Northwestern | Kellogg

CMS-EMS
Center for Mathematical Studies in Economics and Management Sciences

Discussion Paper #1612

"Searching in the Dark and Learning Where to Look"

Suraj Malladi
Northwestern University

May 2025

Searching in the Dark and Learning Where to Look

Suraj Malladi*

First draft: April 2022
This draft: May 2025

Abstract

A searcher knows little about how items' observable attributes map to qualities but wants to avoid missing good discoveries. Knowing only that similar items have similar qualities, she sequentially chooses where to look next or stops to take her best discovery. We characterize forward-looking strategies that maximize her payoff guarantee. Relative to a myopic searcher, she typically searches closer to past discoveries but is more willing to explore items with extreme attributes. She stops if the remaining unexplored items are too similar to what she has discovered. She also stops from choice overload if many unexplored items are too dissimilar. JEL: D83, C72

*Thanks to Martino Banchio, Alex Frankel, Jonathan Libgober, Andrzej Skrzypacz, seminar participants at Cornell, Northwestern, Arizona State University, University of Chicago, UCLA, USC and SITE: Dynamic Games for their feedback and suggestions. I am particularly grateful to Ilya Morozov for so many helpful conversations. Northwestern University. Email: suraj.malladi@kellogg.northwestern.edu.

1 Introduction

Any process of new discovery involves search and learning. An inventor tries out different prototypes before settling on a design. A pharmaceutical company experiments with different drug dosages to find the best one. An amateur tennis player demos rackets to learn her preferences over various features.

We model this discovery process and characterize optimal search. The model captures two key features of "exploring unknown territory".

First, the searcher chooses when to stop and what to search next, and this choice is informed by earlier discoveries. An inventor tinkers with a prototype if it is promising but overhauls the design if fails. If a drug causes bad side effects, a pharmaceutical company runs trials with a much lower dosage.

Second, the searcher knows little about which attributes she likes or what is discoverable a priori. She searches so as not to miss out on good discoveries, if they exist. A firm developing a new invention may search the space of designs carefully to avoid being outdone by competitors. A pharmaceutical company may run clinical trials to minimize maximum regret, i.e., to guarantee that the dosage level eventually chosen is as close to the ideal level (Manski, 2023).

The model considers a searcher who searches items in a 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 quasiconcave). The Lipschitz constant is a measure of search complexity: it is a cap on how quickly she thinks quality can vary with attributes.

The searcher sequentially searches different items to learn their quality, and in the process, narrows down the set of quality indices that are consistent with her discoveries. After each search, she decides whether to continue exploring or stop to take the best item she had discovered so far.

Her payoff increases in the quality of her best discovery, decreases in the quality of the best attainable item, and decreases with costly search. She is forward-looking and evaluates any strategy by its worst-case payoff upon eventually stopping. The worst-case outcome is searching hard but missing out on good discoveries. A sequentially optimal strategy always maximizes the searcher's worst-case continuation payoff to avoid such wild goose chases.

We show that all sequentially optimal strategies are iterative simultaneous search strategies: every period, the searcher acts as if she must irrevocably decide the set of items to explore before making a choice and then searches one of those items. Relative to our forward-looking searcher, a myopic searcher would stop sooner and typically search deeper into unknown territory.

The model can tie out why search effort may increase at first with search complexity (Griffin and Broniarczyk, 2010) but eventually collapse due to 'information/choice overload' (Scheibehenne et al., 2010). It can also reproduce funnel-like search dynamics observed in online shopping data (Bronnenberg et al., 2016; Blake et al., 2016), where searchers initially explore broadly in attribute space but narrowly over time. Finally, sequentially optimal strategies reflect Weitzman's intuition about directed 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 unknown territory, the searcher plans her searches to minimize a measure of distance between the anticipated set of explored and 'relevant' unexplored items at the conclusion of search. And roughly speaking, proximity in search space corresponds to correlation in a Bayesian model.

1.1 Related literature

A broad literature on search with learning considers settings where search order is either irrelevant or exogenous. Rothschild (1974) considers an agent who draws independent samples from an unknown distribution, learns its shape and decides when to stop; Bikhchandani and Sharma (1996) generalize this to allow for recall and different distributions. Schlag and Zapechelnyuk (2021) find stopping policies that achieve a large fraction of achievable surplus regardless of the true distribution. In addition to optimal stopping, Urgun and Yariv (2021) and Wong (2021) study how agents choose the speed of exploration.

Far less is known about optimal search with learning when agents can choose search order. Weitzman (1979) studies directed search without learning: items have independent payoffs, so learning about one is uninformative about another. Introducing correlations generally makes the model intractable, but Adam (2001) solves a special case: items have independent and unknown payoff distributions, but items in the same observable group have identical distributions. In a richer setting, Callander (2011) characterizes where a sequence of short-lived agents (or a long-lived, myopic agent) search when the mapping from choices to payoffs is a Brownian motion path. Garfagnini and Strulovici (2016) study a similar framework in which agents live two periods. We study a fully forward-looking agent, as in Adam (2001), but where the qualities of items are flexibly interrelated, as in Callander (2011).

As Hörner and Skrzypacz (2017) note, understanding how dynamics change in "Callander's model for patient agents is an important open problem." Our model is close enough to shed some light. Relative to a patient agent, a myopic agent searches farther from past discoveries but avoids extreme attributes.

We follow a growing body of work in economics that considers maxmin or minmax regret objectives; see reviews by Carroll (2019) and Banerjee et al. (2017). Few papers in this literature consider dynamics, especially without commitment. Libgober and Mu (2022) is a recent exception. They also consider a sequential maxmin framework in the durable goods monopoly problem and find that their agent behaves in a dynamically consistent fashion.

Some dynamics in Bayesian spatial search models, like funneling in Callander (2011) and Hodgson and Lewis (2020), arise here for similar reasons. Unlike these models, ours rationalizes choice overload, as regret aversion makes finding good outcomes in complex sets seem like 'finding a needle in a haystack'.

Our approach is related to the literature on adversarial multi-arm bandits and optimization (see reviews by Lattimore and Szepesvári (2020) and Hansen, Jaumard and Lu (1992), resp.). Unlike bandits, search does not involve uncertain flow utility. Unlike optimization, search entails optimal stopping. While this literature studies algorithms with good asymptotic properties (e.g., speed of convergence or regret), we study how rational agents optimally search.

2 The model

There is a *searcher*, and $S \equiv [0, 1]$ is the *search space*, i.e., the set of items to be explored. Let $\Omega \subset [0, 1]^S$ be the set of states or *quality indices*—mappings from the search space to quality. There is a true quality index $q \in \Omega$, and each item $x \in S$ has a quality $q(x) \in [0, 1]$.

The searcher does not know the true quality index. She can learn the quality of any item $x \in S$ through costly search. Each item-quality pair is a point on the true quality index. Search narrows down the set of plausible states to those that pass through all the points discovered.

2.1 Timing and strategies

In each period, $t = 0, 1, 2, \ldots$, the searcher takes one of two kinds of actions. She either searches an item $x_t \in S$ to learn its quality, $q(x_t)$. Or she stops search, denoted by $x_t = \emptyset$, to take her best discovery to date.

Let $h_t = \{(x_i, z_i)\}_{i=1}^t$ be a time t history where search has not yet stopped, with $z_i = q(x_i)$. Let $X_{h_t} \equiv \{x_1, \ldots, x_t\}$ be the set of items searched at h_t. Let $X^*_{h_t} \equiv \arg\max_{x_i \in X_{h_t}} q(x_i)$ be the set of *best discoveries* at h_t, with corresponding quality denoted by $z^*_{h_t}$. The *empty history* is h_0, and $z^*_{h_0} = 0$.

The searcher discovers *good news* at h_t if $q(x_t) > z^*_{h_{t-1}}$, *bad news* if $q(x_t) < z^*_{h_{t-1}}$, and *no news* if $q(x_t) = z^*_{h_{t-1}}$.

Let $\Omega_h \subset \Omega$ be the set of quality indices that are *consistent* with what the searcher observes at history h. That is, Ω_h is the set of quality indices $\tilde{q} \in \Omega$ such that $\tilde{q}(x_i) = z_i$ for all $(x_i, z_i) \in h$.

Let H denote the set of all histories h where Ω_h is nonempty. A *strategy* $\sigma : H \to S \cup \{\emptyset\}$ is a deterministic map from histories to actions. From $h \in H$ and for $\tilde{q} \in \Omega_h$, a strategy σ reaches a *terminal history* $h(\sigma, \tilde{q}) \in H$ if this is the first on-path history such that $\sigma(h(\sigma, \tilde{q})) = \emptyset$. Let Σ be the set of strategies that reach a terminal history, starting from any $h \in H$ and for any $\tilde{q} \in \Omega_h$.

2.2 Searching in the dark

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

Assumption **1.** *Either* $\Omega = Q_L$ *or* $\Omega = Q_L^{QC}$.

The searcher knows little about the true quality index, a priori.

First, she knows there are limits to quality, as $0 \leq q \leq 1$.

Next, she knows that proximate items in S cannot be too different in quality. If an item's location in S is interpreted as an index of its observable attributes (e.g., drug dosage), then observably similar items have similar qualities. The Lipschitz constant L, or *search complexity*, is a cap on how quickly quality can vary with observable attributes; it is known to the searcher. If L is low, knowing a couple points on q reveals a lot about its shape. If L is high, finding a good outcome may be like finding a needle in a haystack.

Finally, if $\Omega = Q_L^{QC}$, the searcher also knows that there is an ideal range (or "sweet spot") in S, and that items further from this range are of lower quality. For example, there may be an ideal drug dosage level ('therapeutic window') below which efficacy drops and above which risks increasingly outweigh benefits. Alternatively, a buyer may be aware that she has single-peaked preferences over ski boot sizes, motorcycle horse power, racket string tension, etc., even before she demos different products to discover what she likes.

2.3 Payoffs

The ex post payoff to choosing item x when q is the true quality index is

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

where $U : [0,1] \to \mathbb{R}_+$. The payoff to stopping at h_0 is $U(0, \max_{y \in S} q(y))$.

Assumption **2.** *First, U is differentiable almost everywhere and continuous. Next, U is increasing in the first argument (i.e., quality of the chosen item)* *and decreasing in the second argument (i.e., the quality of the best item in S). Finally, $U_1 + U_2 \geq 0$ almost everywhere (i.e., translating the true quality index q upward leaves the searcher better off).*

An example of payoffs satisfying Assumption 2 is linear regret:

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

The searcher pays $c > 0$ for each search. Her payoff net of search costs to stopping at h_t is:

$$p(h_t, q) = U(z^*_{h_t}, \max_{x \in S} q(x)) - c \cdot t. \tag{1}$$

Restricting to strategies that always reach a terminal node is without loss of generality, as the payoff to searching indefinitely is formally set to $-\infty$.

At any $h \in H$, the searcher evaluates a strategy σ by its worst-case continuation payoff or *payoff guarantee*:

$$\inf_{\tilde{q} \in \Omega_h} p(h(\sigma, \tilde{q}), \tilde{q}).$$

We describe payoff guarantees in examples where U is linear regret and $L = 1$.

Example 1: A strategy σ that stops at h_0 has a payoff $p(h_0, q) = -\max_{y \in S} q(y)$. Therefore, its payoff guarantee is $\inf_{\tilde{q} \in \Omega}(-\max_{y \in S} \tilde{q}) = -1$.

Example 2: Consider a strategy σ where $\sigma(h_0) = 0$, and $\sigma(h_1) = 1$ if $q(0) < \epsilon \approx 0$, and $\sigma(h_1) = \emptyset$ otherwise. The payoff guarantee to σ depends on c. Searching more reduces worst-case regret but is not worth it if costs are high.

If $c = 0.75$, then in the worst case, $q(y) = 0.5 - |0.5 - y|$: the searcher searches twice and face a regret of 0.5, leading to a payoff guarantee of -2.

If $c = \epsilon$, then in the worst case, $q(y) = \min\{y + \epsilon, 1\}$: the searcher searches only once and face regret of $1 - \epsilon$, leading to a payoff guarantee of $1 - 2\epsilon$.

2.3.1 Interpretation of payoffs

The searcher's payoff, for example, can capture a firm whose future profits depend not only on the quality of its chosen technology but also on the technology chosen by a strong competitor who conducts extensive R&D. The second argument of U is the quality of the competitor's technology.

The model can also capture settings where the searcher never learns her payoff but imagines the worst case. For example, a shopper may feel buyer's remorse for potentially missing out on a better purchase, whether or not one exists. She would worry less if she searched extensively before buying, as this reduces the scope for missing a much better product.

2.4 The searcher's problem

A strategy $\sigma^* \in \Sigma$ is (*ex ante*) *optimal* if it maximizes the payoff guarantee at the empty history h_0:

$$\sigma^* \in \arg\max_{\sigma \in \Sigma} \big\{ \inf_{\tilde{q} \in \Omega} p(h_0(\sigma, \tilde{q}), \tilde{q}) \big\}. \tag{2}$$

As the searcher discovers more points on the true quality index, the set of consistent quality indices shrinks, so the worst-case outcome to different strategies may change. A strategy is sequentially optimal if it remains optimal in every contingency. Formally, a strategy $\sigma^* \in \Sigma$ is *sequentially optimal* if at every history $h \in H$,

$$\sigma^* \in \arg\max_{\sigma \in \Sigma} \big\{ \inf_{\tilde{q} \in \Omega_h} p(h(\sigma, \tilde{q}), \tilde{q}) \big\}. \tag{3}$$

A sequentially optimal strategy is also ex ante optimal, because it satisfies eq. (3) at the empty history, which is precisely eq. (2). A searcher following a sequentially optimal strategy would never want to deviate from it at any history. To capture learning, we characterize sequentially optimal strategies.

3 Optimal search

In section 3.1, we define simultaneous search strategies. This is useful for characterizing sequentially optimal strategies in Section 3.2 and Section 3.3. Section 3.4 discusses the intuition for the results and the role of the assumptions. Section 3.5 gives an algorithm for computing optimal search strategies.

3.1 Simultaneous search

Simultaneous search refers to static models where an agent commits to the set of items she will explore (Stigler, 1961; Chade and Smith, 2006). We capture simultaneous search by restricting to strategies that are independent of what is learned along the way. "Search $x_1 = 0.3$, then $x = 0.9$, then stop," is a simultaneous search strategy. "Search $x_1 = 0.3$, then search $x = 0.9$ if $q(x_1) \leq 0.5$ but stop if $q(x_1) > 0.5$," is not a simultaneous search strategy.

Formally, σ is a *simultaneous search strategy* at history $h \in H$ if it is measurable with respect to the calendar date at histories after h. Let $\Gamma_h \subset \Sigma$ denote the set of all simultaneous search strategies at h. A strategy $\sigma_s^* \in \Gamma_h$ is an *optimal simultaneous search strategy at $h \in H$* if

$$\sigma_s^* \in \arg\max_{\sigma \in \Gamma_h} \Big\{ \inf_{\tilde{q} \in \Omega_h} p(h(\sigma, \tilde{q}), \tilde{q}) \Big\}. \tag{4}$$

Next, we say $\sigma \in \Sigma$ *follows optimal simultaneous search strategies* if at every $h \in H$, there exists an optimal simultaneous search strategy $\sigma_{s,h}^* \in \Gamma_h$ such that $\sigma(h) = \sigma_{s.h}^*(h)$. Such a strategy σ need not be a simultaneous search strategy itself: it uses the quality of the latest discovery to recompute an optimal simultaneous search strategy, whereas the search path under a simultaneous search strategy would be independent of the qualities discovered.

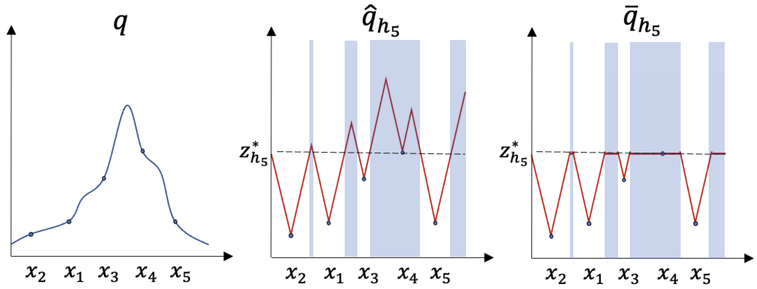

Figure 1: On the left is the true quality index. At h_5, the searcher has only observed qualities at x_1, x_2, x_3, x_4 and x_5. In the middle is the upper envelope of consistent quality indices when $\Omega = Q_L$. The search window is highlighted. On the right is the 'no news' quality index.

3.2 Sequentially optimal search strategies

THEOREM **1.** *Sequentially optimal search strategies exist. A search strategy σ is sequentially optimal if and only if it follows optimal simultaneous search strategies.*

Appendix C.1 and Appendix C.2 prove Theorem 1 in the cases where $\Omega = Q_L$ and $\Omega = Q_L^{QC}$, respectively.

Theorem 1 says that optimal sequential search is iterative, optimal simultaneous search. The searcher solves for an optimal simultaneous search strategy each period, but only executes one step of that plan. If the realized quality from that search is not what she anticipated, she computes a new plan.

3.3 Optimal simultaneous search

Theorem 1 evokes the question: what are optimal simultaneous search strategies? We characterize such strategies here.

Let $\hat{q}_h$ be the upper envelope of Ω_h, i.e., $\hat{q}_h(x) = \max_{\tilde{q}\in\Omega_h} \tilde{q}(x)$ for all $x \in S$.

The *search window*, $S_h \subset S$, is the set of items x for which $\hat{q}_h(x) > z_h^*$; good news is only possible inside the search window. Let $\bar{q}_h(\cdot) \equiv \min\{\hat{q}_h(\cdot), z_h^*\}$ be the quality index where the searcher always finds no news in S_h; see Figure 1.

Starting from any $h \in H$, a strategy $\sigma \in \Gamma_h$ that only explores items in the search window S_h reaches terminal history $h' \equiv h(\sigma, \bar{q}_h)$ if every subsequent discovery yields no news. Let $q_h^\sigma \equiv \hat{q}_{h'}$ be the the quality index which also delivers no news on path but good news everywhere else in the search window.

Proposition 1. *At any $h \in H$, optimal simultaneous search strategies exist, and $\sigma_s^* \in \Gamma_h$ is an optimal simultaneous search strategy if and only if:*

$$\sigma_s^* \in \arg\max_{\sigma \in \Gamma_h} \left\{ p(h(\sigma, q_h^\sigma), q_h^\sigma) \right\}.$$

Proposition 1 says that optimal simultaneous search strategies are those that hedge against the risk that each new discovery is only as good as the previous best, but what is left unexplored is as high quality as possible.

3.4 Discussion

We give intuitions for the results and highlight the role of the assumptions.

3.4.1 Intuition for Proposition 1: fear of missing out

A simultaneous strategy at some history h realizes a low payoff if it discovers only low quality items (the *direct force*). The realized payoff would be lower still if some undiscovered items have high quality (the *indirect force*).

The direct force is maximal if all future discoveries are of quality below z_h^*, the searcher's outside option upon stopping. Whether these discoveries yield bad news or no news does not affect the quality of the item eventually chosen.

The indirect force is greatest when the quality of the best unexplored item is as large as possible. If the searcher's future discoveries are far worse than z_h^*, then continuity of q ensures that unexplored items cannot be so much better.

The perfect storm is that all future discoveries yield no news (maximizing the direct force), and the best unexplored item is as large as possible conditional on this. The indirect force would be even stronger if the searcher ever discovers good news, but she is better off in such scenarios, as $U_1 + U_2 \geq 0$.

Preparing for this perfect storm yields an optimal simultaneous search strategy. The proof shows that the converse also holds.

3.4.2 Intuition for Theorem 1: wild goose chases

In sequential search, strategies can be conditioned on the qualities of discoveries made along the way. Here, a third source of low payoffs is failing to learn where or how good the best undiscovered items are and searching too much.

Upon discovering a low-quality item, the searcher learns that items with similar attributes are also low quality. She would thereafter redirect her attention to a more promising area of the search space or stop search early. Similarly, a good discovery may alert the searcher to try similar items. A discovery yielding no news is least informative as to where to search next.

So at any history, the worst-case scenario when conducting simultaneous search is still the worst-case for sequential search. The worst-case scenario changes if good or bad news arises but stays the same under no news, so following optimal simultaneous strategies is sequentially optimal.

Sequentially optimal strategies do not always exist in decision problems with max-min objectives. Appendix A clarifies why they exist here.

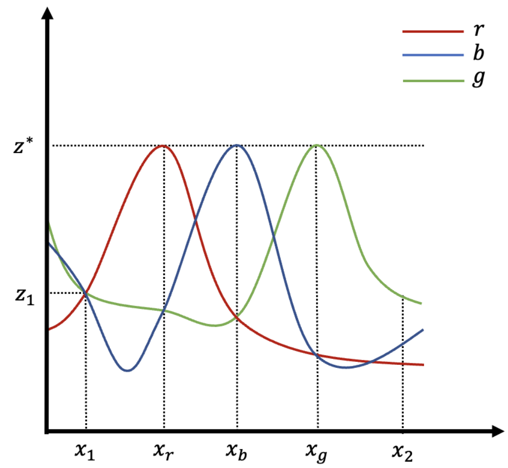

Figure 2: r, b, and g attain maxima at x_r, x_b, and x_g respectively.

3.4.3 The role of searcher's knowledge

The assumption that $\Omega = Q_L$ or $\Omega = Q_L^{QC}$ is crucial, and Theorem 1 may not hold when the searcher knows more *a priori* about the quality index's shape.

For example, let $\Omega = \{r, b, g\}$ as pictured in Figure 2. Consider history $h_1 = \{(x_1, z_1)\}$ where only the quality of x_1 is known. Because $r(x_1) = b(x_1) = g(x_1) = z_1$, all three quality indices are consistent at h_1.

If c is small and $U_1 > 0$, an optimal simultaneous strategy searches x_r, x_b and x_g and stops. One of these is guaranteed to be of quality z^*. Exploring other items is clearly wasteful. Less search gives a lower payoff in the worst case. For example, searching only x_b and x_g yields a low payoff if $q = r$.

However, no optimal (sequential) strategy starts by searching x_r, x_b or x_g. Suppose, for contradiction, the searcher first explores x_r and learns that $q \neq r$. Because $b(x_r) = g(x_r)$, both b and g remain consistent. Suppose she next explores x_b, and learns that $q \neq b$. She deduces that $q = g$ and searches x_g: following optimal simultaneous strategies takes up to three searches. Instead, if the searcher were to first explore x_2, she would immediately identify q, as $r(x_2) \neq b(x_2) \neq g(x_2)$. She would next search the highest quality item. This strategy guarantees z^* with just two searches.

In this example, the potentially high-quality items are not the most informative. Optimal simultaneous strategies only explore the former while sequentially optimal strategies start with the latter.

Note that Theorem 1 holds for $\Omega = Q$, where $Q \subset Q_L$ (or $Q \subset Q_L^{QC}$) and Q contains every piecewise-linear $\tilde{q} \in Q_L$ ($\tilde{q} \in Q_L^{QC}$). For such a Ω, the worst quality index upon stopping is the same as when Assumption 1 holds.

3.4.4 The role of deterministic strategies

Whether the restriction to deterministic strategies is with loss of generality depends on how the searcher would evaluate mixed strategies. Here, we adapt the discussions in Saito (2015) and Ke and Zhang (2020) who make this point generally for maxmin expected utility models.

If the searcher evaluates a mixed strategy $\sigma \in \Delta(\Sigma)$ by the expected worst-case payoffs to each pure strategy in its support, randomizing would not help:

$$\max_{\sigma \in \Sigma} \big\{ \inf_{\tilde{q} \in \Omega} p(h_0(\sigma, \tilde{q}), \tilde{q}) \big\} = \max_{\sigma \in \Delta(\Sigma)} \big\{ \mathbb{E}[\inf_{\tilde{q} \in \Omega} p(h_0(\sigma, \tilde{q}), \tilde{q})] \big\}.$$

This can capture a searcher who uses a randomization device (e.g., a coin flip) to pick a strategy but cannot commit to the result of the randomization (e.g., flips again if the selected strategy had a lower worst-case payoff than another).

If the searcher instead evaluates $\sigma \in \Delta(\Sigma)$ by its worst-case expected payoff, she is weakly better off randomizing, because

$$\max_{\sigma \in \Sigma} \big\{ \inf_{\tilde{q} \in \Omega} p(h_0(\sigma, \tilde{q}), \tilde{q}) \big\} \leq \max_{\sigma \in \Delta(\Sigma)} \big\{ \inf_{\tilde{q} \in \Omega} \mathbb{E}[p(h_0(\sigma, \tilde{q}), \tilde{q})] \big\}.$$

Appendix B shows an example where this inequality is strict. Even under this interpretation, the model can be applied to settings where randomization is simply infeasible, e.g., due to organizational constraints to how firms search.

3.5 Computing optimal strategies

We describe how to compute an optimal action at any history, $h \in H$.

1. Fix a $\bar{k}_h \in \mathbb{N}$ such that the cost of searching $\bar{k}_h + 1$ or more times outweighs any possible marginal benefit.

2. For each of $k = 0, 1, 2, \ldots, \bar{k}_h$, find a k-*center*, i.e., a set of k items in S_h such that no undiscovered item in the S_h is too far from the nearest of these k items.[1] For example, $S_{h_0} = [0, 1]$, so $\{1/2\}$ is a 1-center, $\{1/4, 3/4\}$ is a 2-center, etc.

3. For each k, consider a strategy $\sigma \in \Gamma_h$ that searches items in the k-center in any order and stops. Compute the payoff when $q = q_h^\sigma$.

4. For k^* corresponding to the highest payoff: stop search if $k^* = 0$, and otherwise search any item in the k^*-center.

Each period, the searcher plans to eventually have every unexplored item in the search window be close to some explored item. This procedure reflects Weitzman's intuition (quoted in Section 1) that ordered search with correlated rewards favors exploring items most correlated to unexplored options. Here, proximity is roughly the analogue of correlation. Searching an item $x \in S$ reveals its quality, $q(x)$, but also constrains the maximum and minimum possible

[1] Among all subsets of size k in S_h, a k-center minimizes the Hausdorff distance between (1) S_h and (2) the union of these k items and the boundary points of S_h in $S \setminus S_h$.

values q can take elsewhere in S. These bounds are tighter for items closer to x, so the searcher learns more about the quality of proximate items.

4 Forward-looking versus myopic search

We compare sequentially optimal strategies to myopic strategies.

For any $h \in H$, let $\Sigma_h^M \subset \Sigma$ be the set of strategies that stop at h or search once and stop immediately thereafter. A strategy σ^M is *myopic* at h if:

$$\sigma^M \in \arg\max_{\sigma \in \Sigma_h^M} \big\{ \inf_{\tilde{q} \in \Omega_h} p(h(\sigma, \tilde{q}), \tilde{q}) \big\}.$$

4.1 Propensity to Search

Myopic searchers stop sooner than forward-looking searchers. Formally, if $h \in H$ is a terminal node for some sequentially optimal search strategy σ, then there is a myopic strategy σ^M such that $\sigma^M(h) = \emptyset$.

To see why, suppose all myopic strategies search at h. Then continuing search at h for one period and stopping afterward can be strictly better than stopping immediately, so any sequentially optimal strategy would also search.

4.2 Search Location

If there is a unique longest interval $I \subset S_h$, we refer to it as the *most promising area* at h and denote it by P_h; see Figure 3.

PROPOSITION **2.** *Let $h \in H$. For any myopic strategy at h, σ^M, either $\sigma^M(h) = \emptyset$, or $\sigma^M(h)$ is in most promising area at h.*

Searching exclusively outside of P_h would not teach the searcher anything about how good the best unexplored alternative could be. A forward-looking

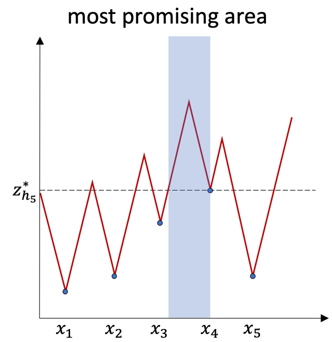

Figure 3: The most promising area at h_5 supports the tallest peak of $\hat{q}_{h_5}$.

searcher may search outside of P_h in anticipation of searching there later, as she is indifferent to search order, but a myopic searcher would not.

4.2.1 Embracing uncertainty versus avoiding extreme attributes

A myopic searcher often ventures deeper into unknown territory while a forward-looking searcher explores closer to previous discoveries. Quality at $x \in S_h$ is *more uncertain* than the quality at $y \in S_h$ if $\hat{q}_h(x) \geq \hat{q}_h(y)$, i.e., if x is farther away from discovered items or the boundary of the search window. If some myopic strategy at h does not stop immediately and if the item x' of most uncertain quality is in the interior of S, searching x' and stopping is a myopic strategy at h (by the definition of a 1-center in Section 3.5). A myopic searcher explores more uncertain items to make the most of one search.

There is a countervailing force if the item of most uncertain quality is at the boundary of S . If, say, $\hat{q}_h(x)$ is maximized at $x = 0$, no myopic strategy searches at 0 (searching slightly to the right of 0 reduces the distance to the best unexplored alternative in the worst case). And if uncertainty outside of P_h is sufficiently low, some sequentially optimal strategy searches weakly closer to 0 (i.e., an item of more uncertain quality) than any myopic strategy.

Intuitively, a patient searcher is open to trying items with extreme attributes (e.g., the heaviest golf club), because she will search again if need be.

5 Complexity and stopping

Search complexity, L, measures the ex ante difficulty of finding good outcomes. For example, the safety or efficacy of some compounds may be known to be very sensitive to dosage (high L), while others are more stable (low L).

We characterize how search complexity affects the amount of search. In light of Proposition 1, a history $h_t \in H$ is *on-path* if it is nonempty and $q(x_1) = \ldots = q(x_t)$. We focus on such histories, as they can arise for any L.

5.1 Choice overload

The first result is that search effort is non-monotonic in complexity: the searcher stops if complexity is sufficiently low or high.

PROPOSITION **3.** *At any on-path $h \in H$, there exist $\underline{L}, \overline{L} \in \mathbb{R}_{++}$ such that if $L < \underline{L}$ or $L > \overline{L}$, any sequentially optimal strategy stops search.*

When L is small, the quality index is fairly flat. Undiscovered items are similar to past discoveries, so the searcher stops.

If L is large, the searcher anticipates that even after considerable search, she may be nowhere near the peak outcome. Making a worthwhile discovery is like finding a needle in a haystack, so she stops due to "choice overload."

Scheibehenne et al. (2010) review the empirical literature on choice overload, suggesting that when the "precondition [of] lack of familiarity with, or prior preferences for, the items in [a] choice assortment" is met, large and complex assortments cause agents to avoid search in favor of known or default options. In a meta-analysis of choice overload experiments, Chernev et al. (2015) conclude that "regret, as an operationalization of individuals' decision goal, was a particularly strong driver of choice overload."

Our model explains choice overload on these exact lines. Section 5.1.2 and Section 5.1.1 discuss the role of regret and large assortments in Proposition 3.

5.1.1 The role of min-max regret

By Assumption 2, U is increasing in the quality of the best discovery and decreasing in the quality of the best available item, $\max_{y \in S} q(y)$. The searcher's objective is a generalization of min-max linear regret.

If U depends only on its first argument, the unique sequentially optimal strategy is to stop at any $h \in H$. In the worst case, the continuation payoff to any amount of search would be negative due to the possibility that searching is never worth the cost (e.g., if $q \leq z_h^* + c$ everywhere). Therefore, regret-like objectives generate search and deter it only when complexity is very high.

5.1.2 The role of large S

Suppose S were finite. For the same reason as in Proposition 3, search stops immediately if complexity is sufficiently low. On the other hand, if S is small, search might not stop even when complexity is very high.

When S is finite, the searcher can eliminate regret by exploring every item. So if complexity is sufficiently high, the optimal strategy is either "leave no stone unturned"—search items in any order until a sufficiently high-quality discovery is made or all of S is explored—or to stop immediately. If search costs are sufficiently low (or if S is sufficiently small), "leaving no stone unturned" is optimal. High search complexity shuts down search only if there are too

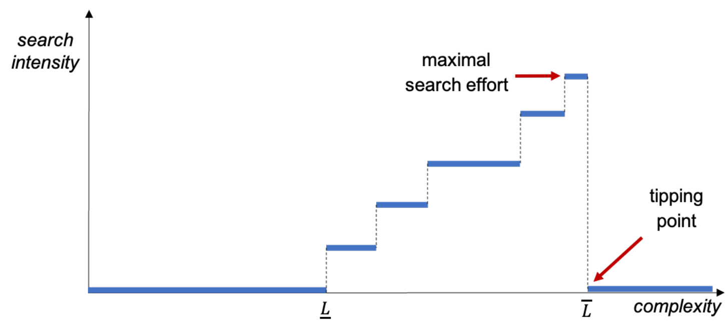

Figure 4: Search intensity correspondence at some on-path history.

many items left to explore.

5.2 Search intensity

We give a converse to Proposition 3 and describe search intensity when there are *decreasing returns to search*, i.e., when $U_{22} < 0$. Let $\mathcal{I}^h(L)$ be the set of $k \in \mathbb{N}$ such that, when search complexity is L, there is an optimal simultaneous search strategy at h that searches k more times in the worst case. $\mathcal{I}^h(L)$ measures the number of intended searches and approximates realized search effort if the true quality index is relatively flat in most places.

Proposition 4. *Let $h \in H$ be on-path and suppose there are decreasing returns to search. Then there exist $\underline{L}, \overline{L} \in \mathbb{R}_{++}$ such that*

1. *if $L < \underline{L}$ or $L > \overline{L}$, every sequentially optimal strategy stops,*
2. *if $L \in (\underline{L}, \overline{L})$, no sequentially optimal strategy stops,*
3. *every selection from $\mathcal{I}^h(\cdot)$ is non-decreasing on $(\underline{L}, \overline{L})$.*

Proposition 4 implies that search intensity is 'cliff-shaped' in complexity (see Figure 4). When complexity is low, there is no search. As complexity rises, search intensity initially ramps up but eventually drops back to zero.

Near the tipping point, small differences in the perceived complexity determine whether the searcher exerts maximal effort or gives up exploration entirely. Therefore, there may be large effects to subsidizing costs if agents are paralyzed by the perceived difficulty of making good discoveries.

These comparative statics also reflect empirically observed patterns of directed search in unfamiliar territories. For example, Griffin and Broniarczyk (2010) study "consumer search in categories in which people have limited knowledge" and find that making the set of available options harder to compare (i.e., more complex) increases search effort. Proposition 4 rationalizes increasing effort and choice overload at different levels of complexity.

6 News and search location

The searcher explores as if she expects no news. But generically, good or bad news arises, and we study how this affects where she looks.

Previous sections showed results that were less tractable or absent in Bayesian spatial search models. Here, we instead highlight parallels to such models.

6.1 Search step size

The next proposition says that after sufficiently bad news, an optimal strategy either stops or jumps past the nearest discovery to another part of S.

Proposition 5. *Let σ be a sequentially optimal strategy and consider any nonempty history $h_t \in H$ at which $\sigma(h_t) = x_{t+1} \in S$ and $\min_{\tilde{q} \in \Omega_{h_t}} \tilde{q}(x_{t+1}) > 0$. Let x be the closest item to x_{t+1} in X_{h_t}. Then, there exists $z < z^*_{h_t}$ such that:*

1. *$\Omega_{h'}$ is non-empty, where $h' \equiv \{(x_1, z_t), \ldots, (x_t, z_t), (x_{t+1}, z)\}$.*

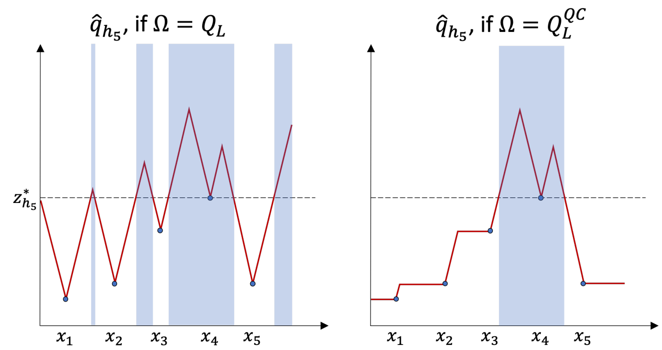

Figure 5: At the same history, $\hat{q}_h$ and S_h can vary with Ω.

2. *If $z_{t+1} \leq z$, then either $\sigma(h_{t+1}) = \emptyset$ or $|\sigma(h_{t+1}) - \sigma(h_t)| > |\sigma(h_t) - x|$.*

Proposition 5 matches the empirical finding in Hodgson and Lewis (2020) that "when consumers view products with "surprisingly low utility" …they jump further away in attribute space." Note that when L is smaller (i.e., items are more similar), the step size following such bad news is larger.

6.2 Funnel shaped search

The path of optimal search varies depending on whether $\Omega = Q_L$ or $\Omega = Q_L^{QC}$, because $\hat{q}_h$ varies with Ω; see Figure 5. News has a more dramatic effect on search location when the quality index is known to be quasiconcave.

Proposition **6.** *Let $\Omega = Q_L^{QC}$, and let $h_t \in H$ be such that $\sigma(h_t) \equiv x_{t+1} \in S$.*

1. *Suppose the searcher finds good news at h_{t+1}: $z_{t+1} > z^*_{h_t}$.*

   (a) *If $x_{t+1} > x_t$, then $S_{h_{t+1}} \subset [x_t + \frac{1}{L}(z_{t+1} - z^*_{h_t}), 1]$.*

   (b) *If $x_{t+1} < x_t$, then $S_{h_{t+1}} \subset [0, x_t - \frac{1}{L}(z_{t+1} - z^*_{h_t})]$.*

2. *Suppose the searcher finds bad news at h_{t+1}: $z_{t+1} < z^*_{h_t}$. Let $x^*_{h_t} \in X^*_{h_t}$.*

(a) If $x_{t+1} > x^*_{h_t}$*, then* $S_{h_{t+1}} \subset [0, x_{t+1} - \frac{1}{L}(z^*_{h_t} - z_{t+1})]$.

(b) If $x_{t+1} < x^*_{h_t}$*, then* $S_{h_{t+1}} \subset [x_{t+1} + \frac{1}{L}(z^*_{h_t} - z_{t+1}), 1]$.

Good and bad news events cause the search window to close in when the searcher knows there is a 'sweet-spot'. Because optimal strategies only explore inside this window, search unfolds in a 'funnel shape'. 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) and Hodgson and Lewis (2020) in ordered-search data from online shoppers.

Callander (2011) finds a similar "triangulating" pattern in a myopic strategy to find a zero of a Brownian motion path. There, too, funneling arises from knowing that there is a sweet spot (i.e., a zero) and learning that one had undershot or overshot it (from seeing a positive or negative draw). These dynamics can, therefore, arise whether the searcher is myopic or forward looking, Bayesian or ambiguity averse, or faces a bandit or optimal stopping problem.

7 Extensions

We highlight some dimensions along which the model and results generalize.

7.1 History-dependent costs

At any history, the searcher is indifferent to the order in which she explores the k^*-center. But in some applications, searching farther from previous discoveries is costlier (e.g., overhauling versus tinkering with a prototype's design).

The model can be extended to let the cost of searching an item x depend on the history h. Suppose the search cost is $C(x, h)$, where (1) $C : S \times H \to \mathbb{R}_{++}$ is continuous and bounded away from 0, and (2) C may depend on the sequence of items $\{x_i\}_{i=1}^t$ explored to date but not their qualities $\{z_i\}_{i=1}^t$. The statement and proof of Theorem 1 hold for this more general model. However, the algorithms for computing optimal strategies vary with C.

7.2 Topology, cardinality and dimension of S

The assumption that $S = [0,1]$ fits settings where there is a continuum of choices along a one-dimensional attribute space (e.g., dosage of a drug). All results would also hold if attributes lie in a 'Salop' circle (e.g., product color).

In motivating applications like consumer search, the sercher faces a finite set of products (e.g., different tennis rackets), each with many attributes (e.g., weight, head size and stiffness). Both Theorem 1 and its proof hold under the assumption that $S \subset \mathbb{R}^n_+$ is any compact set. The algorithm can also be adapted, but generally, finding the k^*-center of a multidimensional search space is tractable only when S is finite and sufficiently small.

7.3 Discontinuous quality indices

In some settings, quality may not vary continuously in observable attributes. For example, two tennis rackets with similar frames may feel different due to factors that become apparent only after using them.

We can capture the idea that observably similar have similar qualities without imposing continuity. Let $\epsilon, \delta > 0$. Let $\mathcal{Q}_{\epsilon,\delta}$ be the set of all quality indices $q : S \to [0,1]$ satisfying *uniform local boundedness*: if $x, y \in S$ and $|x - y| \leq \delta$, then $|q(x) - q(y)| \leq \epsilon$. This is a superset of Q_L for $L = \frac{\epsilon}{\delta}$ that allows for jump discontinuities. Theorem 1 would generalize when $\Omega = \mathcal{Q}_{\epsilon,\delta}$; see Figure 6.

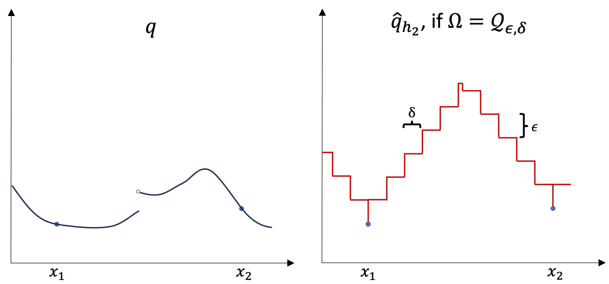

Figure 6: On the left is the true quality index. On the right is a worst-case quality index for a strategy that stops at h_2: it maximizes the quality of the best unexplored alternative under the uniform local boundedness constraint.

8 Conclusion

We study how forward-looking agents explore an unfamiliar set of items. They know little about how item qualities vary with observable attributes but use what they learn about this relationship to guide future searches. We characterize optimal search and rationalize choice overload.

Directed search with learning is a ubiquitous but intractable problem. Economists have proposed studying heuristics, arguing that the complexity of search and bandit problems with learning raises doubts about the positive content of rational theories (Radner, 1975; Francetich and Kreps, 2020).

Here, we insist on optimality, but we change the model to reflect the unstructured environments in which people often search. People may not know how exactly the qualities of various items are jointly correlated. Yet, they may still guess that similar items have similar qualities. In such settings, strategies that maximize guaranteed payoffs are tractable and make intuitive predictions about search. A similar approach may be useful in other complex learning problems as an alternative to Bayesian or boundedly rational models.

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.

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

Chade, Hector and Lones Smith, "Simultaneous search," *Econometrica*, 2006, *74* (5), 1293–1307.

Chernev, Alexander, Ulf Böckenholt, and Joseph Goodman, "Choice overload: A conceptual review and meta-analysis," *Journal of Consumer Psychology*, 2015, *25* (2), 333–358.

Epstein, Larry G and Martin Schneider, "Recursive multiple-priors," *Journal of Economic Theory*, 2003, *113* (1), 1–31.

Francetich, Alejandro and David Kreps, "Choosing a good toolkit, I: Prior-free heuristics," *Journal of Economic Dynamics and Control*, 2020, *111*, 103813.

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

Griffin, Jill G and Susan M Broniarczyk, "The slippery slope: The impact of feature alignability on search and satisfaction," *Journal of Marketing Research*, 2010, *47* (2), 323–334.

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.

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

Ke, Shaowei and Qi Zhang, "Randomization and ambiguity aversion," *Econometrica*, 2020, *88* (3), 1159–1195.

Lattimore, Tor and Csaba Szepesvári, *Bandit algorithms*, Cambridge University Press, 2020.

Libgober, Jonathan and Xiaosheng Mu, "Coasian Dynamics under Informational Robustness," *arXiv preprint arXiv:2202.04616*, 2022.

Manski, Charles F, "Using Limited Trial Evidence to Credibly Choose Treatment Dosage when Efficacy and Adverse Effects Weakly Increase with Dose," *arXiv preprint arXiv:2305.17206*, 2023.

Radner, Roy, "Satisficing," in "Optimization Techniques IFIP Technical Conference: Novosibirsk, July 1–7, 1974" Springer 1975, pp. 252–263.

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

Saito, Kota, "Preferences for flexibility and randomization under uncertainty," *American Economic Review*, 2015, *105* (3), 1246–1271.

Scheibehenne, Benjamin, Rainer Greifeneder, and Peter M Todd, "Can there ever be too many options? A meta-analytic review of choice overload," *Journal of consumer research*, 2010, *37* (3), 409–425.

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

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 Preferences and dynamic consistency

Section 2 describes the searcher's preferences over strategies at any history. A strategy $\sigma \in \Sigma$ and a history $h \in H$ determine the terminal payoff $p(h(\sigma, q), q)$ that the searcher gets upon stopping, for any $q \in \Omega_h$.

We extend the searcher's preferences $\{\succsim_h\}_{h \in H}$ to the set of all *acts* $f : \Omega \to \Delta_s \mathbb{R}$ from states to simple lotteries. For any $f, g \in \mathcal{A}$ and history $h \in H$,

$$f \succsim_h g \iff \inf_{q \in \Omega_h} \mathbb{E}[f(q)] \geq \inf_{q \in \Omega_h} \mathbb{E}[g(q)],$$

with $\succ_h$ and $\sim_h$ denoting the asymmetric and symmetric parts of $\succsim_h$. The set of inconsistent states $\Omega \setminus \Omega_h$ is $\succsim_h$-*null* in that $f \sim_h g$ if $f(\cdot) = g(\cdot)$ on Ω_h.

Preferences $\{\succsim_h\}_{h \in H}$ satisfy a weak form of dynamic consistency. Let $h{+}(x, z) \in H$ be the history where $x \in S$ is searched at $h \in H$ and has quality $z \in [0, 1]$. Let $\mathcal{Z}_x = \{z \in [0, 1] | h + (x, z) \in H\}$.

Proposition 7 (W-DC). *Let $h \in H$ and $x \in S \setminus X_h$. If $f \succsim_{h+(x,z)} g$ for all z such that $h{+}(x, z) \in H$, then $f \succsim_h g$.*

Proof. By assumption,

$$\inf_{q \in \Omega_{h+(x,z)}} \mathbb{E}[f(q)] \geq \inf_{q \in \Omega_{h+(x,z)}} \mathbb{E}[g(q)],$$

for each $z \in \mathcal{Z}_x$, and

$$\Omega_h = \bigcup_{z \in \mathcal{Z}_x} \Omega_{h+(x,z)},$$

so,

$$\inf_{q \in \Omega_h} \mathbb{E}[f(q)] \geq \inf_{q \in \Omega_h} \mathbb{E}[g(q)].$$

□

There are differences between our framework and that of Epstein and Schneider (2003), but preferences here satisfy W-DC for the reason they identify. The searcher can be interpreted as having a "rectangular" set of priors that she updates prior-by-prior. For brevity, we highlight the connection without defining a probability space and filtration. The set of priors at $h \in H$ is

$$C_h = \{\delta_q | q \in \Omega_h\},$$

where δ_q is the Dirac delta measure on Ω concentrated at $q \in \Omega_h$. The set of posteriors at $h + (x, z)$ for $z \in \mathcal{Z}_x$ is

$$C_{h+(x,z)} = \{\delta_q | q \in \Omega_{h+(x,z)}\}.$$

If $q \in \Omega_{h+(x,z)}$, then the prior δ_q at h is also the posterior at $h+(x, z)$, by Bayes rule. If $q \notin \Omega_{h+(x,z)}$, it is dropped from $C_{h+(x,z)}$. Therefore, C_h is *rectangular*:

$$C_h = \left\{ \int_{z \in \mathcal{Z}_x} \delta_q(\Omega_{h+(x,z)}) \delta_{q_z} : q \in C_h, q_z \in C_{h+(x,z)} \right\}$$

In Epstein and Schneider (2003), rectangularity and prior-by-prior updating implies full dynamic consistency: in addition to W-DC, if $f \succ_{h+(x,z)} g$ for some $z \in \mathcal{Z}$, then $f \succ_h g$. This does not hold in our model. The discrepancy is because Epstein and Schneider (2003) assume that priors have full support. Rectangularity then implies that some worst-case belief guides the searcher's choices. Here, the priors are Dirac measures, so choices are driven by the same worst-case belief until it is shown to be inconsistent at some $h \in H$. Then, a new worst-case belief takes over. But because h is a zero-probability event under the previous worst-case belief, any plan following h can be part of an optimal strategy at earlier histories. Therefore, sequentially optimal strategies exist, though not every ex ante optimal strategy is sequentially optimal.

B When randomizing helps the searcher

Consider the extension where the searcher can use mixed strategies and solves:

$$\max_{\sigma \in \Delta(\Sigma)} \big\{ \inf_{\tilde{q} \in \Omega} \mathbb{E}[p(h_0(\sigma, \tilde{q}), \tilde{q})] \big\}.$$

We show by example that, sometimes, an optimal strategy must be mixed.

Suppose the search space consists of two items, $S = \{0, 1\}$. Let $L = 0.5$ and $\Omega = [0, 1]^S$. Let $c = 0.3$ and let payoffs be given by linear regret:

$$U(q(x), \max_{q \in \Omega} q(y)) = q(x) - \max_{q \in \Omega} q(y).$$

A strategy σ that stops at h_0 can incur a regret of up to 1 (e.g., if $q(0) = q(1) = 1$) but incurs no search costs. The worst-case net payoff of σ is -1.

The strategy σ that searches at h_0 and stops at any subsequent history can incur a regret of up to 0.5 (e.g., if $\sigma(h_0) = 0$ and $q(0) = 0$ while $q(1) = 0.5$) and a search cost of 0.3. The worst-case net payoff of σ is -0.8.

A $\sigma \in \Sigma$ that searches twice for some $q \in \Omega$ has no regret and 0.6 in search costs if the state is q. The worst-case net payoff for such a σ is at most -0.6.

Therefore, the strategy σ that searches at h_0 and searches at every h_1 is ex ante optimal in Σ, but there are mixed strategies that do better.

Consider a random strategy σ that searches once at h_0 and stops afterwards, so it incurs a search cost of 0.3. At h_0, $\sigma(h_0) = 0$ or $\sigma(h_0) = 1$, each with probability 0.5. If $q(0) = q(1)$, then σ has zero regret. If $q(0) \neq q(1)$, then σ has a regret of 0 or a regret of up to 0.5, each with probability 0.5 (e.g, if $q(0) = 0$ and $q(1) = 0.5$). Therefore, the worst-case net payoff of σ is -0.55.

C Proofs

C.1 Proposition 1 and Theorem 1 when $\Omega = Q_L$

In this section, we maintain the assumption that $\Omega = Q_L$.

Let $h_t = \{(x_i, z_i)\}_{i=1}^t \in H$. For each $x_i \in X_{h_t}$ and $y \in S$, let

$$f_{h_t,x_i}(y) = L||y - x_i|| + z_i.$$

For each $y \in S$, let

$$f_{h_t}(y) = \min_{x \in X_{h_t}} f_{h_t,x}(y).$$

LEMMA **1.** *For any $h_t \in H$, $\hat{q}_{h_t} \in \Omega_{h_t}$. Moreover, $\hat{q}_{h_t}(x) = \min\{f_{h_t}(x), 1\}$ for all $x \in S$.*

Proof of Lemma 1. Let $g(x) \equiv \min\{f_{h_t}(x), 1\}$ for all $x \in S$. We proceed by proving three claims.

Claim 1: g 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 so g, is L-Lipschitz.

Claim 2: Every quality index in Ω_{h_t} is bounded above pointwise by g.

Let $h_t = \{(x_i, z_i)\}_{i=0}^{t-1} \in H$, and let $\tilde{q} \in \Omega_{h_t}$. For any $x_i \in T$ and $y \in S$, $|\tilde{q}(y) - \tilde{q}(x_i)| \leq L||y - x_i||$; Moreover, note that $|f_{h_t,x_i}(y) - f_{h_t,x_i}(x_i)| = f_{h_t,x_i}(y) - \tilde{q}(x_i) = L||y - x_i||$. Together, this implies that $\tilde{q}(y) \leq f_{h_t,x_i}(y)$ for all $x_i \in T$. Therefore, $\tilde{q}(y) \leq \min\{f_{h_t}(y), 1\}$ for all $\tilde{q} \in \Omega_{h_t}$.

Claim 3: g is consistent at h_t.

For any $x_i \in X_{h_t}$ and $\tilde{q} \in \Omega_{h_t}$, $z_i = \tilde{q}(x_i) \leq f_{h_t}(x_i) \leq f_{h_t,x_i}(x_i) = z_i$, 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 g is consistent.

The first and third claims imply that $g \in \Omega_{h_t}$. So by the second claim, $\hat{q}_{h_t} = g$. □

Lemma **2.** *Let histories* $h'_t = \{(x_i, z'_i)\}_{i=1}^{t}$ *and* $h''_t = \{(x_i, z''_i)\}_{i=1}^{t}$ *be such that* $z'_i \geq z''_i$ *for all* i. *Then*

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

Proof of Lemma 2. Let $T \equiv X_{h'_t} = X_{h''_t}$ be the set of searched items 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.** *For any* $h \in H$,

$$\inf_{\tilde{q} \in \Omega_h} p(h, \tilde{q}) = p(h, \hat{q}_h).$$

Proof.

$$\begin{aligned}
\inf_{\tilde{q}\in\Omega_h} p(h,\tilde{q}) &= \inf_{\tilde{q}\in\Omega_h} U(z_h^*, \max_{x\in S} \tilde{q}(x)) - c\cdot t \\
&= U(z_h^*, \max_{x\in S, \tilde{q}\in\Omega_h} \tilde{q}(x)) - c\cdot t \\
&= U(z_h^*, \max_{x\in S} \hat{q}_h(x)) - c\cdot t \\
&= p(h, \hat{q}_h),
\end{aligned}$$

where the second equality is because U is decreasing in its second argument, and the third equality is by Lemma 1. □

LEMMA **4.** *Suppose that* $\Omega = Q_L$. *Fix some history* $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1} \in H$ *and let* x *be a best discovery at* h'_t. *Let* $T \equiv X^*_{h'_t} \setminus \{x\}$ *be the remaining best discoveries. 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* $\Omega_{h'_t}$ *and* $\Omega_{h''_t}$ *are nonempty. Then,*

$$\inf_{\tilde{q}\in\Omega_{h''_t}} p(h''_t, \tilde{q}) \geq \inf_{\tilde{q}\in\Omega_{h'_t}} p(h'_t, \tilde{q}).$$

Proof of Lemma 4. If $T = \emptyset$, the statement is trivially true. Suppose $T \neq \emptyset$.

Case 1: If $\epsilon_i \leq 0$ for each $x_i \in T$, then x is a best discovery at h''_t and h'_t, so

$$z^*_{h''_t} = z^*_{h'_t},$$

and by Lemma 2,

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

Then, by Assumption 2,

$$p(h''_t, \hat{q}_{h''}) \geq p(h'_t, \hat{q}_{h'}).$$

Therefore, by Lemma 3,

$$\inf_{\tilde{q}\in\Omega_{h''_t}} p(h''_t, \tilde{q}) \geq \inf_{\tilde{q}\in\Omega_{h'_t}} p(h'_t, \tilde{q}).$$

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 history $h'''_t = \{(x_i, z'_i + \epsilon)\}_{i=0}^{t-1}$ where the quality of *all* items searched in h'''_t is higher than in h'_t by ϵ. Then,

$$\max_{x\in S} \hat{q}_{h'''_t}(x) \leq \max_{x\in S} \hat{q}_{h'_t}(x) + \epsilon.$$

By Assumption 2, $U(a+\epsilon, b+\epsilon) \geq U(a, b)$ for any $a, b \in \mathbb{R}_+$, so

$$p(h'''_t, \hat{q}_{h''}) \geq p(h'_t, \hat{q}_{h'}).$$

Now, by construction,

$$z^*_{h''_t} = z^*_{h'''_t}.$$

So by Lemma 2 and Assumption 2 again,

$$p(h''_t, \hat{q}_{h''_t}) \geq p(h'''_t, \hat{q}_{,\hat{q}_{h'''_t}}).$$

Therefore, $p(h''_t, \hat{q}_{h''}) \geq p(h'_t, \hat{q}_{h'})$, so by Lemma 3,

$$\inf_{\tilde{q}\in\Omega_{h''_t}} p(h''_t, \tilde{q}) \geq \inf_{\tilde{q}\in\Omega_{h'_t}} p(h'_t, \tilde{q}).$$

□

Proof of Proposition 1. Let $\sigma \in \Gamma_h$, let $h' \equiv h(\sigma, \bar{q}_h)$, and let $h'' \equiv h(\sigma, q)$ for some $q \in \Omega_h$. Note that $X_{h'} = X_{h''}$, because σ is a simultaneous search strategy. Therefore, by Lemma 4,

$$\inf_{\tilde{q}\in\Omega_{h''}} p(h'', \tilde{q}) \geq \inf_{\tilde{q}\in\Omega_{h'}} p(h', \tilde{q}).$$

By Lemma 3,

$$\inf_{\tilde{q}\in\Omega_{h'}} p(h',\tilde{q}) = p(h',\hat{q}_{h'}).$$

Recall that $q_h^\sigma \equiv \hat{q}_{h'}$, so

$$p(h',\hat{q}_{h'}) = p(h(\sigma,q_h^\sigma),q_h^\sigma).$$

Combining the previous steps,

$$\inf_{\tilde{q}\in\Omega_h} p(h(\sigma,\tilde{q}),\tilde{q}) \geq p(h(\sigma,q_h^\sigma),q_h^\sigma)$$

By Lemma 1, q_h^σ is consistent at $h \in H$, so

$$\inf_{\tilde{q}\in\Omega_h} p(h(\sigma,\tilde{q}),\tilde{q}) = p(h(\sigma,q_h^\sigma),q_h^\sigma).$$

It only remains to be shown that an optimal simultaneous search strategy exists. Let $\bar{k}_h \in \mathbb{N}$ be such that $(\bar{k}_h + 1)\cdot c > 1 - z_h^*$. Let $\Gamma_h' \subset \Gamma_h$ be the strategies σ_s that search no more than $\bar{k}_h$ times after h, i.e., the set of simultaneous strategies at h where $|X_{h(\sigma_s,\cdot)}| - |X_h| \leq \bar{k}_h$. Then by construction,

$$\max_{\sigma_s\in\Gamma_h} p(h(\sigma_s,\tilde{q}),\tilde{q}) = \max_{\sigma_s\in\Gamma_h'} p(h(\sigma_s,\tilde{q}),\tilde{q}).$$

Note that Γ_h' is compact, as each strategy in Γ_h' can be identified with an element in $\prod_{i=1}^{\bar{k}_x} S$. Next, p is continuous in both its arguments (under, say, the topology induced by the sup norm metric on the space of quality indices). Therefore, optimal simultaneous search strategies exist at any $h \in H$.

□

LEMMA **5.** *If $\sigma_s^* \in \Gamma_h$ is an optimal simultaneous search strategy at $h \in H$, then*

$$\sigma_s^* \in \arg\max_{\sigma\in\Sigma} \big\{ \inf_{\tilde{q}\in\Omega_h} p(h(\sigma,\tilde{q}),\tilde{q}) \big\}.$$

Proof. Because $\Gamma_h \subset \Sigma$,

$$\max_{\sigma' \in \Gamma_h} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}) \leq \sup_{\sigma' \in \Sigma} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}),$$

so it only remains to show the reverse inequality.

Let $\sigma \in \Sigma$ be any strategy. Let $\sigma_s \in \Gamma_h$ be such that $h(\sigma, q_h^\sigma) = h(\sigma_s, q_h^\sigma)$; such a simultaneous search strategy can be constructed by searching the same items as σ on path, starting from h when $q = q_h^\sigma$. Then

$$p(h(\sigma, q_h^\sigma), q_h^\sigma) = p(h(\sigma_s, q_h^\sigma), q_h^\sigma) = \inf_{\tilde{q} \in \Omega_h} p(h(\sigma_s, \tilde{q}), \tilde{q}) \leq \max_{\sigma' \in \Gamma_h} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}),$$

where the second equality is by Proposition 1. Therefore,

$$\inf_{\tilde{q} \in \Omega_h} p(h(\sigma, \tilde{q}), \tilde{q}) \leq \max_{\sigma' \in \Gamma_h} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}).$$

Because σ is an arbitrary strategy in Σ,

$$\sup_{\sigma' \in \Sigma} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}) \leq \max_{\sigma' \in \Gamma_h} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}),$$

from which the result follows. □

LEMMA **6.** *If $\sigma \in \Sigma$ is sequentially optimal, then σ follows optimal simultaneous search strategies.*

Proof. Suppose σ does not follow optimal simultaneous search strategies, i.e., there is a $h \in H$ such that $\sigma(h) \neq \sigma_s^*(h)$ for any optimal $\sigma_s^* \in \Gamma_h$.

Let $\sigma_s \in \Gamma_h$ be such that $h(\sigma, q_h^\sigma) = h(\sigma_s, q_h^\sigma)$; such a simultaneous search strategy can be constructed by searching the same items as σ on path, starting from h when $q = q_h^\sigma$. Then

$$p(h(\sigma, q_h^\sigma), q_h^\sigma) = p(h(\sigma_s, q_h^\sigma), q_h^\sigma) < \max_{\sigma' \in \Gamma_h} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}).$$

Therefore, an optimal simultaneous strategy obtains a strictly larger worst-case payoff at h than σ. Such a strategy exists by Proposition 1, contradicting the assumption that σ is sequentially optimal. □

Lemma 7. *If $\sigma \in \Sigma$ follows optimal simultaneous search strategies, then σ is sequentially optimal.*

Proof. Let $h \in H$ and $q \in \Omega_h$. If h is a terminal history, let $\hat{h} \equiv h$. If h is a non-terminal history under σ, let $\hat{h}$ be the history following h where $\sigma(h)$ is searched and has quality $q(\sigma(h))$. Let $A = \{\sigma' \in \Gamma_h | \sigma'(h) = \sigma(h)\}$. Then,

$$\begin{aligned}
\max_{\sigma' \in \Sigma} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}) &= \max_{\sigma' \in \Gamma_h} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}) \\
&= \max_{\sigma' \in A} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}) \\
&\leq \max_{\sigma' \in \Gamma_{\hat{h}}} \inf_{\tilde{q} \in \Omega_{\hat{h}}} p(\hat{h}(\sigma', \tilde{q}), \tilde{q}) \\
&\leq \max_{\sigma' \in \Sigma} \inf_{\tilde{q} \in \Omega_{\hat{h}}} p(\hat{h}(\sigma', \tilde{q}), \tilde{q}),
\end{aligned}$$

where the first equality is by Lemma 5, the second equality is by the definition of following optimal simultaneous strategies, the first inequality is because $\Omega_{\hat{h}} \subset \Omega_h$, and the second second inequality is because $\Gamma_{\hat{h}} \subset \Sigma$. Therefore, the worst-case payoff to σ is weakly increasing on path, for any $h \in H$ and $q \in \Omega_h$.

Suppose that for some $h \in H$ and $q \in \Omega_h$, σ does not reach a terminal history. Then the worst-case payoff must decrease at some history on path when following σ, because payoff tends to $-\infty$ if the searcher never stops, a contradiction. Therefore, σ terminates, so $\sigma \in \Sigma$.

We conclude that,

$$\sigma \in \arg\max_{\sigma' \in \Sigma} \inf_{\tilde{q} \in \Omega_h} p(h(\sigma', \tilde{q}), \tilde{q}).$$

Because this holds for any $h \in H$, σ is sequentially optimal. □

Proof of Theorem 1. Lemma 6 and Lemma 7 together imply that $\sigma \in \Sigma$ is sequentially optimal if and only if it follows optimal simultaneous search strategies. Proposition 1 implies that optimal simultaneous search strategies exist at any $h \in H$. Therefore, there also exist strategies that follow optimal simultaneous search strategies. $\square$

C.2 Theorem 1 in the $\Omega = Q_L^{QC}$ case

In this section, we maintain the assumption that $\Omega = Q_L^{QC}$. The analogues of Lemma 1 and Lemma 2 no longer hold in this case.

To see this, consider the following counter-example: Let $S = [0, 4]$ and $L = 1$. Denote by h_3 the 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 Ω_{h_3} is a saw-tooth shaped function and therefore not quasiconcave.

Next, note that the highest possible quality for some technology under some $q \in \Omega_{h_3}$ 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 $\Omega_{h_3'}$ is quasiconcave, it must now be the case that $q'(1) = 0$ for every $q' \in \Omega_{h_3}$. The highest possible quality for some technology under some $q' \in \Omega_{h_3'}$ 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, weaker forms of Lemma 1 and Lemma 2 hold and suffice for the proof of Theorem 1 when $\Omega = Q_L^{QC}$. Recall that $\bar{q}_{h_t} \equiv \min\{\hat{q}_{h_t}, z^*_{h_t}\}$.

Lemma 8. *For all $h_t \in H$, $\bar{q}_{h_t} \in \Omega_{h_t}$.*

Proof. The argument that $\bar{q}_{h_t}$ is L-Lipschitz and consistent at h_t is exactly as in the proof of Lemma 1. It only remains to be shown that $\bar{q}_{h_t}$ is quasiconcave.

If $\tilde{q} \in \Omega_{h_t}$, then $\min\{\tilde{q}, z^*_{h_t}\}$ is non-decreasing on $[0, \min X^*_{h_t})$ and non-increasing on $(\max X^*_{h_t}, 1]$, and equals $z^*_{h_t}$ on $[\min X^*_{h_t}, \max X^*_{h_t}]$. Therefore, the same is true for $\bar{q}_{h_t}$, so it is quasiconcave. □

Lemma 9. *Suppose histories $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1}$ and $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1}$ are such that $z'_i \geq z''_i$ for all i and $z^*_{h'_t} = z^*_{h''_t}$. Then*

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

Proof of Lemma 9. First, note that $X^*_{h''_t} \subset X^*_{h'_t}$. For $x \in [\min X^*_{h'_t}, \max X^*_{h'_t}]$, then, $\bar{q}_{h'_t}(x) = z^*_{h'_t} = z^*_{h''_t} \geq \bar{q}_{h''_t}(x)$. Next, it follows from Lemma 8 that $\bar{q}_{h'_t}$ and $\bar{q}_{h''_t}$ are non-decreasing on $[0, \min X^*_{h'_t})$ and non-increasing on $(\max X^*_{h'_t}, 0]$. Therefore, $\max\{\bar{q}_{h'_t}, \bar{q}_{h''_t}\} \in \Omega_{h'_t}$, which implies that $\bar{q}_{h'_t} = \max\{\bar{q}_{h'_t}, \bar{q}_{h''_t}\} \geq \bar{q}_{h''_t}$.

For any $\tilde{q} \in \Omega_{h''_t}$,

$$\begin{aligned}\{x \in [0,1] | \tilde{(q)}(x) \geq z^*_{h'_t}\} &\subset \{x \in [0,1] | \bar{q}_{h''_t}(x) = z^*_{h''_t}\} \\ &\subset \{x \in [0,1] | \bar{q}_{h'_t}(x) = z^*_{h'_t}\},\end{aligned}$$

because $\bar{q}_{h'_t} \geq \bar{q}_{h''_t}$. Moreover, these sets are intervals by Lemma 8. Therefore, $q \equiv \max\{\tilde{q}, \bar{q}_{h'_t}\}$ is quasiconcave, and $q \in \Omega_{h'_t}$. Therefore, $\tilde{q} \leq q$ pointwise, which proves the result.

□

Proof of Proposition 1. Let $\sigma \in \Gamma_h$. The only difference with the proof in the case where Ω Q_L is that q_h^σ is not consistent at $h \in H$, because it is not necessarily quasiconcave. Therefore, Lemma 1 cannot be invoked.

However, $h(q_h^\sigma, \sigma) = h(\bar{q}_h, \sigma)$, so $\Omega_{h(q_h^\sigma, \sigma)}$ is nonempty. Let $q \in \Omega_{h(q_h^\sigma, \sigma)}$ be a quality index with $\max_{x \in S} q(x) = \max_{x \in S} q_h^\sigma(x)$. Then $p(h(\sigma, q_h^\sigma), q_h^\sigma) = p(h(\sigma, q_h^\sigma), q)$. Therefore, the worst-case payoff to any simultaneous search strategy is $p(h(\sigma, q_h^\sigma), q_h^\sigma)$. □

Proof of Theorem 1 in the $\Omega = Q_L^{QC}$ case. The proofs of the analogous lemmas to those in Appendix C.1 are identical, with Lemma 9 in place of Lemma 2 whenever the latter is referenced. □

C.3 Proofs for Section 4

Proof of Proposition 2. Search outside of S_h is wasteful. Suppose that searching at $x' \in S_h \setminus P_h$ reveals that $q(x') = z_h^*$. At this history, h',

$$\max_{x \in S} \hat{q}_{h'}(x) = \max_{x \in S} \hat{q}_{h'}(x),$$

so the payoff to stopping at h is better than the worst-case payoff to stopping after searching once in $S \setminus P_h$. □

C.4 Proofs for Section 5

Proof of Proposition 3. If $U_2 = 0$, search stops immediately for any level of search complexity, so for the remainder of the proof, assume $U_2 < 0$.

First we construct $\underline{L}$. If $L = 0$ and Ω_h is nonempty, then clearly there is no value in search, as Ω_h is a singleton containing only a constant function. Let $\epsilon > 0$ be small enough so that $U(z_h^*, z_h^*) - U(z_h^*, z_h^* + \epsilon) < c$. Because S is compact, there exists $\underline{L} > 0$ small enough so that $\hat{q}_h < z_h^* + \epsilon$ when $L \leq \underline{L}$. Let

$L \leq \underline{L}$ and let $\sigma \in \Sigma$ be any strategy such that $\sigma(h) \neq \emptyset$. Then by Theorem 1 and by Proposition 1,

$$\begin{aligned}
\inf_{\tilde{q} \in \Omega_h} p(h(\sigma, \tilde{q}), \tilde{q}) &= p(h(\sigma, \bar{q}_h), q_h^\sigma) \\
&\leq U(z_h^*, z_h^*) - c \\
&< U(z_h^*, z_h^* + \epsilon) \\
&< p(h, \max_{x \in S} \hat{q}_h(x)).
\end{aligned}$$

Therefore, concluding search is optimal at h if $L \leq \underline{L}$.

Next we construct $\overline{L}$. Let $\delta = U(1,1) - U(z_h^*, 1)$. Let $n = \lceil \frac{\delta}{c} \rceil$. Consider any strategy $\sigma \in \Sigma$ for which there is a $q \in \Omega_h$ such that $|X_h(\sigma, q)| - |X_h| \geq n$. Then

$$\begin{aligned}
\inf_{\tilde{q} \in \Omega_h} p(h(\sigma, \tilde{q}), q_h^\sigma) &\leq p(h(\sigma, q), q) \\
&\leq U(1,1) - n \cdot c \\
&< U(z_h^*, 1).
\end{aligned}$$

Therefore, any strategy which searcher n or more times for some $q \in \Omega_h$ cannot be optimal at h. Let $\overline{L}$ be such that if $L \geq \overline{L}$, for any $\sigma \in \Sigma$ such that $|X_h(\sigma, h(\sigma, q_h^\sigma))| - |X_h| < n$,

$$\max_{x \in S} \hat{q}_{h(\sigma, q_h^\sigma)}(x) = 1.$$

For such L, the searcher is better off stopping immediately. □

When considering comparative statics with respect to different levels of search complexity, say L' and L'', we subscript variables to indicate the level.

Lemma **10.** *Let* $0 < L' < L''$. *Let* $h = \{(x_i, z)\}_{i=0}^t \in H$ *be an on-path history, and* $\max_{x \in S} \hat{q}_{h,L''}(x) < 1$. *Let* $y \in S / X_h$, *and let* $h' = h \cup \{(y, z)\}$ *be the* *on-path history at which y is searched. Then,*

$$\max_{x\in S}\hat{q}_{h,L'}(x)-\max_{x\in S}\hat{q}_{h',L'}(x)\le\max_{x\in S}\hat{q}_{h,L''}(x)-\max_{x\in S}\hat{q}_{h',L''}(x).$$

If $\max_{x\in S}\hat{q}_{h,L'}(x)-\max_{x\in S}\hat{q}_{h',L'}(x)>0$, *the preceding inequality is strict.*

Proof. Because $\Omega_{h'}\subset\Omega_h$, $\max_{x\in S}\hat{q}_{h,L''}(x)-\max_{x\in S}\hat{q}_{h',L''}(x)\ge 0$. So, the result holds when $\max_{x\in S}\hat{q}_{h,L'}(x)-\max_{x\in S}\hat{q}_{h',L'}(x)=0$.

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 items to the left and right of y (and $\min S$ or $\max S$, respectively, if there are no such items). Define $\underline{x_s},\overline{x_s}'\in X_h$ similarly as the endpoints in $X_h\cup\{\min S,\max S\}$ of the subinterval containing the second largest peak of q^u_h.

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

It is readily verified (for example, by Lemma 1 and an analogous result for the $\Omega=Q^{QC}_L$ 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', $\max_{x\in S}\hat{q}_{h',L'}(x)=g(L')$ and $\max_{x\in S}\hat{q}_{h',L''}(x)=g(L'')$.

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 $f(L')-g(L')\le f(L'')-g(L'')$, which is the desired result.

Case 2: At history h',

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

and

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

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

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

and

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

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

Case 3: At history h',

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

and

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

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

LEMMA **11.** *Suppose there are decreasing returns to search. Let $0 < L' < L''$, $h \in H$ be on-path, and $\hat{q}_{h,L''} < 1$. If search stops at h under some optimal strategy at h when search complexity is L'', then search stops at h under any optimal strategy at h when search complexity is L'.*

Proof of Lemma 11. Suppose for contradiction that there is an optimal strategy at h, say σ, that does not stop at h when complexity is L'. Let $h' \equiv h(\sigma, q_h^\sigma)$.

Now, $\max_{x\in S} \hat{q}_{h,L'}(x) - \max_{x\in S} \hat{q}_{h',L'}(x) > 0$, or else stopping at h would have been a strict improvement. But then by Lemma 10 and induction on the number of searches,

$$\max_{x\in S} \hat{q}_{h,L'}(x) - \max_{x\in S} \hat{q}_{h',L'}(x) < \max_{x\in S} \hat{q}_{h,L''}(x) - \max_{x\in S} \hat{q}_{h',L''}(x). \tag{5}$$

Next, it is obvious that

$$\max_{x\in S} \hat{q}_{h,L''}(x) \geq \max_{x\in S} \hat{q}_{h,L'}(x). \tag{6}$$

Finally, because both histories are on path,

$$z_h^* = z_{h'}^*. \tag{7}$$

Putting eq. (5), eq. (6) and eq. (7) together, using that $U_2 \leq 0$ and $U_{22} < 0$,

$$\begin{aligned} &U(z_{h'}^*, \max_{x\in S} \hat{q}_{h',L'}(x)) - U(z_h^*, \max_{x\in S} \hat{q}_{h,L'}(x)) \\ &< U(z_{h'}^*, \max_{x\in S} \hat{q}_{h',L''}(x)) - U(z_h^*, \max_{x\in S} \hat{q}_{h,L''}(x)). \end{aligned}$$

So at search complexity L'', no optimal strategy at h stops at h, contradiction.

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

LEMMA **12.** *Let $0 < L' < L''$, $h \in H$, and $\hat{q}_{h,L'} = 1$. If search stops at h under some optimal strategy at h when search complexity is L', then search stops at h under any optimal strategy at h when search complexity is L''.*[2]

Proof of Lemma 12. Suppose for contradiction that there is an optimal strategy at h, say σ, that does not stop at h when complexity is L''. Let $h' \equiv h(\sigma, q_h^\sigma)$.

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

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

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

$$U(z_{h'}^*, \max_{x\in S} \hat{q}_{h',L'}(x)) - U(z_h^*, 1) > U(z_{h'}^*, \max_{x\in S} \hat{q}_{h',L''}(x)) - U(z_h^*, 1).$$

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

[2]We need not assume h is on path or that there are decreasing returns to search.

Proof of Proposition 4, parts 1 and 2. Let h be some on-path history.

If search stops 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 stop immediately in some equilibrium. Define L^τ such that when $L < L^\tau$, $\max_{x \in S} \hat{q}_{h,L}(x)) < 1$ and when $L > L^\tau$, $\max_{x \in S} \hat{q}_{h,L}(x)) = 1$.

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

Let $\underline{L} = \inf \mathcal{L}$. By Proposition 3, $\underline{L} > 0$. Moreover, by Lemma 12, $\underline{L} \leq L^\tau$. Finally, by Lemma 11, 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 11, $\overline{L} \geq L^\tau$. By Lemma 12, if $\overline{L} > L^\tau$, then search continues in all any equilibrium for $L \in [L^\tau, \overline{L})$. □

Proof of Proposition 4, part 3. Let $L' < L''$, and let σ' and σ'' be optimal search strategies at h at complexity levels L' and L'', respectively. Let $h' \equiv h(\sigma', q_h^{\sigma'})$ and $h'' \equiv q_h^{\sigma''}$. Let $|X_{h',L'}| - |X_{h,L'}| = k'$, and $|X_{h'',L''}| - |X_{h,L''}| = k''$. For contradiction, suppose that $k' > k'' > 0$.

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

Next, under the constraint of searching exactly k more times, an optimal strategy at h when complexity is L'' is also optimal when complexity is L'. This is easy to see, for example, from the description of an optimal search algorithm in Section 3.5.

By an argument analogous to that in the proof of Lemma 11,

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

where the first inequality is by the optimality of concluding at h' over searching fewer times and concluding at h''.[3] This 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 stop at h'' when complexity is L''. □

C.5 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 Ω_{h_t} and let $z' \equiv q^l_{h_t}(\sigma(h_t))$. Note that by definition, $\Omega_{h'_{t+1}}$ is nonempty, and $z_{t+1} \geq z'$. By construction and the assumption that $\min_{\tilde{q} \in \Omega_{h_t}, y \in S} \tilde{q}(y) > 0$,

$$\hat{q}_{h_{t+1}}(y) = L|\sigma(h_t) - y| + z' \leq \hat{q}_{h_t}(y),$$

for all $y \in [x-d, x+d]$, where $d = |x - \sigma(h_t)|$. By Proposition 1, any search in $[x-d, x+d]$ could not be a part of an optimal search strategy at h_t, 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 $q \in \Omega_{h_{t+1}}$. Otherwise, q is not quasiconcave. Moreover, $\hat{q}_{h_t}(x) \leq z_t + L(x - x_t)$ for any $x \in S$, by Lemma 1 and the fact that $Q^{QC}_L \subset Q_L$. Therefore, if $x < x_t + \frac{1}{L}(z_{t+1} - z_t)$, then $\hat{q}_{h_t}(x) < z_{t+1}$. By Proposition 1, any any search in $[\min S, x_t + \frac{1}{L}(z_{t+1} - z_t)]$ cannot be a part of a sequentially optimal strategy.

The proof of the remaining cases follow identical arguments. □

[3] Unlike in Lemma 11, h' need not follow h'', but this does not change the argument.