Retrospective Search: Exploration and Ambition on Uncharted Terrain*

Can Urgun[†] Leeat Yariv[‡]

This version: December 7, 2020, Latest version available here

Abstract

We study a model of retrospective search in which an agent—a researcher, an online shopper, or a politician—tracks the value of a product. Discoveries beget discoveries and their observations are correlated over time, which we model using a Brownian motion. The agent decides both how ambitiously, or broadly, to search, and for how long. We fully characterize the optimal search policy and show that it entails constant scope of search and a simple stopping boundary. We also show the special features that emerge from contracting with a retrospective searcher.

Keywords: Retrospective Search, Drawdown Stopping Boundary, Contracting.

JEL: C61, C73, D25, D83

*We thank Roland Benabou, Steven Callander, Doruk Cetemen, Alexander Frankel, Faruk Gul, Yingni Guo, Alessandro Lizzeri, Pietro Ortoleva, Wolfgang Pesendorfer, and Richard Rogerson for helpful comments and suggestions. We also thank seminar audiences at Caltech, University of Chicago, Princeton University, and Org Econ workshop 2020. We gratefully acknowledge financial support from the National Science Foundation through grants SES-1629613 and SES-1949381.
[†]Department of Economics, Princeton University.
[‡]Department of Economics, Princeton University, CEPR, and NBER.

1 Introduction

The Japanese philosophy of *kaizen* undergirds the search process many successful companies, including Toyota, Lockheed Martin, the Mayo Clinic, and Ford, take towards product releases. Kaizen literally means "change for the better," and suggests continuous improvement, whereby innovation occurs through incremental advances. This search entails two types of decisions: first, how much effort—workforce, innovation scope, and the like—to dedicate to the improvement; and second, when to declare an improvement sufficient for release.

Kaizen is a specific example of a broad set of search settings in which good outcomes depend on these two decisions, and where tomorrow's results depend on (are correlated with) today's advances. Discovery spurs discovery, achievement begets achievement. A governmental policy maker may revise a program incrementally or significantly until the right feature combination meets the challenge. Regional mineral exploration proceeds and halts based on discovery—geological survey teams choose their radius of operation throughout the process, with correlation across observations driven by the similarity of adjacent plots. A shopper looking for the best lightbulb on Amazon.com will begin her search with a simple query, then expand her search as related suggestions pop up until she finally buys.

In this paper, we consider a continuous search setting in which observations are correlated over time and a searcher makes two critical decisions: how broadly to search at any point, and when to cease search. Once search stops, the agent is rewarded for the best outcome observed throughout her search. We call this search process *retrospective search*. How ambitious should search be depending on where discoveries stand? When should search stop and what outcomes would it deliver? We characterize the optimal search policies and analyze how they affect agency relationships.

We consider a continuous-time setting in which a Brownian motion governs the process of discovery, as in Callander (2011). This structure naturally introduces intertemporal correlation between discoveries: one moment's observation forms the expectation for any future moment.[1] At every moment, the agent chooses her search scope or ambition, cap-

[1] In most search applications, the mere passage of time does not improve outcomes. Therefore, in our

tured by the momentary variance of the underlying process. We think of more ambitious agents—be it entrepeneurs, researchers, politicians, or shoppers—as ones that inspect a broad set of alternatives, captured by the diversity of outcomes they might observe. The cost of search is strictly positive and depends on search scope. While not strictly necessary, it is convenient to think of costs that are increasing and log-convex in search scope. Whenever the agent stops searching, she can recall all her observations retrospectively and pick the highest, which constitutes her search payoff.[2] To summarize, at each moment, the agent decides whether to continue searching and, if so, how intensely. Once she stops searching, she gets rewarded with the maximal observation from her search minus the cumulative search costs spent.

We show that the agent optimally chooses a constant search scope, maintaining an ambition level independent of her discoveries. That search scope depends on the search costs, with more log-convex costs associated with less ambitious search. In practice, this result implies that product development teams, political task forces, and the like need not change in volume or work intensity in response to findings, good or bad. Set properly, the optimal search scope need not respond to the ebb and flow of instantaneous discoveries. Indeed, research and development expenditures in many U.S. companies appear fairly constant over time. Similarly, while mineral exploration teams can choose the scope of well digging on their path, standard practices since the 50's dictate constant and pre-prescribed scope, independent of prior observations, see Zhilkin (1961).

The optimal stopping policy takes a fairly simple form. The agent stops searching whenever she sees sufficiently bad realizations. The threshold for stopping evolves over time. As the agent accumulates greater discoveries, she becomes more demanding and stops more willingly. Specifically, the agent halts exploration whenever an observation falls below some *fixed* distance from the current maximal observation, where that fixed distance is, again, constant over time. In particular, this result suggests that product development should stop when innovations do not appear sufficiently promising, where

benchmark model we assume no drift. Nonetheless, we describe impacts of drift in some of our discussions.

[2]This implicitly implies that the agent is risk neutral when it comes to her search outcomes. This simplifies our exposition greatly. We discuss how our analysis can be extended to risk-averse agents in Section 6.1 and in the Online Appendix.

promise is assessed in relation to the best option yet and does not change with time or with the realized discovery path.

Certainly, many search processes, particularly in the realm of research and development, occur within organizations. CEOs manage R&D teams, voters manage politicians, financiers contract with mining firms, etc. In the last part of the paper we therefore embed our retrospective-search model within a principal-agent interaction. We consider simple contracting instruments that we term *commission contracts*, reminiscent of the sharing rules first considered by Aghion and Tirole (1994). A commission contract entails a flow wage while the agent is searching, and a share of the ultimate maximal discovery value, the so-called commission. For example, university licensing agreements commonly include both fixed fees and royalties, see Jensen and Thursby (2001). Joint venture contracts between investors and mining companies often specify fixed or flow transfers, in addition to commissions on findings, see Root (1979).

We characterize the optimal contract, which inherits some of the features of the single-agent optimal search. In particular, the agent is induced to search at a constant scope. The optimal wage and commission depend on the level, marginals, and curvature of the search costs in a non-trivial manner. For the special case of linear search costs, we show that the agent searches with lower scope when under a commission contract relative to when searching alone. In fact, we show that contractual frictions come at a substantial cost of one ninth, or 11%, of the surplus.

Formally, retrospective search combines a stopping problem with a control problem. Such problems have received surprisingly limited attention in both the social sciences and the mathematics literature. The paper therefore also provides a methodological contribution by offering techniques for solving these types of problems, which can be useful for a variety of other contexts.

2 Related Literature

In what has become a classic paper, Weitzman (1979) introduced the basic search with recall model. In his model, an agent sequentially inspects a finite set of options with different reward distributions. Each inspection is costly. Hence, the agent stops her search at some point. The optimal strategy turns out to be simple. The agent assigns a reservation price to each option. These prices govern the order of inspection. The agent terminates her search whenever the maximum sampled reward is above the reservation price of every unsampled option. This basic model has been used in a wide array of applications, ranging from job search, see, e.g., Miller (1984) and references that followed, to real estate markets, see Quan and Quigley (1991). The fundamental model has been extended in many ways, see, for instance, Olszewski and Weber (2015) and references therein. There are two important differences between our setting and Weitzman (1979)'s canonical search model. First, we consider samples that are correlated over time. Second, we allow the agent to choose her search scope at any point.

The idea of modeling correlation over time using a Brownian path in a search setting is inspired by Callander (2011).[3] However, in Callander's setting, the features of the process are exogenous. Agents are short-lived, and need to decide whether to choose the optimal historical action or experiment with a new one, for which they get rewarded. Experimentation is then costly only in so much as it impacts rewards. Utilities are negatively proportional to the distance of the sample from 0. Callander and Matouschek (2019) offers a generalization allowing for some risk aversion. Garfagnini and Strulovici (2016) consider a related setting in which each risk-neutral short-lived agent effectively chooses a timed product, where product values follow a Brownian path. Later-timed products are then associated with greater variances. Thus, as in our setting, agents have some control over the variance, though it is restricted to a particular functional form. Wong (2020) studies the tension between exploration and exploitation when a firm searches for its ideal production scale, the returns to which follow a Brownian path. Firms get flow utilities from their samples throughout and pay a quadratic cost for their exploration speed. While closed-form solutions for the optimal policy are challenging to obtain, Wong illustrates that exploration hastens after poor outcomes and that search stops when observed outcomes are sufficiently poor.[4]

[3]See Jovanovic and Rob (1990) for axiomatic foundations of search discoveries following a Brownian path.

[4]The idea that variance might be a control variable associated with costs appears in other experimentation models. For instance, Moscarini and Smith (2001) consider a sequential sampling setting in which an agent can control the precision of the signals she receives.

As Callander (2011) aptly describes, Brownian motion represents a bandit model with a continuum of correlated, deterministic arms. The classic bandit problem, going back to Robbins (1952), assumes arms are independent. In particular, the dynamic experimentation applications assume intertemporal independence, see the survey by Bergemann and Valimaki (2008). In big part, this is due to the known difficulties correlations introduce. We believe allowing for correlation across discoveries is important for capturing the accumulation of expertise and knowledge that are at the heart of many applications.

The labor literature, going back to Pissarides (1984), has considered models of labor search in which firms or workers can invest in their *search intensity*, which affects their probability of finding potential matches; for a review, see Pissarides (2000). Our consideration of the scope of search highlights a different dimension of search efforts when observations are correlated.

Our analysis of contracts relates to the budding literature on contracts for experimentation, which has thus far focused on the (independent) one- or two-armed bandit setting, see Manso (2011), Halac, Kartik, and Liu (2016), and Guo (2016).

Our results utilize techniques from the mathematics literature on optimal stopping in which the objective is related to the maximum seen so far, see e.g. Peskir (1998) and Peskir (2005). The techniques we develop allow for the analysis of such optimal stopping problems with the inclusion of a control—in our case, the costly search scope. We hope these will open the door for further studies in the area.

3 A Model of Retrospective Search

Consider an agent who is searching in continuous time from a diffusion with no drift. At every time instance, the agent first decides whether to continue the search or not. If the agent decides to continue searching, she also decides on the instantaneous scope of search, which is costly. If the agent decides to stop search, she receives the best of her observations until, and including, the stopping time, net of her accumulated search costs.

Specifically, we assume the agent searches across a process dX_t and receives the value

M_t with

$$dX_t = \sigma_t dB_t$$
$$M_t = (\max_{0\le r\le t} X_r \vee M_0),$$

where B_t is the standard Brownian motion, σ_t is the controlled breadth of search, and M_0 and X_0 are exogenously given. For presentation simplicity, we assume $M_0 = X_0 = 0$. The correlation inherent in the process governing the evolution of X_t captures the idea that current search outcomes provide some indication for future outcomes: innovation begets innovation in research, items linked on an online shopping platform through reviews are similar in nature, etc. We assume no drift since, in many of these applications, the mere passage of time does not provide search improvements. Nonetheless, in Section 6.3 we discuss the impacts of drift.

At any point in time t, the searching agent controls the breadth, or scope, of her search. Specifically, the agent chooses a continuous measurable mapping $\sigma_t \in [\underline{\sigma}, \bar{\sigma}]$. We assume $\underline{\sigma} > 0$ so that instead of idling, the agent terminates search. The agent pays a cost $c(\sigma_t)$ for any instantaneous search scope σ_t. We assume that c is twice continuously differentiable, increasing, and log-convex. If $\underline{\sigma} = \bar{\sigma}$, our setting boils down to one in which the agent has fixed search scope she cannot control.

We assume the agent can recall past observations, so that when she stops her search, the maximal obtained value M_t is her reward. Consequently, the agent's problem can thus be written as follows:

$$\sup_{\tau,\{\sigma_t\}_{t=0}^{\tau}} E\left(M_\tau - \int_0^\tau c(\sigma_t)dt\right).$$

We discuss the impacts of risk attitudes when describing our main results, and specifically in Section 6.1.

The agent's ability to recall past realizations is relevant for our applications and important for our analysis. Absent recall, the agent would stop searching immediately, regardless of her current observation. Indeed, since we consider a driftless process, the expectation of any future value of the process coincides with the current observed value, but comes at a cost. In expectation, it is therefore not worthwhile.[5]

[5]Risk attitudes would only make search outcomes less appealing and would not alter this conclusion.

4 Optimal Retrospective Search

4.1 The Optimal Policy

In principle, the agent has two dimensions to consider at any point in time t: the maximum observed so far M_t, and the current outcome X_t. Her chosen search scope may therefore depend on both. The following proposition suggests that, in fact, the agent optimally conditions her choices only on the current outcome.

Proposition 1 (Reducing Dimensionality): *At any t, an infinitesimal change in the search scope, the control, σ_t has no effect on the current maximum M_t.*

To glean some intuition for the proposition, consider panel (a) of Figure 1, depicting a situation in which, at time τ, $X_\tau < M_\tau$. A small perturbation in the search scope at τ cannot impact the observed maximum M_τ. It follows that the chosen search scope should depend only on the local features of the process, namely X_τ. What happens at times τ in which $X_\tau = M_\tau$, as depicted in panel (b) of Figure 1? Our continuous-time formulation implies that, within any infinitesimal interval of time, with probability 1 the agent reaches a new value different than the current maximum.[6] The proposition then follows. Furthermore, the value of M_τ cannot affect the agent's choice of search scope.

The reduction in dimensionality suggested by Proposition 1 simplifies dramatically the analysis of optimal retrospective search, which we can now fully characterize. In general, optimal search is characterized by the search scope over time, as well as the *stopping boundary*: the function $g(M)$ that determines, for each observed maximal value, how low outcomes can go before the agent becomes sufficiently pessimistic to stop searching. That is, the agent continues searching as long as she observes outcomes above $g(M)$ and stops at τ^* given by:

$$\tau^* = \inf\{t \geq 0 : g(M_t) \geq X_t\}.$$

Proposition 2 fully characterizes the optimal search policy.

Adding drift to the governing process would still generate extreme behaviors. For drift lower than the search costs, there would be immediate stopping. For drift high enough, the agent would search indefinitely.

[6] In fact, since we assume linear utility, when $X_\tau = M_\tau$, the agent is facing an analogous problem to that faced by the agent at the outset of the process, at time 0. Whichever scope of search was chosen at time 0 is then optimal at time τ.

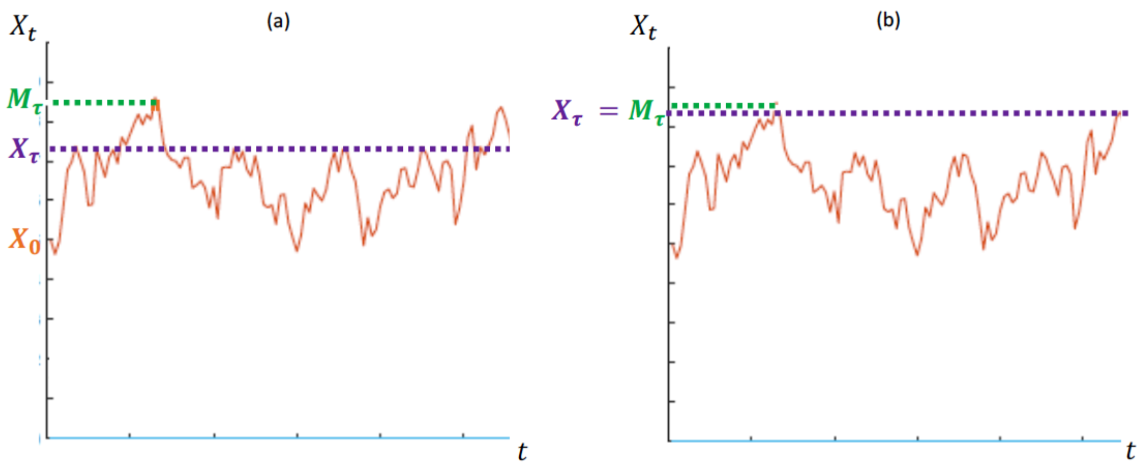

Figure 1: Independence of Local Choices on Established Maxima

Proposition 2 (Optimal Retrospective Search): *In the optimal policy, search scope σ^* is constant and, when interior, solves:*

$$\frac{2c(\sigma^*)}{c'(\sigma^*)} = \sigma^*.$$

Furthermore, the stopping bonudary at any point t with a previously observed maximum M_t is given by:

$$g(M_t) = M_t - \frac{(\sigma^*)^2}{2c(\sigma^*)}.$$

Proposition 2 has two implications. First, while in our setting search scope can be controlled, the agent *optimally* chooses it to be fixed at a level independent of both global features of the process—the maximum value observed—and the local features of the process—the samples recently observed. The fact that both costs and marginal costs impact the optimal solution is to be expected. Indeed, there are effectively two margins our agent considers, one corresponding to the instantaneous search scope, which needs to account for the marginal cost, the other corresponding to the length of search, which needs to account for the flow cost. In fact, the optimal search scope depends only on the underlying costs, with more log-convex costs generating lower search scope. Technically, this result suggests that, while the common assumption that search scope is fixed is without loss of generality, its precise level is naturally tied to the costs of search and their marginals.

In terms of applications, the result indicates that innovations do not require ongoing changes in development teams in terms of either size or work intensity. Interestingly, many companies do not fluctuate in their research and development expenditures over quarters.[7] One specific example for which detailed discovery data is more easily available is that of mineral prospecting. As described in Zhilkin (1961), geologists have substantial leeway in their search for a variety of minerals. Nonetheless, since the 50's, standard practice dictates sampling fixed amounts from a fixed depth along exploration paths, independent of discoveries already made on the path.[8]

The second message of the proposition is that the stopping threshold corresponds to the currently-held maximum of the process minus a *fixed* amount, which depends, again, on search costs. This fixed amount is often referred to as the *drawdown* of the process. Going back to our example of product development, this result suggests that product release follows innovations that do not appear sufficiently promising, where promise is assessed in relation to the best option observed. The distance between the two that leads to the termination of search does not change with time or with the realized discovery path.

The intuition underlying the proposition is as follows. The crucial observation is that the optimal policy is indeed governed by a stopping boundary $g(M)$. The agent stops searching whenever her current observation is sufficiently low relative to the maximal search outcome she has observed and can go back to. Naturally, $g(M)$ depends on the features of the process: both exogenous and endogenous. That is, holding all of the parameters in our environment fixed, $g(M)$ depends only on the (endogenous) choice of the search scope, the agent's control. That dependence is non-trivial: the optimal stopping boundary would change were search scope very low or very high. Indeed, for very low scopes of search the agent would want to stop quickly as there is not much to gain from costly search. For very high search scope, continuation is prohibitively costly and the agent would terminate search rapidly as well. For intermediate levels of search scopes, the agent may benefit from a non-trivial search.

From Proposition 1, the search scope depends only on local conditions and not on the

[7] This observation holds for many companies, including Twitter, Facebook, Pinterest, Honda, and Toyota; see `https://ycharts.com` and `https://www.statista.com/`. Many companies exhibit a trend in their expenditures, in line with inflation, expansion, etc., but occasional spikes in expenditures appear rare.

[8] The search for minerals most commonly occurs over a one-dimensional path, often termed a *vein*, for details see U.S. Department of Agriculture (1995). As mentioned, correlations between adjacent plots make this example particularly germane for our setting.

maximal value observed thus far. Changing the search scope from 1 to σ at any small interval of time is tantamount to "speeding up" the process by a factor of σ^2. On a given path, the agent then effectively finds the efficient *speed* to hit the boundary $g(M_t)$. This amounts to minimizing the cost per speed, or $c(\sigma)/\sigma^2$. The corresponding first-order condition yields the expression for σ^*.

The intuition for the drawdown nature of the optimal stopping policy is more straightforward. Indeed, consider a process $Y_t = X_t + a$, where a is a constant. Since utility is linear, marginal considerations remain the same for this process and the optimal solution should echo the one we analyze. In particular, the optimal search scope is constant. Importantly, the optimal stopping boundary $\tilde{g}(M)$ must satisfy $\tilde{g}(M) = g(M-a)$. In other words, the optimal stopping boundary depends only on the distance from the observed maximum.

Since Proposition 2 establishes that the optimal scope of search is constant, we can use the Reflection Principle (see, e.g., Rogers and Williams (2000)) to infer that M_t follows the same distribution as $|\sigma^* B_t|$. It then follows that, for any t,

$$\mathbf{E}(M_t) = \sigma^* \sqrt{2t/\pi}.$$

That is, the scope of search directly affects the expected value of the observed maximum.[9]

As mentioned in our literature review, models of labor search often consider firms' or workers' choice of search intensity, which influence the likelihood of encountering potential matches and consequent expected payoffs, see Pissarides (1984, 2000). Those models usually assume encounters in the market are independent. In our model, with correlated samples, the scope of search captures a rather different aspect of search efforts—the breadth, or speed of search. Nonetheless, the impact on the expected returns to search are similar, with higher search scope yielding a higher expected maximal value.

[9] Formally, let T_a denote the first time X_t hits some level a. Then,

$$\Pr(M_t \geq a) = \Pr(T_a \leq t) = 2\Pr(X_t \geq a) = \Pr(|X_t| \geq a),$$

where the second equality follows since, if the process hits the level a at time T_a, it has equal probability of moving above or below a. Since B_t is normally distributed with mean 0 and variance t, we have that $\mathbf{E}(M_t) = \mathbf{E}(\sigma^* |B_t|) = \sigma^* \sqrt{2t/\pi}$.

While the intuition, and proof, that search scope is constant does not depend on the linearity of the agent's utility, the argument for a fixed drawdown stopping boundary most certainly does. In Section 6.1 and in the Online Appendix, we discuss an extension to general concave utilities. The analysis becomes far more intricate technically. However, for constant relative risk aversion (CRRA) utilities with parameter ρ, that is, assuming the utility from a maximal value of M is captured by $u(M) = \frac{M^{1-\rho}}{1-\rho}$, a closed-form solution for the stopping threshold can be derived. Intuitively, as the agent becomes more risk averse, increasing ρ, the marginal value of improving the already attained maximum declines, and the agent demands superior outcomes to continue searching. Furthermore, the stopping boundary is no longer a fixed-drawdown boundary. In particular, as the attained maximum increases, the marginal value of an improvement decreases, and the agent is less likely to continue searching.

Proposition 2 also suggests the optimal stopping boundary when search scope is exogenously fixed at a constant $\hat{\sigma}$ with associated search cost of $\hat{c}$. The optimal boundary at any time t with observed maximum value M_t is then

$$g(M_t) = M_t - \frac{\hat{\sigma}^2}{2\hat{c}}.$$

Indeed, we can always find a log-convex cost function c such that $\hat{c} = c(\hat{\sigma})$ and $\hat{\sigma}$ is the optimal (endogenous) search scope for search costs given by $c(\cdot)$. The result then follows directly from the second part of Proposition 2.[10]

4.2 Retrospective Search Outcomes

We now turn to the value of search. Let $d = \frac{(\sigma^*)^2}{2c(\sigma^*)}$ denote the drawdown size governing the optimal stopping policy. Since the process is symmetric, and can go up or down with equal probability, this would suggest the agent might hope to gain up to d while searching. How long it takes to achieve the maximum should naturally depend on the search scope: the greater the search scope, the lower the expected time. The following proposition provides the precise expectation of accomplished project values and expected search times.

[10]This special case is reminiscent of that studied in Peskir (1998). For the uncontrolled problem, he offers an ordinary differential equation governing the optimal stopping boundary. It is difficult to use his formulation for general utilities. Our analysis in the Appendix and Online Appendix offer an alternative approach.

Proposition 3 (Expected Values of Retrospective Search): *The expected project value upon optimal stopping is* $\mathbf{E}(M^*_\tau) = \frac{\sigma^{*2}}{2c(\sigma^*)}$ *and the expected optimal search duration is* $\mathbf{E}(\tau^*) = \frac{\sigma^{*2}}{4c^2(\sigma^*)}$*. Consequently, the expected payoff from optimal retrospective search is:*

$$\mathbf{E}(\Pi^*) = \mathbf{E}(M^*_\tau) - c(\sigma^*)\mathbf{E}(\tau^*) = \frac{\sigma^{*2}}{4c(\sigma^*)}.$$

To gain some more technical intuition, consider the moving maximum of our process and let $d = \frac{(\sigma^*)^2}{2c(\sigma^*)}$ denote the drawdown governing the optimal stopping policy. Consider the difference between the record-high level at time t and the observed value at time t, $M_t - X_t$. From continuity of the Brownian motion, the optimal stopping policy implies that, if the agent stops at time τ^*, $M_{\tau^*} - X_{\tau^*} = d$. It follows that, for any realized stopping time τ^*,

$$\mathbf{E}(M_{\tau^*} - X_{\tau^*} \mid \tau^*) = d.$$

Now, X_t is a martingale with expectation 0. Therefore, for any stopping time τ, $\mathbf{E}(X_\tau \mid \tau) = 0$. It then follows that, for any t,

$$\mathbf{E}(M^*_\tau) = \mathbf{E}(M^*_\tau | \tau^* = t) = d.$$

The proposition's claim then follows. Furthermore, how long it takes for retrospective search to run its course is not indicative of the resulting expected value of the project.

Consider the expected search duration. As d goes up, the agent is more willing to continue search, and search duration increases. A scaling of Brownian motion by σ is tantamount to a "speeding up" of time by a factor of σ^2. Search scope therefore proxies for the "speed" of search and search duration is inversely proportional to it. In fact, expected search time is precisely the ratio of the squared drawdown size and the speed of the process σ^2; that is, $\mathbf{E}(\tau^*) = \frac{d^2}{\sigma^2}$.

This result also highlights a subtle connection between the efficient search scope and the optimal drawdown. Given any drawdown d, the efficient scope is constant and minimizes the cost per speed. Conversely, for any fixed search scope σ, reorganizing the expectation in propositon 3 yields $\mathbf{E}(\Pi(d,\sigma)) = d - d^2\frac{c(\sigma)}{\sigma^2}$. The optimal drawdown size then depends on the search scope and, in turn, is maximized at the optimal search scope. In this formulation, the optimization problem is reminiscent of a monopolist choosing a "quantity" d, where the price is fixed at 1 and production costs are quadratic and given by $d^2\frac{c(\sigma)}{\sigma^2}$.

Viewed through this lens, investment in search scope is analogous to investment in a reduction of production costs, which impacts quantity.[11]

Going back to the case in which search scope is exogenously fixed at a constant $\hat{\sigma}$ with associated search cost of $\hat{c}$, natural comparative statics emerge. Keeping costs fixed, as search scope increases, expected payoffs go up. Keeping the search scope fixed, as search costs go up, expected payoffs go down. In what follows, we discuss comparative statics for general search costs.

4.3 Comparative Statics

As noted after Proposition 2, the optimal search scope declines as costs become more log-convex. For any fixed stopping boundary, the time to stop would then increase as search scope declines. However, costs naturally impact the stopping boundary as well. Let $c_1(\cdot)$ and $c_2(\cdot)$ be two cost functions such that $c_2(\cdot)$ is more log-convex than $c_1(\cdot)$ and denote by σ_i^* and $\mathbf{E}(M_i^*)$ the optimal search scope and expected project value under cost function $c_i(\cdot)$. We know that $\sigma_1^* > \sigma_2^*$. The ultimate impact on $\mathbf{E}(M_i^*)$ or, equivalently, on the expected search payoff, depends on how the drawdown size changes. If $c(\sigma_2^*) > c(\sigma_1^*)$, then search is truly inhibited: it is optimally less ambitious and more costly per unit of time. In this case, $\mathbf{E}(M_1^*) > \mathbf{E}(M_2^*)$. However, if $c(\sigma_2^*) < c(\sigma_1^*)$, the comparison is inconclusive in general.

To see the nuanced impacts of cost changes, consider the particular class of log-convex cost functions, $c(\sigma) = \exp(\sigma^\gamma)$, with $\gamma > 1$. As already mentioned, the optimal search scope declines as costs become more log-convex. Here, the optimal search scope $\sigma^* = \left(\frac{2}{\gamma}\right)^{1/\gamma}$ indeed declines in γ and shown on panel (a) of Figure 2. Notice that $\sigma^* > 1$ whenever $\gamma < 2$ and $\sigma^* < 1$ whenever $\gamma > 2$. Since project values are governed by a Brownian motion, for any stopping boundary characterized by a fixed drawdown, the resulting maximal value is proportional to σ^*, while the expected search time is inversely proportional to the "speed," captured by σ^{*2}. Thus, for small γ, small increases in γ have a greater impact on the expected search time, and therefore expected search costs, than on the expected maximal

[11] With a unit price of 1, It is well-known from the monopolist's quantity setting problem with quadratic costs that the optimal profits correspond to precisely half the optimal quantity, which is reflected in the proposition. Indeed, $\mathbf{E}(\Pi^*) = \frac{1}{2}\mathbf{E}(M_\tau^*)$.

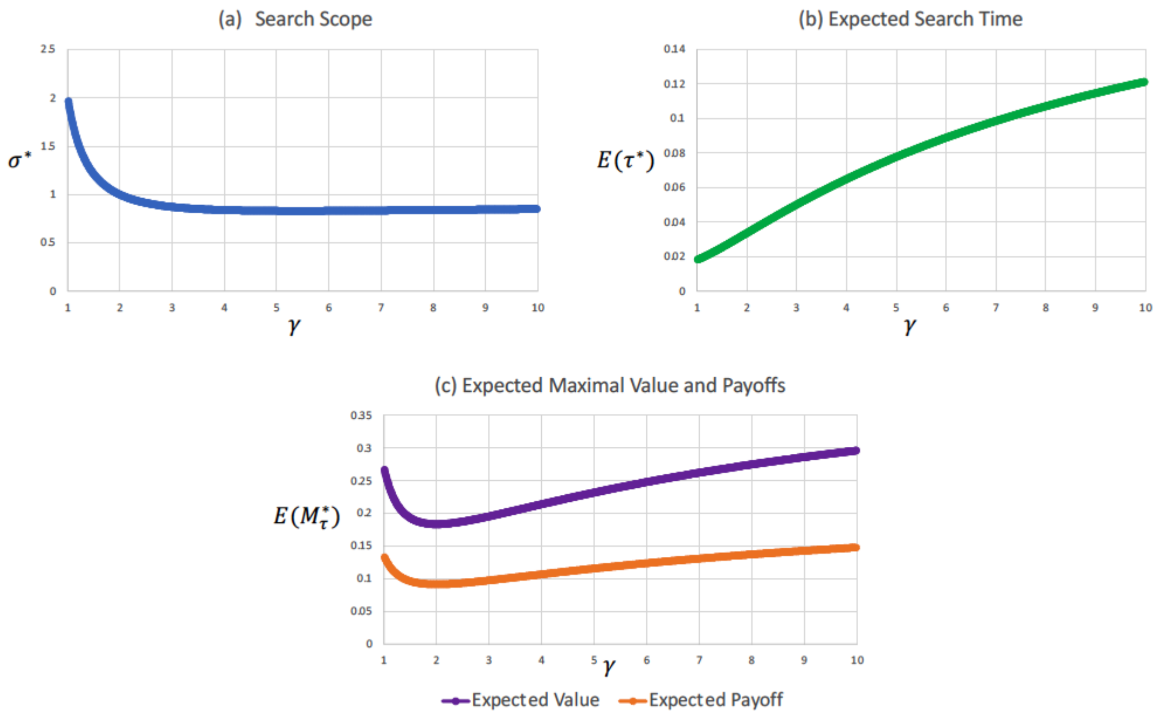

FIGURE 2: Impacts of Cost Changes when $c(\sigma) = \exp(\sigma^\gamma)$

value of the project. In particular, for small γ, the expected payoff from retrospective search declines. The reverse occurs for larger γ.

Formally, from Proposition 3, the expected search time is given by $\mathbf{E}(\tau^*) = \frac{1}{4}\left(\frac{2}{e^2\gamma}\right)^{2/\gamma}$, which is increasing in γ and depicted in panel (b) of Figure 2. The expected project value is given by $\mathbf{E}(M_\tau^*) = \frac{1}{2}\left(\frac{2}{e\gamma}\right)^{2/\gamma}$, which is double the expected payoff from search, and is non-monotonic in γ. It is decreasing initially and then increasing indefinitely, as depicted in panel (c) of Figure 2.

In other words, increasing log-convexity makes the agent less ambitious in terms of search scope, but also prolongs the search. Payoffs are higher with greater log-convexity. However, importantly, they are non-monotonic. For low levels of γ, the benefits of ambition overwhelm those of the length of search.

In general, a point-wise increase in the cost function or its marginals can lead to either an increase or decrease in search scope. For example, suppose $\bar{c}(\cdot)$ is defined by $\bar{c}(\sigma) = c(\sigma) - w$ for all σ, where $w \geq 0$. This would correspond to a case in which, say, a constant flow wage of w is paid to the agent as long as she searches. The optimal search scope corresponding to $\bar{c}(\cdot)$ would then satisfy:

$$\sigma^*(w) = \frac{2\left(c(\sigma^*(w)) - w\right)}{c'(\sigma^*(w))}.$$

If $\sigma c'(\sigma)$ decreases in σ, then $\sigma^*(w)$ increases in w. In other words, a shift downward in the costs increases search scope. If $\sigma c'(\sigma)$ is non-decreasing, the impact of wages on search scope need not be monotonic and depend on the curvature of the cost function.

One useful family of costs that we will return to when we discuss principal-agent interactions is that of linear costs. Namely, assume that $c_{a,b}(\sigma) = -a + b\sigma$, where $a, b \geq 0$. In this case, $\sigma c'_{a.b}(\sigma)$ increases in σ, but this family of costs is special in that we can sign the impacts of shifts in both the level of cost parameter a and the marginal cost parameter b. In particular, the optimal search scope is:

$$\sigma^*_{a,b} = \frac{2a}{b},$$

which is monotonically increasing in the level of cost parameter a and monotonically decreasing in the marginal cost parameter b.

5 Commissioned Retrospective Search

We now incorporate retrospective search into a moral hazard problem, considering commissioned search. We assume that a principal (she) contracts with an agent (he) who has access to our retrospective search technology. The principal is the residual claimant of search outcomes, but cannot conduct the search herself. This is often the case with research and development teams that are separate from the main shareholders of a company. The principal then cares about the outcome of the search, but does not experience its direct costs. Similarly, artists often commission the sale of their pieces to galleries, which can access a pool of potential buyers they can search through; and home-owners frequently use the help of real-estate agents, who search for a buyer on their behalf.

In such settings, the principal does not know what effort the agent exerts in his search. That is, the principal does not observe the search scope σ. We now think of $\sigma_t = \underline{\sigma}$, with the agent exerting the minimal scope, as the agent shirking. We do, however, assume that the principal sees whether the agent is on the job or not.

We consider contracts that are comprised of a fixed wage $w \geq 0$ and a fraction $\alpha \in (0,1]$ of the final search outcome. We call the combination of wage and fraction a *commission contract* and denote it by (w,α). Commission contracts correspond to sharing rules first considered by Aghion and Tirole (1994) and are commonplace in the field. For instance, Jensen and Thursby (2001) report the licensing practices of 62 U.S. universities. Their data suggests the prevalence of commission contracts, namely agreements based on fixed fees and royalties. In the realm of mineral exploration, the Securities and Exchange Commission (SEC) reports on thousands of joint venture agreements between investors and mining companies each year. These contracts often specify fixed and flow fees throughout the search process, in addition to pre-agreed upon shares of the findings.[12]

If the agent does not work for the principal, he can take an outside option that offers him $\underline{u}$. For expositional simplicity, we assume this outside offer vanishes once the agent contracts with the principal. While our analysis does not hinge on this assumption, we believe it is realistic for many applications—e.g., employees who turn down offers cannot reconsider them soon thereafter. As will hopefully become clear from our analysis, however, if the agent prefers pursuing search over his outside option at the outset, he will maintain this preference throughout the search process.

For simplicity, we assume $X_0 = M_0 = 0$. The principal's problem is then:

$$max_{w,\alpha}\mathbf{E}((1-\alpha)M_{\tau_{w,\alpha}} - \tau_{w,\alpha}w)$$
$$\text{s.t.}\tau_{w,\alpha} \in \arg\max_{\tau,\{\sigma_t\}_{t=0}^{\tau}} \mathbf{E}(\alpha M_\tau - \int_0^\tau [c(\sigma_t) - w]dt),$$

where $dX_t = \sigma_t dB_t$ and $M_t = \max_{0\leq s\leq t}(X_s \vee M_0)$ as before.

5.1 The Agent's Problem

We start by analyzing the agent's optimal choices for any commission contract (w,α). In any optimal contract, we must have $w < c(\underline{\sigma})$; otherwise, the agent would shirk indefinitely. We therefore maintain that as an assumption.

[12]See https://www.sec.gov/edgar/search/ for a full set of reported contracts since 2001.

Our benchmark case of an agent searching on his own, which we analyzed in Section 4.1, can be seen as a special case of an agent responding to a commission contract with wage $w = 0$ and full remuneration for efforts in the form of an $\alpha = 1$ share of the ultimate maximal value found. Using similar techniques, we can find the characterization generalizing Proposition 2:

Proposition 2* (Optimal Comissioned Search): *With commission contract (w, α), the agent's optimal search scope is constant and solves:*

$$\sigma^* = \frac{2(c(\sigma^*) - w)}{c'(\sigma^*)}$$

if such a $\sigma^* \geq \underline{\sigma}$ exists, and otherwise satisfies $\sigma^* = \underline{\sigma}$. Furthermore, the stopping boundary under contract (w, α) at any point t with a previously observed maximum M_t is given by:

$$g(M_t) = M_t - \frac{\alpha(\sigma^*)^2}{2(c(\sigma^*) - w)}$$

Intuitively, a wage effectively shifts downward the search costs by a constant amount. That is, the agent effectively considers a cost $\tilde{c}(\cdot) = c(\cdot) - w$. The formula for the optimal search scope then follows directly from that identified in Proposition 2. The stopping boundary, however, needs to be adjusted according to the commission rate α. As α decreases, the agent would be more keen to stop. In fact, as α becomes vanishingly small, the agent stops searching immediately.

One important feature of the optimal policy is that the search scope is fixed and independent of the commission rate α offered. It solely depends on the wage w and the agent's private cost, which together define the flow search expenses. To some extent, this is to be expected since, as we explain following Proposition 2, the optimal search scope is independent of the maximal value obtained at any moment. Therefore, it should not be sensitive to the commission awarded. Getting a smaller share of the pie still limits search, however. The impact manifests only through the stopping boundary—if the agent gets a small share, he is likely to stop searching sooner, but does not alter his search scope.

The impacts of wages are intimately connected to the curvature of the cost function, as suggested by our discussion in Section 4.3. If $c'(\sigma)\sigma$ is decreasing in σ, then increasing wages increases search scope and the drawdown size. That is, greater wages induce the agent to search more ambitiously and more extensively.[13] If $c'(\sigma)\sigma$ is non-decreasing in σ, the impacts of wages depend on the precise shape of the cost function.

In principle, the agent may choose a corner solution in terms of his search scope. However, regardless of search costs, the principal can always set the wages sufficiently high so that the agent is induced to search at a greater, interior, scope. We note, however, that, in general, wages may be low, even negative. Indeed, we do not impose any limited-liability constraints. The inclusion of a lower boundary on admissible wages would not alter the methods we present and, if anything, would make interior search intensities easier to sustain optimally as wages would naturally be forced to be higher. We maintain no such constraints both for presentation simplicity and for realism. Indeed, it is not uncommon for commissioned researchers to rent labs or testing equipment, cover various experimental outlays, etc.[14] Such flow expenses would formally translate into negative wages.

We now turn to the returns of commissioned search resulting from the agent's optimal policy. For each commission contract (w,α), we can identify the expected payoff to both the principal, denoted by $\mathbf{E}(\Pi^P_{w,\alpha})$, and the agent, denoted by $\mathbf{E}(\Pi^A_{w,\alpha})$. The precise formulation of these expected payoffs naturally generalizes our characterization in Proposition 3. Naturally, the principal and agent get complementary shares of the pie. Furthermore, their flow costs differ—the principal experiences a flow cost of w, while the agent experiences a flow cost of $c(\sigma^*)-w$.

Proposition 3* (Outcomes of Commissioned Search): *The expected project value under a commission contract* (w,α) *is* $\mathbf{E}(M_{\tau_{w,\alpha}}) = \frac{\alpha(\sigma^*)^2}{2(c(\sigma^*)-w)}$ *and the expected search duration is* $\mathbf{E}(\tau_{w,\alpha}) = \left(\frac{\alpha(\sigma^*)}{2(c(\sigma^*)-w)}\right)^2$. *Let*

$$G(\beta,\phi) = \frac{\beta\alpha(\sigma^*)^2}{2(c(\sigma^*)-w)} - \phi\left(\frac{\alpha(\sigma^*)}{2(c(\sigma^*)-w)}\right)^2.$$

The expected returns the agent and principal get from a commission contract (w,α) *are:*

$$\mathbf{E}(\Pi^A_{w,\alpha}) = G(\alpha, c(\sigma^*)-w) \quad \text{and} \quad \mathbf{E}(\Pi^P_{w,\alpha}) = G(1-\alpha, w).$$

[13] Indeed, when considering the drawdown size, if $c'(\sigma)\sigma$ decreases in σ, so does $c'(\sigma)/\sigma$ and our conclusion follows directly from substituting the optimal search intensity in the expression of the stopping boundary.

[14] Labs, both in the hard sciences, such as the Marine Biological Lab at the University of Chicago, as well as in the social sciences, such as the Oxford Experimental Lab, offer access to research resources at a fee. Platforms such as scienceexchange.com offer online marketplaces for various aspects of research.

5.2 The Principal's Problem

The characterization of the agent's problem and the resulting expected payoffs to the principal enter the principal's optimization problem. Using the characterization of Propositions 2* and 3*, we can write the principal's problem as follows:

$$max_{w,\alpha}\mathbf{E}(\Pi^{P}_{w,\alpha}) \text{ subject to } \sigma^* = \frac{2(c(\sigma^*)-w)}{c'(\sigma^*)}.$$

Notice that σ^* is pinned down uniquely by a choice of w. In particular, if the principal wants to induce a search scope of $\sigma \geq \underline{\sigma}$, the wage she needs to offer is:

$$w(\sigma) = c(\sigma) - \frac{\sigma c'(\sigma)}{2}.$$

We can then convert the principal's problem into an unconstrained problem in which she selects the commission share α and the constant search scope σ:

$$max_{\sigma,\alpha}\frac{(1-\alpha)\alpha\sigma}{c'(\sigma)} - \left(c(\sigma) - \frac{\sigma c'(\sigma)}{2}\right)\left(\frac{\alpha}{c'(\sigma)}\right)^2.$$

If the principal engages with the agent at all, it must be the case that $\alpha \in (0,1)$ and the optimal share chosen should satisfy a first-order condition. The search scope the principal optimally targets should satisfy a first-order condition whenever interior. As a consequence, we have the following characterization:[15]

Proposition 4 (Optimal Commission Contract): *Whenever the principal's optimal commission contract* (w^*, α^*) *guarantees an interior search scope, it satisfies:*

$$w^* = c(\sigma^*) - \frac{\sigma^* c'(\sigma^*)}{2} \quad \textit{and} \quad \alpha^* = \frac{\sigma c'(\sigma^*)}{2c(\sigma^*) + \sigma c'(\sigma^*)},$$

where

$$\sigma^* = \frac{4c(\sigma^*)}{c'(\sigma^*) + \sigma^* c''(\sigma^*)}.$$

As discussed in Section 4.3 and the previous subsection, general comparative statics depend on details of the cost function and are challenging to characterize in general. In what follows, we consider the special case of linear costs in order to illustrate simply how the optimal commission contract responds to the environment's features.

[15] The proof involves simple algebraic manipulations of the corresponding first-order conditions and is therefore ommitted.

5.3 Contracting with Linear Search Costs

Suppose search costs are linear, $c(\sigma) = -a + b\sigma$, where $a, b > 0$ and $c(\underline{\sigma}) = -a + b\underline{\sigma} > 0$. That is, costs are positive and increasing over the relevant range of search intensities. Were the agent searching on his own, the optimal constant search scope would be given by $\sigma^{NC} = \frac{2a}{b}$, see our discussion in Section 4.3.

Using Proposition 4 above, the optimal commission contract (w^*, α^*) satisfies

$$w^* = -\frac{a}{3} \quad \text{and} \quad \alpha^* = \frac{2}{3}$$

with induced search scope of

$$\sigma^* = \frac{4a}{3b}.$$

The agent searches with lower scope when under a commission contract relative to when searching on his own. The linearity of costs simplifies dramatically the structure of the optimal contract. In particular, the optimal commission is independent of the precise parameters of the costs, both their level captured by a and their margins captured by b. The optimal wages are negative and increasing in cost levels, but independent of their marginals. Intuitively, wages act as cost shifters and are therefore closely linked to the level of costs captured by a. As search costs shift up, wages are set so that the agent pays smaller fees to participate in the search.

Proposition 3* suggests that the expected maximal value of the project is $E[M_\tau^*] = \frac{8}{9}\frac{a}{b^2}$, while the expected search duration is $E[\tau^*] = \frac{4}{9b^2}$. Thus the overall expected maximum is decreasing in both the level of costs and their marginals. Search duration, however, depends only on the marginal cost of search, with greater marginal costs yielding shorter expected searches. The resulting expected payoffs to the agent and the principal then coincide and can be calculated to equal:

$$\mathbf{E}(\Pi^A_{w^*,\alpha^*}) = \mathbf{E}(\Pi^P_{w^*,\alpha^*}) = \frac{4}{9}\frac{a}{b^2}.$$

In other words, with linear costs, the agent and the principal split the overall search surplus evenly.

If the agent and principal contemplate different projects, characterized by different linear search costs, their assessments of the projects would be fully aligned and ultimately boil down to the comparison of the fixed costs to squared marginal costs.

With a single-agent's retrospective search, Proposition 3 suggests an expected payoff of $\mathbf{E}(\Pi^*) = \frac{a}{b^2}$. In contrast, the overall surplus generated by the agent and the principal is $\frac{8}{9}\frac{1}{b^2}$. Thus, With linear costs, contractual frictions come at a cost of *one ninth*, or about 11%, of the surplus.

Linear costs are certainly special and, as can directly be seen from Proposition 4, second-order terms of the cost function can generally play a role in the contractual solution. In the Online Appendix, we also analyze the case of quadratic costs. While the analysis becomes a bit more intricate, the basic messages remain qualitatively similar.

6 Conclusions and Discussion

This paper proposes a simple model of retrospective search. Agents—product developers, politicians, geological survey teams, or online shoppers—observe evolving options that are correlated over time. They have two decisions at each point: their scope of search and whether or not to stop and collect the maximal observed value. In our model, the scope of search is closely tied to the speed with which search is conducted and directly affects the expected maximal value that is generated through search. The optimal search policy entails a constant scope of search and a simple stopping boundary: a retrospective searcher ceases search whenever the observed value is a certain fixed and constant distance below the maximal value observed. Our characterization of the optimal policy offers an array of comparative statics and is amenable to embedding in a variety of applications. Specifically, we illustrate a principal-agent application in which the principal—an innovator, an artist, a home seller, etc.—cannot perform the search herself, but can contract with an agent—a university, an art gallery, a real-estate agent, and the like—to conduct the search. We fully characterize the optimal commission contract, comprised of a search wage and a commission, a pre-specified fraction of the final project returns. The resulting scope of search depends only on the wage, while the induced stopping boundary depends on both the wage and the commission. When scope of search entails linear costs, we show that contractual frictions come at a cost of one ninth of the surplus.

We hope our framework is useful for many search processes that exhibit intertemporal correlations, from labor search, to policy experimentation, to online commerce. In what follows, we discuss several natural extensions of our benchmark model, inspecting the role of risk attitudes, discounting, drift, and the possibility of resetting.

6.1 Risk Aversion

As mentioned in Section 4.1, when considering optimal retrospective search, the intuition, and proof, that search scope is fixed does not depend on the linearity of the agent's utility. The characterization of the optimal scope remains virtually identical to that pertaining to the risk-neutral case. Nonetheless, the characterization of the optimal stopping boundary changes and, in general, need not be characterized by a drawdown.

Absent a particular functional form for the utility, it is difficult to analytically characterize the optimal stopping boundary. Why is that? The characterization of the optimal stopping boundary emerges from the following. When the agent hits the boundary and stops search, her continuation value coincides with the precise value of the stopping boundary. In fact, standard results imply what is often termed a smooth-pasting condition, whereby the stopping boundary and the continuation value coincide smoothly, with all their derivatives agreeing. That restriction generates an ordinary differential equation (ODE) of the following form. For any well-behaved utility function u, and achieved maximum M, the optimal stopping boundary $g(\cdot)$ satisfies:

$$g'(M) = \frac{(u'(M))(\sigma^*)^2}{2c(\sigma^*)(M - g(M))}.$$

In fact, analysis following Peskir (1998) illustrates that the optimal stopping boundary is the maximal solution $g(M) \leq M$ satisfying this ODE. This is a non-linear and non-homogeneous ODE. When $u'(M)$ is a constant, as in most of the paper, it is easily solvable. In general, we show in the Online Appendix that it ties to a well-known class of ODEs, referred to as Abel's equation of the second kind (see, e.g., Murphy (2011)). The general solution for this class of ODEs has been an open question for nearly 200 years. Nonetheless, for particular functional forms, such as those corresponding to constant relative risk aversion (CRRA) utilities, we can derive analytical solutions using recent developments in mathematics, particularly the parametric solution in Panayotounakos and Kravvaritis

(2006). In the Online Appendix, we offer some general guidance on the techniques required to solve for optimal retrospective search with non-linear utilities.

For illustration, consider the class of CRRA utilities with parameter ρ, where utility from a maximal value of M is captured by $u(M) = \frac{M^{1-\rho}}{1-\rho}$ and assume $M_0 \geq 1$ so that agents are indeed risk averse. There is always a level $\bar{M}$ such that whenever the maximal observed value M exceeds $\bar{M}$, the agent stops her search immediately. Intuitively, at $\bar{M}$ the marginal returns from increasing the reward are overwhelmed by the marginal costs that search entails. Furthermore, as the degree of risk aversion ρ increases, the corresponding level $\bar{M}$ decreases—as the agent becomes more risk averse, increasing ρ, the marginal value of improving the already attained maximum, declines and the agent becomes more demanding when deciding whether to continue search. The stopping boundary for levels $M \leq \bar{M}$ can also be derived. Naturally, it is no longer characterized by a fixed drawdown. As the attained maximum increases, the marginal value of an improvement decreases, and the agent is less likely to continue searching.

6.2 Discounting

We assumed search costs are captured through the function $c(\cdot)$. For a constant search scope, search costs translate into constant flow costs. We believe this captures first-order frictions in various search applications. Nonetheless, one could also contemplate other frictions. For instance, suppose that the agent experiences exponential discounting costs in lieu of flow search costs and, for simplicity, suppose that the search scope is fixed. Formally, assume the agent's objective takes the form of $e^{-rt}M_t$, where $r > 0$ is the discount rate, with a given search scope σ that governs the underlying path. Since $\ln(\cdot)$ is monotonically increasing, we can equivalently consider the agent's objective as $\ln(M_t) - rt$. This objective coincides with the one we study when search scope is constant (with a cost of r per unit of time). However, rewards are now concave in the achieved maximal value. The resulting optimization problem is then a special case of the analysis of our original problem allowing for risk aversion, as discussed above. As it turns out, when utility is logarithmic, a closed-form solution for the optimal stopping boundary can be derived. We provide some details in the Online Appendix.

We could also allow for discounting of utilities on top of the costs entailed by the scope of search. However, the full characterization of the optimal policy, in particular the stopping boundary, is challenging to derive in closed-form. The Online Appendix provides one approach for tackling such settings as well.

6.3 Allowing for Drift

Throughout the paper, we assume the path of discovery has no drift. We do so for two reasons. First, it seems reasonable for the search applications we have in mind: when looking for new ideas, products or services to acquire, etc., the mere passage of time need not improve values absent any effort. Second, the no-drift assumption simplifies our presentation. Having said that, with a fixed drift and search scope, our analysis remains virtually identical. The stopping boundary does need to be adjusted by a constant, however, with greater drifts associated with more lenient stopping boundaries. Indeed, when the drift is high, the agent has a stronger incentive to continue searching.

Naturally, one could also contemplate drift as the instrument the agent controls, instead of the search scope, as considered by Peskir (2005). That turns out to generate straightforward results with a bang-bang nature. Namely, the agent always prefers higher drift, and the problem boils down to a simple calculus comparing costs with benefits. For small enough costs, a searching agent always chooses the maximal possible drift. For high enough costs, the agent dispenses with drift.

6.4 Resetting the Process

One might wonder about the possibility of resetting the process upon observing bad outcomes. Certainly, if such resetting were possible, an agent performing retrospective search would want to go back to the highest observed value whenever observing a lower project value. Given the continuous-time nature of our environment, the problem would then be ill-defined. Another way to view our setting is as follows. Suppose that the path describing the evolution of values is realized at the outset, as in Callander (2011). The scope of search σ is then simply a proxy of the speed σ^2 at which the agent traverses the process. The characterization we offer would then still hold. Since the path is realized once and for all, however, “resetting” the process by returning to a point at which observed values were high would be of no value. Our example of mine prospecting is instructive. If an exploratory well is not deemed good enough for further development, there is no point in returning to an already-explored well and inspecting it, or those that followed it on the path, yet again. As an analogy, consider driving through a new city seeking a restaurant. You can certainly vary your driving speed to cover more ground, and you might decide to go back to a previously-seen restaurant and stop your search. However, there is no point in going back to a restaurant and driving through the same streets again. The path of values—here, the sequence of restaurants—does not change with every drive through.

7 Appendix

7.1 Background to General Stopping and Control Problems

In this section, we provide a heuristic derivation of the Hamilton-Jacobi-Bellman (HJB) equation for the general stopping and control problem.

We consider a slightly more general underlying Weiner process X_t than that discussed throughout the paper and allow it to have a potentially non-trivial drift. The process X_t is therefore defined by its starting point X_0, its drift μ, and its standard deviation σ, which is controlled by the agent. For simplicity, we assume, as in the paper, that $X_0 = 0$.

Let $Z_t = (M_t, X_t)$ and let $V(Z_t)$ denote the continuation value for a slightly more general problem, where the utility function u is not necessarily linear, but is a uniformly Lipschitz continuous function that is twice differentiable with $u(0) = 0$.

$$V(Z_t) = \max_{\tau, \{\sigma_t\}_{t=0}^{\tau}} \mathbf{E}(u(M_\tau) - \int_0^\tau c(\sigma) dt | Z_t = Z).$$

Since the Brownian motion has independent increments, excluding the point in time t of consideration from the state description is without loss of generality. In particular, it suffices to consider the optimization at $t = 0$.

At any instance, the agent has two options: stop or continue. If the agent stops, she receives $u(M_t)$; if she continues, she receives $V(Z_t) = V(M_t, X_t)$. Thus, it is optimal to stop whenever $u(M) \geq V(M, X)$. If the agent does not stop, she chooses a search scope σ at a cost $c(\sigma)$. Now, for a heuristic derivation, assume that the agent chooses either an optimal fixed σ for a small amount of time dt, or stops immediately. Then, the dynamic programming principle will yield:

$$V(M_t, X_t) = \max \left\{ u(M_t), \max_\sigma \{-c(\sigma)dt + \mathbf{E}(V(M_{t+dt}, X_{t+dt} | \sigma, M_t, X_t))\} \right\}.$$

Equivalently,

$$V(M_t, X_t) = \max \left\{ u(M_t), \max_\sigma \{-c(\sigma)dt + \mathbf{E}(V(M_t, X_t) + d(V | \sigma, M_t, X_t))\} \right\}.$$

Let B_t denote the standard Brownian motion, with no drift, and instantaneous variance of 1. The drift of the underlying process and the choice of search scope σ_t at any point t induce instantaneous drift and standard deviations of the maximum value observed, denoted by $\mu_M(M_t, X_t, \sigma_t)$ and $\tilde{\sigma}_M(M_t, X_t, \sigma_t)$.[16] Furthermore, we denote by $\tilde{\sigma}_{M,X}(M_t, X_t, \sigma_t)$ the induced instantaneous covariance between M_t and X_t. By Ito's lemma, and dropping arguments whenever no confusion is caused, we have

$$dV(Z_t) = \left[\frac{\partial V}{\partial M}\mu_M + \frac{\partial V}{\partial X}\mu + \frac{1}{2}\left(\frac{\partial^2 V}{\partial M^2}\tilde{\sigma}_M^2 + 2\frac{\partial^2 V}{\partial M \partial X}\tilde{\sigma}_{M,X}(M_t, X_t, \sigma_t) + \frac{\partial^2 V}{\partial X^2}\sigma^2\right)\right]dt + \left(\frac{\partial V}{\partial M}\tilde{\sigma}_M + \frac{\partial V}{\partial X}\sigma\right)dB_t.$$

The multiplier of dt is generally called the *controlled infinitesimal generator* of the process Z applied to the function V, and denoted by $\mathcal{A}_Z^\sigma V(Z_t)$. In what follows, it will be useful to denote $\mathcal{A}_Z^\sigma V(Z_t) = \mathcal{A}_M^\sigma V(Z_t) + \mathcal{A}_X^\sigma V(Z_t) + \frac{\partial^2 V}{\partial M \partial X}\tilde{\sigma}_{M,X}(M_t, X_t, \sigma_t)$, where

$$\mathcal{A}_M^\sigma V(Z_t) = \frac{\partial V}{\partial M}\mu_M + \frac{1}{2}\frac{\partial^2 V}{\partial M^2}\tilde{\sigma}_M^2, \quad \text{and}$$
$$\mathcal{A}_X^\sigma V(Z_t) = \frac{\partial V}{\partial X}\mu + \frac{1}{2}\frac{\partial^2 V}{\partial X^2}\sigma^2.$$

Since the Brownian motion has expectation of 0 at any instance, the second term in the sum above falls out in expectation, and we can write the equation succinctly as follows:

$$V(M_t, X_t) = \max\left\{u(M_t), \max_\sigma\left[-c(\sigma) + V(M_t, X_t) + \mathcal{A}_Z^\sigma V(Z_t)\right]\right\}.$$

Now, subtracting $V(M, X)$ from both sides and noticing that maximization over σ has no bearing on the already observed maximum value M, allows us a simplification:

$$0 = \max_{\sigma_t}\{u(M_t) - V(Z_t), \mathcal{A}_Z^{\sigma_t} V(Z_t) - c(\sigma_t)\}. \tag{1}$$

This last equality is Hamilton-Jacobi-Bellman (HJB) equation.

Looking at the HJB, if $V(Z_t) > u(M_t)$, it is strictly optimal to continue. Therefore, in that region, the term $\mathcal{A}_Z^{\sigma_t} V(Z_t) - c(\sigma_t)$ governs the agent's decisions. If, on the other hand, $V(Z_t) < u(M_t)$, it is strictly optimal to stop. The region in which $V(Z_t) = u(M_t)$ defines the *stopping boundary*. This equality implicitly defines X as a function of M at the stopping boundary. It will be useful for us to write the stopping boundary as the set $\{(X, M) : X = g(M)\}$ for the corresponding function $g(\cdot)$. By definition, at the boundary, we have $V(g(M_t), M_t) = u(M_t)$. This is often referred to as *value matching*.

[16] Standard arguments imply that the instantaneous drift and variance of the maximum do not depend on historical levels of the agent's control, the past values of the observed process, or prior maximum values.

Since u is Lipschitz continuous and σ is chosen from a compact interval, it follows that $V(Z) = V(M, X)$ is smooth (see, e.g., Yong and Zhou (1999)). This implies what is often termed *smooth pasting*, namely $V_x(g(M), M) = u_x(M) = 0$.[17] In particular, this implies that the stopping boundary $g(\cdot)$ is differentiable.

While the HJB necessarily holds at an optimal continuous solution, the reverse is not guaranteed in general. For the cases analyzed in this paper, the reverse indeed holds using standard, textbook verification results (see, e.g., Fleming and Rishel (2012)).

7.2 Proofs of Results

Proof of Proposition 1: In order to prove the proposition, we show that:

1. *If $M_t > X_t$, then $A_Z^{\sigma_t} = A_X^{\sigma_t} = \frac{1}{2}(\sigma_t)^2 \frac{\partial^2}{\partial X^2}$.*
2. *If $M_t = X_t$, then $\frac{\partial V}{\partial M} = 0$.*

For part 1, whenever $M_t > X_t$, an infinitesimal change in X_t has no effect on M_t and the formula for $\mathcal{A}_Z^{\sigma_t}$ follows. The formula for $\mathcal{A}_X^{\sigma_t}$ follows directly from the definition since, in our environment, the governing process has no drift.

For part 2, we denote, for any function W of Z, for any t:

$$\mathbf{E}_{m,m}\left(W(M_t, X_t)\right) = \lim_{\substack{r \to t \\ r < t}} \mathbf{E}\left(W(M_t, X_t) | M_r = X_r = m\right).$$

It is well known that (see, e.g., Fleming and Rishel (2012)):

$$\mathbf{E}_{m,m} V(M_t, X_t) = \mathbf{E}_{m,m}\left(\int_0^t \mathcal{A}_Z^{\sigma_r} V(M_r, X_r) dr\right).$$

Using Ito's Lemma,

$$\begin{aligned}
\frac{\mathbf{E}_{m,m} V(M_t, X_t) - V(m,m)}{t} &= \mathbf{E}_{m,m}\left(\frac{1}{t}\int_0^t \mathcal{A}_Z^{\sigma_r} V(M_r, X_r) dr\right) \\
&+ \mathbf{E}_{m,m}\left(\frac{1}{t}\int_0^t \frac{\partial V}{\partial M}(M_r, X_r) dM_r\right. \\
&+ \frac{\partial V(m,m)}{\partial M}\left(\lim_{t \to 0} \frac{\mathbf{E}_{m,m}(M_t - m)}{t}\right).
\end{aligned}$$

[17] The smooth-pasting condition, together with our smoothness assumptions on the utility u, imply that, in fact, the function g is differentiable, which we use below.

Notice that $\mathbf{E}_{m,m}(M_r - m) = \sigma_r$ for all r. Since $\sigma_r \geq \underline{\sigma} > 0$ for all r, it follows that $\lim_{t\to 0} \frac{E_{m,m}(M_t - m)}{t}$ goes to infinity. Therefore, it must be the case that $\frac{\partial V(m,m)}{\partial M} = 0$. ■

Recall the HJB identifying the solutions to our problem, equation (1). The second term corresponds to the continuation choice of search. From Proposition 1 above, we can substitute $\mathcal{A}_X$ for $\mathcal{A}_Z$. Our HJB can then be written as follows:

$$0 = \max_{\sigma_t}\{u(M_t) - V(Z_t), \frac{1}{2}(\sigma_t)^2 \frac{\partial^2 V(Z_t)}{\partial X^2} - c(\sigma_t)\}. \tag{2}$$

We proceed in two steps. First, we identify a recursive formulation of the value function. Using Proposition 1, we identify the optimal control in Claim 1. Then, in Claim 2, we show that the optimal stopping boundary can be derived as the solution of an ordinary differential equation and provide its characterization. An alternative formulation of the ordinary differential equation offered in Claim 3 is useful for a variety of classes of utility functions. The optimal stopping boundary for linear utilities is described in Claim 4, while the solutions for CRRA utilities and logarithmic utilities, which relate to the analogous discounted problem, are relegated to the Online Appendix.

Let $\left\{\mathcal{F}_t^X\right\}_t$ denote the filtration generated by X. A *control adapted to* $\left\{\mathcal{F}_t^X\right\}_t$, also termed feedback control, is a control that is measurable with respect to the filtration $\left\{\mathcal{F}_t^X\right\}_t$. We often omit the explicit reference to the filtration generated by X and refer to such a control as an *adapted control*. Denote by $X_{[0,t]}$ the full path of X_s in the time interval $[0,t]$, namely $\{X_s | s \in [0,t]\}$.

Let σ_t be an arbitrary adapted control. Consider the following problem of choosing an optimal control and optimal stopping, allowing here for general utility functions satisfying the smoothness restrictions imposed in Section 7.1:

$$V(M,X) = \sup_{\tau} \mathbf{E}(u(M_\tau) - \int_0^\tau c(\sigma_t)dt)$$

$$\text{subject to}$$

$$dX_t = \sigma_t dB_t$$

As described in our background section, at the point of stopping, the agent's utility from the achieved maximum coincides with her continuation value: $u(M_\tau) = V(M_\tau, X_\tau)$.

Furthermore, the stopping time has to be of the form $\tau^* = \inf\{t \geq 0 : X_t \leq g(M_t)\}$ for some differentiable function g and the optimal control will be of the form $\sigma(M,X)$. We now use our simplified HJB equation, captured in (2) together with the smooth pasting restrictions to establish the following three constraints:

$$\begin{aligned} &\frac{(\sigma_t)^2}{2}\frac{\partial^2 V}{\partial X^2} = c(\sigma_t) \text{ for } g(M) < x < M && \text{(Continuation Region)} \\ &V(M,X)|_{x=g(M)} = M && \text{(Value Matching)} \\ &\frac{\partial V(M,g(M))}{\partial X} = 0 && \text{(Smooth Pasting).} \end{aligned}$$

Our next goal is to characterize $\sigma(M,X)$ and $g(\cdot)$. As standard, define the scale function of a diffusion, characterized by drift $\mu(X)$ and $\sigma(X)$ as:

$$S(x) = \int_0^x e^{-\int_0^y \frac{2\mu(z)}{(\sigma(z))^2}dz} dy.$$

Define the speed measure of such a diffusion as:

$$m(dx) = \frac{2dx}{S'(x)(\sigma(x))^2}.$$

Since our process has no drift, the corresponding scale function $S(x)$ satisfies $S(x) = x$ for any adapted control. That is, our process is on the so-called *natural scale*. From Revuz and Yor (2013), for any stopping rule of the form $\tau_{a,b} = \tau_a \wedge \tau_b$, where $\tau_a = \inf\{t \geq 0 : X_t = a\}$ and $\tau_b = \inf\{t \geq 0 : X_t = b\}$, for any $a \leq x \leq b$, we have

$$\begin{aligned} P_{X=x}(X_{\tau_{a,b}=a}) &= \frac{S(b)-S(x)}{S(b)-S(a)} = \frac{b-x}{b-a}, \text{ and} \\ P_{X=x}(X_{\tau_{a,b}=b}) &= \frac{S(x)-S(a)}{S(b)-S(a)} = \frac{x-a}{b-a}. \end{aligned} \tag{3}$$

From the scale function, we can also identify the Green function on $[a,b]$ as

$$\begin{aligned} G_{a,b}(x,y) &= \begin{cases} \frac{(S(b)-S(x))(S(y)-S(a))}{S(b)-S(a)} & \text{if } a < y < x < b \\ \frac{(S(b)-S(y))(S(x)-S(a))}{S(b)-S(a)} & \text{if } a < x < y < b \end{cases} \\ &= \begin{cases} \frac{(b-x)(y-a)}{b-a} & \text{if } a < y < x < b \\ \frac{(b-y)(x-a)}{b-a} & \text{if } a < x < y < b. \end{cases} \end{aligned}$$

It is well known—again, see, e.g. Revuz and Yor (2013)—that for any function f we have

$$\mathbf{E}_{X=x}\left(\int_0^{\tau_{a,b}} f(x)dt\right) = \int_a^b f(y)G_{a,b}(x,y)m(dy), \tag{4}$$

where $m(dx)$ is the speed measure of the diffusion X defined above.

Now, consider a stopping time of the form

$$\tau_{g(M),M} = \inf\{t \geq 0 : X_t \notin (g(M), M)\}.$$

This stopping time involves an upper bound, which we will use for a recursive description of the value function. The lower bound corresponds to our stopping boundary. For any current pair (M, X), we are interested in

$$V(M,X) = \mathbf{E}_{X=x}\left(V(M_{\tau_{g(M),M}}, X_{\tau_{g(M),M}}) - \int_0^{\tau_{g(M),M}} [c(\sigma)]dt\right).$$

Start with the first term in this formulation, which captures the expected value from stopping. In the stopping rule identified above, if the upper bound is reached, the agent continues her search and receives $V(M,M)$. If the lower bound is reached, the agent receives $u(M)$. Until one of the bounds is reached M remains constant so $\sigma(M,X)$ can only vary according to X. Multiplying these outcomes with their respective probabilities we have:

$$\mathbf{E}_{X=x}(V(M_{\tau_{g(M),M}}, X_{\tau_{g(M),M}})) = P_{X=x}(X_{\tau_{g(M),M}} = M)V(M,M) + P_{X=x}(X_{\tau_{g(M),M}} = g(M))u(M).$$

Using the formulations from equations (3), we can write:

$$E_{X=x}(V(M_{\tau_{g(M),M}}, X_{\tau_{g(M),M}})) = V(M,M)\frac{X-g(M)}{M-g(M)} + u(M)\frac{M-X}{M-g(M)}.$$

From equation (4), we have that the second term in the formulation of $V(M,X)$ can be written as:

$$E_{X=x}(-\int_0^{\tau_{g(M),M}} [c(\sigma(x_t))]d_t) = -\int_{g(M)}^{M} G_{g(M),M}(x,y)(c(\sigma(M,y)))m(dy).$$

Thus, we have that

$$V(M,X) = u(M)\frac{M-X}{M-g(M)} + V(M,M)\frac{X-g(M)}{M-g(M)} - \int_{g(M)}^{M} G_{g(M),M}(x,y)(c(\sigma(M,y)))m(dy). \quad (5)$$

Reorganizing the above we get

$$V(M,M) - u(M) = \frac{M-g(M)}{X-g(M)}\left(V(M,X) - u(M)\int_{g(M)}^{M} G_{g(M),M}(x,y)(c(\sigma(M,y)))m(dy)\right).$$

We can now use the smooth-pasting conditions to pin down $V(M,X)$.

First, letting x approach $g(M)$, we have

$$\lim_{X \to g(M)} \frac{(V(M,X) - u(M))}{(x - g(M))} X - M = V_X(M, g(M))(M - g(M)).$$

By smooth pasting, we have $V_X(M, g(M)) = 0$.

Similarly, as x approaches $g(M)$, the Green function $G_{g(M),M}(x,y)$ reduces to $\frac{(M-y)(y-g(M))}{(M-g(M))}$. Thus,

$$V(M,M) - u(M) = \int_{g(M)}^{M} (M - y)(c(\sigma(M,y)))m(dy).$$

Plugging $V(M,M)$ back into the equation (5) yields, after algebraic manipulations,

$$V(M,X) = u(M) + \int_{g(M)}^{X} (X - y)(c(\sigma(M,y))) \frac{2}{(\sigma(M,y))^2} dy.$$

The proof of Proposition 2 will follow from the ensuing four claims.

Claim 1: *In the optimal policy, the search scope σ^* is constant and, when interior, solves:*

$$\frac{2c(\sigma^*)}{c'(\sigma^*)} = \sigma^*.$$

Proof of Claim 1: Consider the continuation part of the HJB,

$$\sup_{\sigma_t} \{\mathcal{A}_Z^{\sigma_t} V(Z_t) - c(\sigma_t)\}.$$

Using Proposition 1, this reduces to

$$\sup_{\sigma_t} \{\frac{\sigma_t^2}{2} \frac{\partial^2 V(M,X)}{\partial X^2} - c(\sigma_t)\}.$$

Replacing the sup with the appropriate first-order condition we have

$$0 = \sigma(x) \frac{\partial^2 V(M,X)}{\partial X^2} - c'(\sigma(M,x)).$$

To fully characterize the first-order condition, we differentiate our expression for $V(M,X)$ with respect to X to get:

$$\frac{\partial V(M,X)}{\partial X} = \int_{g(M)}^{X} \frac{2c(\sigma(M,y))}{\sigma^2(M,y)} dy,$$

$$\frac{\partial^2 V(M,X)}{\partial X^2} = \frac{2c(\sigma(M,X))}{\sigma^2(M,X)}.$$

Plugging in the second derivative to the maximized HJB yields

$$\sigma(M,X) = \frac{2c(\sigma(M,X))}{c'(\sigma(M,X))}$$

This last equality does not depend on M or x. It follows that optimal search scope σ^* is constant and a solution to the above equation whenever interior. ■

Claim 2: *The optimal stopping boundary solves the following ordinary differential equation (ODE):*

$$g'(M) = \frac{u'(M)(\sigma^*)^2}{2(c(\sigma^*))(M-g(M))}.$$

Proof of Claim 2: Since $g(M)$ is differentiable, equation (5) is satisfied if and only if $g(M)$ solves the following ODE

$$g'(M) = \frac{u'(M)(\sigma(M,g(M)))^2}{2(c(\sigma(M,g(M))))(M-g(M))}.$$

Plugging in σ^* from Claim 1 implies that the optimal stopping boundary satisfies the following ODE:

$$g'(M) = \frac{u'(M)(\sigma^*)^2}{2(c(\sigma^*))(M-g(M))}.$$

■

As mentioned in the text, analysis following Peskir (1998) shows that the optimal stopping boundary is the maximal solution $g(M) \leq M$ satisfying this ODE. This is a non-linear ODE, which is generally not straightforward to solve in analytical form.[18] Nonetheless, in Claim 3 below, we illustrate this ODE's reduction to an alternative ODE that is more amenable to various classes of utilities studied in the economics literature. In the Online Appendix, we provide the proof of Claim 3 as well as its application to the case of CRRA utilities, where stopping boundaries can be characterized analytically. Nonetheless, the case of linear utilities turns out to be remarkably simple and is described in Claim 4.

Claim 3: *Let $w(\cdot)$ be the solution of the following Abel equation of the second kind:*

$$w(M)w'(M) - w(M) = \frac{\sigma^2}{2c}u'(M).$$

[18] Peskir (1998) identified an equivalent ODE for the case of search without control. He notes the difficulty in providing a general solution and states "to the best of our knowledge the equation... has not been studied before, and... we want to point out the need for its investigation."

The optimal stopping boundary is given by:

$$g(M) = M - H'(M)\frac{(\sigma^*)^2}{2c(\sigma^*)},$$

where $\frac{4c}{\sigma^2}w(M) = 2\sqrt{\frac{4c}{\sigma^2}(H(M) - u(M))}$.

Claim 4: When $u(M) = M$, the optimal stopping boundary is given by:

$$g(M) = M - \frac{(\sigma^*)^2}{2c(\sigma^*)}.$$

Proof of Claim 4: From Claims 1 and 2, it follows that the optimal stopping boundary satisfies:

$$g'(M) = \frac{(\sigma^*)^2}{2(c(\sigma^*))(M - g(M))}.$$

This ODE has a unique solution that satisfies the value-matching condition. It is straightforward to verify the claim.[19] ■

Proof of Proposition 2: The proof follows directly from Claims 1 and 4. ■

Proof of Proposition 3: First, observe that given Proposition 2, we can write the optimal stopping time as:

$$\tau^* = \inf\left\{t \geq 0 : X_t \leq M_t - \frac{(\sigma^*)^2}{2c(\sigma^*)}\right\}.$$

Rearranging terms leads to:

$$\tau^* = \inf\left\{t \geq 0 : M_t - X_t \geq \frac{(\sigma^*)^2}{2c(\sigma^*)}\right\}.$$

This stopping time is commonly referred to as a "drawdown stopping time" of a Brownian Motion, where the size of the "drawdown" is $d \equiv \frac{\sigma^*}{2c^*}$. From Taylor et al. (1975), the joint moment generating function and Laplace transform of X_{τ^*} and τ^* is given by:

$$E[e^{X_{\tau^*} - c(\sigma^*)\tau^*}] = \frac{\beta e^{-d}}{\beta \cosh(\beta d) - \sinh(\beta)},$$

where $\beta = \sqrt{2c(\sigma^*)/\sigma^2}$.

[19] One can explicitly derive the solution by consider the ODE corresponding to the inverse function $M(g)$, which is a first-order linear and homogeneous ODE that has well-known solutions.

The characterization of the distribution of M_{τ^*} then follows since $M_{\tau^*} = X_{\tau^*} + \frac{\sigma^*}{2c(\sigma^*)}$. Again, from Taylor et al. (1975), following conventional techniques, some moments as well as the distributions of M_τ^* and τ^*, are readily identified. In particular, we have

$$\mathbf{E}(\tau^*) = \frac{d^2}{(\sigma^*)^2} \textit{ and } \mathbf{E}(M_\tau^*) = d.$$

In addition, the distribution of M_τ^* is a standard exponential distribution with mean d. Furthermore, the distribution of the maximal value does not depend on the calendar time at which search stops. That is, for any t_1 and t_2,

$$\mathbf{E}(M_\tau^*) = \mathbf{E}(M_\tau^* | \tau^* = t_1) = \mathbf{E}(M_\tau^* | \tau^* = t_2) = d.$$

■

Proof of Proposition 2*: The proof is analogous to the proof of Proposition 2. Since the agent faces a flow expense of $c(\sigma) - w$ at any point in time, similar arguments to those in the proof of Claim 1 imply that

$$\sigma^* = \frac{2(c(\sigma^*) - w)}{c'(\sigma^*)}.$$

From Claim 2, we have that:

$$g'(M) = \frac{\alpha(\sigma^*)^2}{2(c(\sigma^*) - w)(M - g(M))}$$

The stopping boundary is again linear and similar analysis to that in Claim 4 yields:

$$g(M_t) = M_t - \frac{\alpha(\sigma^*)^2}{2(c(\sigma^*) - w)}.$$

■

Proof of Proposition 3*: The expressions describing $\mathbf{E}(M_\tau)$ and $\mathbf{E}(\tau)$ follow directly from the proof of Proposition 3, where the drawdown size is now given by $\frac{\alpha(\sigma^*)^2}{2(c(\sigma^*)-w)}$. The expected returns that the principal and agent receive follow from immediately. ■

Proof of Proposition 4: The principal's problem, taking the agent's solution from proposition 2* as given, can be written as:

$$\max_{\alpha,w} \frac{(1-\alpha)\alpha\sigma^{*2}}{2(c(\sigma^*) - w)} - w\left(\frac{\alpha\sigma^*}{2(c(\sigma^*) - w)}\right)^2$$

subject to

$$\sigma^* = \frac{2(c(\sigma^*) - w)}{c'^*)}$$

The optimal scope of search σ^* is pinned down uniquely by the choice of w. It follows that if the principal induces a search scope of σ, the wages she needs to offer are given by:

$$w(\sigma) = c(\sigma) - \frac{\sigma c'(\sigma)}{2}.$$

Using this induced wage, we can rewrite the objective of the principal as a standard optimization problem:

$$\max_{\alpha,\sigma} \frac{(1-\alpha)\alpha\sigma}{c'(\sigma)} - \left(c(\sigma) - \frac{\sigma c'(\sigma)}{2}\right)\left(\frac{\alpha}{c'(\sigma)}\right)^2.$$

We soon show the conditions under which the first-order condition approach is valid. When it is, taking the first-order conditions and setting them to 0 simplifies to:

$$\begin{aligned}
&\text{w.r.t} \quad \sigma: \quad 4\alpha c(\sigma)c''(\sigma) + (2-3\alpha)c'^2 + (\alpha-2)\sigma c'(\sigma)c''(\sigma) = 0,\\
&\text{w.r.t} \quad \alpha: \quad \alpha = \frac{\sigma c'(\sigma)}{2c(\sigma) + \sigma c'(\sigma)}.
\end{aligned}$$

Using the expression generated for α in the constraint pertaining to σ and simplifying yields:

$$\sigma = \frac{4c(\sigma)}{(\sigma c''(\sigma) + c'(\sigma))}.$$

The components of the Hessian corresponding to the principal's objective are given by:

$$\begin{aligned}
f_{11} &\equiv -\frac{2\left(c(\sigma) - \frac{1}{2}\sigma c'(\sigma)\right)}{c'^2} - \frac{2\sigma}{c(\sigma)},\\
f_{22} &\equiv \frac{1}{2}\alpha\left(-\frac{4(\alpha-1)\sigma c'^2}{c(\sigma)^3} + \frac{2(\alpha-1)(\sigma c''(\sigma) + 2c'(\sigma))}{c(\sigma)^2}\right.\\
&\quad \left. + \frac{4\alpha c(\sigma)\left(c'''(\sigma)c'(\sigma) - 3c''^2\right)}{c'^4} + \frac{\alpha\left(2\sigma c''^2 + c'(\sigma)(4c''(\sigma) - \sigma c'''(\sigma))\right)}{c'^3}\right),\\
f_{12} &\equiv \frac{(2\alpha-1)\sigma c'(\sigma)}{c(\sigma)^2} + \frac{4\alpha c(\sigma)c''(\sigma)}{c'^3} - \frac{\alpha(\sigma c''(\sigma) + c'(\sigma))}{c'^2} + \frac{1-2\alpha}{c(\sigma)}.
\end{aligned}$$

The first-order approach is valid whenever

$$f_{11}f_{22} - (f_{12})^2 \geq 0.$$

In particular, when costs are linear, $c(\sigma) = -a + b\sigma$, with $a, b > 0$, this condition holds as long as $b > a/3$

Whenever the first-order approach is not valid, the principal chooses a boundary solution that, in turn, induces a boundary search scope for the agent. ■

References

Aghion, P. and J. Tirole (1994): "The management of innovation," *The Quarterly Journal of Economics*, 109, 1185–1209.

Bergemann, D. and J. Valimaki (2008): *Bandit problems*, The New Palgrave Dictionary of Economics, 2nd ed. Macmillan Press.

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

Callander, S. and N. Matouschek (2019): "The risk of failure: Trial and error learning and long-run performance," *American Economic Journal: Microeconomics*, 11, 44–78.

Fleming, W. H. and R. W. Rishel (2012): *Deterministic and stochastic optimal control*, vol. 1, Springer Science & Business Media.

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

Guo, Y. (2016): "Dynamic delegation of experimentation," *American Economic Review*, 106, 1969–2008.

Halac, M., N. Kartik, and Q. Liu (2016): "Optimal contracts for experimentation," *The Review of Economic Studies*, 83, 1040–1091.

Jensen, R. and M. Thursby (2001): "Proofs and prototypes for sale: The licensing of university inventions," *American Economic Review*, 91, 240–259.

Jovanovic, B. and R. Rob (1990): "Long waves and short waves: Growth through intensive and extensive search," *Econometrica: Journal of the Econometric Society*, 1391–1409.

Manso, G. (2011): "Motivating innovation," *The Journal of Finance*, 66, 1823–1860.

Miller, R. A. (1984): "Job matching and occupational choice," *Journal of Political economy*, 92, 1086–1120.

MOSCARINI, G. AND L. SMITH (2001): "The optimal level of experimentation," *Econometrica*, 69, 1629–1644.

MURPHY, G. M. (2011): *Ordinary differential equations and their solutions*, Courier Corporation.

OLSZEWSKI, W. AND R. WEBER (2015): "A more general Pandora rule?" *Journal of Economic Theory*, 160, 429–437.

PANAYOTOUNAKOS, D. E. AND D. C. KRAVVARITIS (2006): "Exact analytic solutions of the Abel, Emden–Fowler and generalized Emden–Fowler nonlinear ODEs," *Nonlinear analysis: real world applications*, 7, 634–650.

PESKIR, G. (1998): "Optimal stopping of the maximum process: The maximality principle," *Annals of Probability*, 1614–1640.

——— (2005): "Maximum process problems in optimal control theory," *International Journal of Stochastic Analysis*, 2005, 77–88.

PISSARIDES, C. A. (1984): "Search intensity, job advertising, and efficiency," *Journal of labor Economics*, 2, 128–143.

——— (2000): *Equilibrium unemployment theory*, MIT press.

QUAN, D. C. AND J. M. QUIGLEY (1991): "Price formation and the appraisal function in real estate markets," *The Journal of Real Estate Finance and Economics*, 4, 127–146.

REVUZ, D. AND M. YOR (2013): *Continuous martingales and Brownian motion*, vol. 293, Springer Science & Business Media.

ROBBINS, H. (1952): "Some aspects of the sequential design of experiments," *Bulletin of the American Mathematical Society*, 58, 527–535.

ROGERS, L. C. G. AND D. WILLIAMS (2000): *Diffusions, Markov processes and martingales: Volume 2, Itô calculus*, vol. 2, Cambridge university press.

Root, T. E. (1979): "Economic Influences on the History and Development of the Mining Joint Venture as a Vehicle for Hard Mineral Exploration and Development," *Natural Resources Lawyer*, 491–510.

Taylor, H. M. et al. (1975): "A stopped Brownian motion formula," *The Annals of Probability*, 3, 234–246.

U.S. Department of Agriculture, F. S. (1995): "Anatomy of a mine from prospect to production," Tech. rep.

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

Wong, Y. F. (2020): "Spatial Experimentation," *Available at SSRN 3611751*.

Yong, J. and X. Y. Zhou (1999): *Stochastic controls: Hamiltonian systems and HJB equations*, vol. 43, Springer Science & Business Media.

Zhilkin, N. (1961): "Feasible New Methods of Prospecting and Exploration," *International Geology Review*, 3, 71–77.