proof in this case is identical to case 2. □

LEMMA **12.** *Suppose there are decreasing returns to search. Let* $0 < L' < L''$, $h \in H$ *be on-path, and* $q^u_{h,L''} < 1$. *If search concludes immediately under some optimal sequential search procedure when search complexity is* L'', *then search concludes immediately under any optimal sequential search procedure when search complexity is* L'.

Proof of Lemma 12. Suppose for contradiction that there is an optimal sequential search procedure that does not conclude search immediately when complexity is L'. Let h' be the history at which search ends along the equilibrium path.

Now, $\max_{x \in S} q^u_{h,L'}(x) - \max_{x \in S} q^u_{h',L'}(x) > 0$, or else concluding search immediately at h would have been a strict improvement. But then by Lemma 11 and induction on the number of searches,

$$\max_{x \in S} q^u_{h,L'}(x) - \max_{x \in S} q^u_{h',L'}(x) < \max_{x \in S} q^u_{h,L''}(x) - \max_{x \in S} q^u_{h',L''}(x). \tag{1}$$

Next, it is obvious that

$$\max_{x \in S} q^u_{h,L''}(x) \geq \max_{x \in S} q^u_{h,L'}(x). \tag{2}$$

Finally, because both histories are on path,

$$z^*_h = z^*_{h'}. \tag{3}$$

Putting eq. (1), eq. (2) and eq. (3) together, along with the facts that $U_2 \leq 0$ and $U_{22} < 0$, we have

$$U(z^*_{h'}, \max_{x \in S} q^u_{h',L'}(x)) - U(z^*_h, \max_{x \in S} q^u_{h,L'}(x)) < U(z^*_{h'}, \max_{x \in S} q^u_{h',L''}(x)) - U(z^*_h, \max_{x \in S} q^u_{h,L''}(x)).$$

So when search complexity is L'', no optimal sequential search procedure concludes search at h. This is a contradiction.

Therefore, continuing search is not a part of any optimal sequential search procedure when complexity is L'. □

LEMMA **13.** *Let* $0 < L' < L''$, $h \in H$, *and* $q^u_{h,L'} = 1$. *If search concludes immediately under some optimal sequential search procedure when search complexity is* L', *then search concludes immediately under any optimal sequential search procedure when search complexity is* L''.[28]

[28]We need not assume h is on path for this result, as Lemma 11 is not invoked, nor do we need to assume that there are decreasing returns to search.

Proof of Lemma 13. Suppose for contradiction that there is an optimal sequential search procedure that does not conclude search immediately when complexity is L''. Let h' be the history at which search ends along the equilibrium path.

Now, $1 - \max_{x \in S} q^u_{h',L''}(x) > 0$, or else concluding search immediately at h would have been a strict improvement. But then

$$\max_{x \in S} q^u_{h',L''}(x) > \max_{x \in S} q^u_{h',L'}(x)$$

Along with the facts that $U_2 \leq 0$, and $z^*_h = z^*_{h'}$, we have

$$U(z^*_{h'}, \max_{x \in S} q^u_{h',L'}(x)) - U(z^*_h, 1) > U(z^*_{h'}, \max_{x \in S} q^u_{h',L''}(x)) - U(z^*_h, 1).$$

So when search complexity is L', no optimal sequential search procedure concludes search at h. This is a contradiction. □

Proof of Proposition 3. Let h be some on-path history.

If search concludes immediately at h for any L (e.g., if $z^* = 1$ at h), then the result holds for any $\underline{L} = \overline{L} \in \mathbb{R}_{++}$.

Suppose there is an L at which search does not conclude immediately in some equilibrium. Let L^τ be such that when $L < L^\tau$, $\max_{x \in S} q^u_{h,L}(x)) < 1$ and when $L > L^\tau$, $\max_{x \in S} q^u_{h,L}(x)) = 1$.

Let $\mathcal{L}$ be the set of complexity levels L such that, in some equilibrium, search does not conclude when complexity is L.

Let $\underline{L} = \inf \mathcal{L}$. By Proposition 2, $\underline{L} > 0$. Moreover, by Lemma 13, $\underline{L} \leq L^\tau$. Finally, by Lemma 12, if $\underline{L} < L^\tau$, then search continues in all any equilibrium for $L \in (\underline{L}, L^\tau]$.

Similarly, let $\overline{L} = \sup \mathcal{L}$. By Lemma 12, $\overline{L} \geq L^\tau$. By Lemma 13, if $\overline{L} > L^\tau$, then search continues in all any equilibrium for $L \in [L^\tau, \overline{L})$. □

Proof of Proposition 4. Let $L' < L''$ and suppose there is an equilibrium where the searcher searches $k' > 0$ more times when complexity is L' (and on-path concludes search at h') and an equilibrium where the searcher searches $k'' > 0$ more times when complexity is L'' (and on path concludes search at h''). For contradiction, suppose that $k' > k''$.

Note first that $\max_{x \in S} q^u_{h'',L''}(x) < 1$ (and, therefore, $\max_{x \in S} q^u_{h'',L'}(x) < 1$); otherwise, the searcher would have been better off concluding search immediately at h.

Next, note that conditional on searching k more times (i.e., among all on-path histories that conclude after after h), the location of an optimal set of k searches when complexity is L'' is also optimal when complexity is L'. This is easy to see, for example, from the description of the optimal search algorithm in Appendix B.

By an argument analogous to that in the proof of Lemma 12,

$$0 < \max_{x \in S} q^u_{h'',L'}(x) - \max_{x \in S} q^u_{h',L'}(x) < \max_{x \in S} q^u_{h'',L''}(x) - \max_{x \in S} q^u_{h',L''}(x),$$

where the first inequality is by the optimality of concluding at h' over searching fewer times and concluding at h''.[29]

This however implies that the marginal benefit of concluding at history h' rather than history h'' is higher when complexity is L'' than when complexity is L'. Since the benefits net of costs of searching the additional $k' - k''$ times are non-negative when complexity is L', they are strictly positive when complexity is L''. This contradicts the assumption that at h, the searcher optimally plans to conclude search at h'' when complexity is L''. □

F.4 Proofs for Section 6

Proof of Proposition 5. We prove this result by constructing a candidate z'. Let $q^l_{h_t}$ be the lower envelope of Q_{h_t} and let $z' \equiv q^l_{h_t}(\sigma(h_t))$. Note that by definition, $Q_{h'_{t+1}}$ is nonempty, and $z_{t+1} \geq z'$. By construction, and the assumption that $\min_{q \in Q_{h_t}, y \in S} q(y) > 0$, $q^u_{h_{t+1}}(y) = L|\sigma(h_t) - y| + z' \leq q^u_{h_t}(y)$ for all $y \in [x-d, x+d]$, where $d = |x - \sigma(h_t)|$. By Proposition 1, any any search in $[x-d, x+d]$ could not be a part of an optimal search procedure, proving the result. □

Proof of Proposition 6. If the searcher learns good news at $\sigma(h_t)$ and $\sigma(h_t) > x_t$, then $q(x) \leq z_t$ for any $q \in Q_{h_{t+1}}$. Otherwise, q is not quasiconcave. Moreover, $q^u_{h_t}(x) \leq z_t + L(x - x_t)$ for any $x \in S$, by Lemma 1 and the fact that $Q^{QC} \subset Q^{MP}$. Therefore, if $x < x_t + \frac{1}{L}(z_{t+1} - z_t)$, then $q^u_{h_t}(x) < z_{t+1}$. By Proposition 1, any any search in $[\min S, x_t + \frac{1}{L}(z_{t+1} - z_t)]$ cannot be a part of an optimal search procedure.

The proof of the remaining cases follow identical arguments. □

F.5 Proofs for Appendix A

Proof of Proposition 8. Fix some history $h_t \in H$. Recall σ^A_d is the 'no news' strategy of the adversary (see Section 4.2). Let σ^C denote the searcher's optimal strategy in the commitment game, should the adversary respond with σ^A_d. Since the adversary observes the searcher's action after each history, it is without loss of generality to take σ^C to be a pure strategy.

Let T be the sequence of items that are explored on path under the pair of strategies (σ^C, σ^A_d). Then the simultaneous strategy that explores T in the same order is an optimal simultaneous search strategy, by Proposition 1, in the simultaneous search game at h_t. But by Lemma 6, any strategy that follows an optimal simultaneous search strategy in the sequential search game achieves the same payoff as an optimal simultaneous search strategy does in the simultaneous search game. By Theorem 1, this implies that an optimal sequential search

[29] A difference from the claim in Lemma 12 is that history h' need not follow h'', but this does not change the argument.

procedure (which by Lemma 6 and Proposition 1 is a best response to σ_d^A in the sequential search game) achieves the same payoff as under (σ^C, σ_d^A) in the commitment game. This implies that σ_d^A is optimal for the adversary in the commitment game, since she is no worse off than she would be in the sequential search game by playing it. It also then implies that playing an optimal sequential search strategy is also optimal in the commitment game. □

Proof. Suppose the commitment game starts at history h_0. Let σ^C denote the searcher's strategy in the commitment game. Let h_t be the history where search concludes under the pair of strategies $(\sigma^C, q_{h_0}^A)$.

By playing $q_{h_t}^u$, the adversary is weakly better off than she would be in the simultaneous search game at h_0 (by Proposition 1). So an optimal strategy σ^C in the commitment game must coincide on path with an optimal simultaneous strategy at h_0 should the adversary play $q_{h_t}^u$. This means $q_{h_t}^u$ is optimal for the adversary, since she is no worse off than she would be in the simultaneous search game. The searcher can also achieve the same payoff by committing to an optimal sequential search procedure. □

F.6 Proofs for Appendix E

For the minimal prior case (i.e., $\kappa \in (Q^{MP})^k$), it suffices to adapt Lemma 3 to the present case.

LEMMA **14.** *Consider an attribute learning model where S is a (k-dimensional) rectangular search space, and attributes are substitutes. Let $Q = (Q^{MP})^k$. Fix some history $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1} \in H$, where $z'_i = (z_i^{1\prime}, \dots, z_i^{k\prime})$, and let x be an optimal technology at h'_t. Let $T = X^*_{h'_t} \setminus \{x\}$. For each $x_i \in T$, let $\epsilon_i = (\epsilon_i^1, \dots, \epsilon_i^k) \in \mathbb{R}^k$. Consider an alternate history $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1} \in H$ where $z''_i = z'_i + \epsilon_i$ for $x_i \in T$ and $z''_i = z'_i$ otherwise. Suppose moreover that $Q_{h'_t}$ and $Q_{h''_t}$ are nonempty . The searcher's payoff to concluding search in h''_t is greater than or equal to the searcher's payoff to concluding search at h'_t.*

Proof. The nontrivial case is when T is nonempty and some $\epsilon_i \geq 0$. Let $\epsilon = (\epsilon^1, \dots, \epsilon^k) \in \mathbb{R}^k$ denote the attribute-wise maximum of all the ϵ_i, and let h''' be the history $\{(x_i, z'''_i)\}_{i=0}^{t-1}$, where $z''_i = z'_i + \epsilon$. Let $z', z''' \in [0,1]^k$ denote the attribute-wise maximum of all z'_i's and z'''_i's at histories h' and h''', respectively. Since S is a rectangular search space, there exist $x', x''' \in S$ such that $\kappa(x') = z'$ and $\kappa(x''') = z'''$. Clearly, x' and x''' are optimal at histories h' and h''', respectively (and the searcher, by assumption, can conclude search with x' (or x'''), even if they were previously unexplored).

Recall, $q(x) = f \circ \kappa(x)$ where f is submodular. Let x'_A denote an adversary optimal unexplored alternative in S, with $z'_A = \kappa(x'_A)$, should the searcher conclude at history h'. Define z'''_A analogously. Now $z'''_A \leq z'_A + \epsilon$ (with equality if the upper bound on quality is not binding).

Now $f(z' + \epsilon) - f(z') > 0$, since f is increasing. Next $z' \leq z'_A$. By submodularity, $f(z' + \epsilon) - f(z') \geq f(z'_A + \epsilon) - f(z'_A)$. In other words, $f(z''') - f(z') \geq f(z'''_A) - f(z'_A)$. The left side of this inequality is the increase in the quality of searcher's best discovery, should he conclude search at history h''' instead of h'. The right side is the increase in quality of the best unexplored option, in the worst case. Along with the assumption that $U_1 \geq -U_2$, this implies that the searcher is better off at history h''' than at h'.

The searcher is weakly better off concluding search at history h'' than at h''', by the same reasoning as in Lemma 3. □

Proof of Corollary 1. Applying Lemma 14 in place of Lemma 3, all other parts of the proof are identical to the proof of Theorem 1. □