NBER WORKING PAPER SERIES

RETROSPECTIVE SEARCH:
EXPLORATION AND AMBITION ON UNCHARTED TERRAIN

Can Urgan
Leeat Yariv

Working Paper 29127
<http://www.nber.org/papers/w29127>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2021

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, University of Oxford, University College London, Toulouse School of Economics, Org Econ workshop 2020, and EC'21. We gratefully acknowledge financial support from the National Science Foundation through grants SES-1629613 and SES-1949381. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Can Urgan and Leeat Yariv. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Retrospective Search: Exploration and Ambition on Uncharted Terrain
Can Urgan and Leeat Yariv
NBER Working Paper No. 29127
August 2021
JEL No. C61,C73,D25,D83

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, a standard exponential discounter, decides the breadth and length of search. We fully characterize the optimal search policy. The optimal search scope is U-shaped, with the agent searching most ambitiously when approaching a breakthrough or when nearing search termination. A drawdown stopping boundary is optimal, where the agent ceases search whenever current observations fall a constant amount below the maximal achieved alternative. We also show special features that emerge from contracting with a retrospective searcher.

Can Urgan
Department of Economics
Princeton University
Julis Romo Rabinowitz
Princeton, NJ 08544
curgun@princeton.edu

Leeat Yariv
Department of Economics
Princeton University
Julis Romo Rabinowitz
Princeton, NJ 08544
and NBER
lyariv@princeton.edu

A data appendix is available at <http://www.nber.org/data-appendix/w29127>

1 Introduction

The theory of search takes center stage in the economics of information and uncertainty. Many economic decisions are preceded by a dynamic process of discovery that generates an alternative deemed best. When should search stop and what outcomes would it deliver? How ambitious should search be depending on where discoveries stand? The canonical models used to answer these key questions mostly assume independent search outcomes.¹ Independence simplifies dramatically the analysis of search—in its basic form, the distribution of possible discoveries is assumed fixed over time. Recall has a limited role and the agent need only look one step ahead, making the optimization problem effectively static. Nonetheless, in many settings, search follows a path, and prior discoveries not only serve as a valuable fallback, but also pave the way for future discoveries. Without the technical convenience afforded by independence, models allowing for correlation assume short-lived, myopic agents and an exogenously fixed search scope.² We provide a framework for studying search by long-lived agents when observations are correlated over time, and when search scope can be adjusted dynamically. We fully characterize the optimal search policy and analyze how it affects agency relationships.

Correlation over the path of search is relevant for numerous applications. 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 such settings, 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 such a process *retrospective search*.

Formally, 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

¹For a survey of the literature, see McCall and McCall (2007).

²For example, see Callander (2011) and Garfagnini and Strulovici (2016).

between discoveries: one moment’s observation forms the expectation for any future moment.³ At every moment, the agent, who is a standard exponential discounter, chooses her search scope or ambition, captured by the momentary variance of the underlying process. We think of more ambitious agents—be it entrepreneurs, researchers, politicians, or shoppers—as ones that inspect a broader set of alternatives, generating greater diversity of outcomes they might observe. Investing more time at work or more financial resources on hardware and workforce increase the variance of outcomes: they either lead to greater accomplishments, or greater losses. We assume the cost of search is increasing and convex in search scope. Whenever the agent stops searching, she can recall all her observations retrospectively and pick the highest, which constitutes her search outcome.⁴

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 discounted maximal observation from her search minus the cumulative discounted search costs spent.

In the optimal policy, the agent chooses a search scope that depends on the distance between the current discovery and her best discovery yet. The agent’s optimal search scope is U-shaped. The optimal search scope is highest when approaching a transition point: either accomplishing a new peak discovery—a breakthrough—or when nearing search termination. This observation is in line with the vast literature on the so-called *goal-gradient hypothesis*. First introduced by [Hull \(1932\)](#) using rat experiments, the goal-gradient hypothesis suggests a tendency for animals and humans to increase effort as a goal approaches. [Kivetz et al. \(2006\)](#) resurrected the hypothesis with an array of human experiments. It has been used to explain effort patterns in a variety of contexts, from marketing to athletics. While much of the literature on the goal-gradient hypothesis views the effect as arising from deeply-ingrained psychological instincts, it emerges in our model as the result of an optimal search protocol.

The optimal search scope depends on search costs, with more log-convex costs associated with less ambitious search. More patient agents are associated with flatter search scopes. In fact, absent discounting, a constant search scope is optimal. In practice, many product development teams, political task forces, and the like consider long activity horizons and arguably exhibit less pronounced

³In most search applications, the mere passage of time does not improve outcomes. Therefore, in our benchmark model we assume no drift. Nonetheless, we describe impacts of drift in some of our discussions.

⁴This implicitly implies that the agent is risk neutral when it comes to her search outcomes. This greatly simplifies our exposition greatly. We discuss how our analysis can be extended to risk-averse agents in Section 6.1 and in the Supplementary Materials.

discounting. Set properly, the optimal search scope for such searching units 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 mine 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 remarkably simple form. Specifically, the agent halts exploration whenever her current discovery falls below some *fixed* distance from the best discovery she's made, where that fixed distance is constant over time. The threshold for stopping evolves over time: as the agent accumulates greater discoveries, she becomes more demanding and stops more willingly. This result suggests that product development, for example, should stop when innovations do not appear sufficiently promising, where promise is assessed in relation to the best option yet. The distance between a current innovation that leads to search termination and the best innovation idea the search has produced does not change with time or with the realized discovery path.

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. Absent discounting, 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 prob-

lems have received surprisingly limited attention in both the social sciences and the mathematics literature, particularly in the presence of discounting. 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

Weitzman (1979)’s classic paper 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. Under the optimal strategy, 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 model has also 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.

Modeling correlation over time using a Brownian path in a search setting is inspired by Callander (2011).⁵ 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. Urgan and Yariv (2021) consider a similar setting to this paper’s in which an agent is constrained to search within a fixed amount of

⁵See Jovanovic and Rob (1990) for axiomatic foundations of search discoveries following a Brownian path.

time. [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 \(2020\)](#) illustrates that exploration hastens after poor outcomes and that search stops when observed outcomes are sufficiently poor.⁶

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 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 et al. \(2016\)](#), 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 and there is no control, see e.g. [Peskir \(1998\)](#) and [Peskir \(2005\)](#). In this literature, agents experience flow costs and no discounting. The techniques we develop allow for the analysis of such optimal stopping problems with both discounting and the inclusion of a control—in our case, the costly search scope. We hope the methods we introduce open the door for further studies in the area.

⁶The 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.

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 X_t and assess the rewards from search through M_t with

$$dX_t = \sigma_t dB_t$$

$$M_t = (\max_{0 \leq r \leq t} X_r \vee M_0),$$

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 observation X_t can stand for the expected value of the discovery made at time t . If search ends at time t , the best discovery, the one yielding M_t , gets implemented.

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 6.2 we discuss the impacts of drift.⁷

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 convex. If $\underline{\sigma} = \bar{\sigma}$, our setting boils down to one in which the agent has fixed search scope she cannot control.

When the search scope is constant over time, $\sigma_t = \sigma^*$ for all t , we can use the Reflection Principle (see, e.g., [Rogers and Williams \(2000\)](#)) to infer that M_t follows the same distribution as

⁷One could also consider alternative processes with additional features: Lévy processes allowing for discrete breakthroughs; Ornstein-Uhlenbeck processes implying mean-reversion; or even Brownian motions with ex-ante uncertain features, compounding learning on top of search. Our approach may be useful for the study of such processes as well.

$|\sigma^* B_t|$. Therefore, 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.⁸ Furthermore, the times at which M_t increases by fixed amounts, hitting $\Delta, 2\Delta$, etc. for any $\Delta > 0$ —which can be thought of as times at which *breakthroughs* are made—follow an exponential distribution. This is in line with assumptions made in the innovation literature, see for example [Kortum \(1997\)](#).

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 \(1984\)](#) and [Pissarides \(2000\)](#). 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. Nonetheless, the impact on the expected returns to search are similar, with higher search scope yielding a higher expected outcomes.

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 be written as:

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

We discuss the impacts of risk attitudes when describing our main results, and in [Section 6.1](#).⁹

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.¹⁰ This highlights an important contrast with the case of independent search outcomes. There, the option value of search continuation is fixed over time. Thus, an

⁸Formally, 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 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 $\mathbf{E}(M_t) = \mathbf{E}(\sigma^* |B_t|) = \sigma^* \sqrt{2t/\pi}$.

⁹In the Supplementary Materials, we also analyze a version of the model in which search scope is fixed and there are no corresponding flow costs.

¹⁰Risk attitudes would only make search outcomes less appealing and would not alter this conclusion. 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.

outcome rejected in the past is never more appealing in the future. Recall is never utilized—the searcher optimally *satisfices* and stops search as soon as she observes a sufficiently high outcome, with no regard to prior realizations.

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 agent’s optimal policy is governed by a stopping boundary $g(M)$ 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:

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

Naturally, $g(M)$ depends on the features of the process: both exogenous and endogenous. That is, holding all of our environment’s parameters fixed, $g(M)$ depends 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. For very low search scopes, the agent would want to stop quickly as there is not much to gain from costly search. For very high search scopes, continuation is prohibitively costly and the agent would terminate search rapidly as well. For intermediate levels of search scopes, the agent may benefit from non-trivial search.

Let $\mathbb{T}_{[a,b]}(Y)$ denote the first time a standard Brownian motion (with instantaneous variance of 1) escapes an interval $[a, b]$ with an initial point $Y \in [a, b]$. With a slight abuse of notation, we suppress the dependence on the starting point Y in what follows. Proposition 1 fully characterizes the optimal search policy.

Proposition 1 (Optimal Search Scope). *For any continuous stopping boundary $g(M)$, the optimal search scope $\sigma^r(M, X)$ solves:*

$$\frac{2c(\sigma^r(M, X))}{c'(\sigma^r(M, X))} \mathbf{E} \left(e^{-r\mathbb{T}_{[g(M), M]}} | X \right) = \sigma^r(M, X).$$

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. The first corresponds to the instantaneous search scope, minute changes in which affect the marginal cost. The second corresponds to the length of search, minute changes in which affect the flow cost. When costs are strictly convex, the agent’s objective is strictly concave, and there is a unique optimal policy. Nonetheless, in general, there could be multiple solutions of the fixed-point equation in the proposition.¹¹ When the left-hand-side of the equation is monotonically decreasing, say when costs are log-convex, a unique fixed-point solution is guaranteed.

The expected discounted time a standard Brownian takes to escape the interval $[g(M), M]$ starting from observation X can also be described explicitly:¹²

$$\mathbf{E} \left(e^{-r\mathbb{T}_{g(M),M}} | X \right) = \frac{\sinh \left[(M - X)\sqrt{2r} \right] + \sinh \left[(X - g(M))\sqrt{2r} \right]}{\sinh \left[(M - g(M))\sqrt{2r} \right]}$$

Thus, $\sigma^r(M, X)$ is symmetric around $(M + g(M))/2$. In fact, we have the following corollary to Proposition 1.

Corollary 1 (Features of Optimal Search Scopes). *For any continuous stopping boundary $g(M)$ and any $r > 0$, the optimal search scope $\sigma^r(M, X)$ symmetric around $\frac{M+g(M)}{2}$, maximized at the boundaries M and $g(M)$ and minimized at $\frac{M+g(M)}{2}$, and decreasing in r . When $r = 0$, the optimal search scope is constant, $\sigma^0(M, X) = \sigma^0$, and when interior, solves:*

$$\frac{2c(\sigma^0)}{c'(\sigma^0)} = \sigma^0.$$

Furthermore, $\sigma^r(M, M) = \sigma^r(M, g(M)) = \sigma^0$ for any $r > 0$ and M .

With non-trivial discounting, the optimal search-scope is U -shaped. The agent responds to details of her environment, searching more expansively when nearing a breakthrough—namely, when a new maximum is likely to be achieved in short horizon—or when recent results are discouraging and search termination is near. This observation provides a new rationale for the so-called *goal-gradient hypothesis*, suggesting a tendency to exert greater efforts when approaching a goal. A

¹¹The value function can then be used to identify the optimal search scope among those solutions, see the Appendix for details.

¹²The hyperbolic functions $\sinh$ and $\cosh$ are defined as follows:

$$\sinh = \frac{e^x - e^{-x}}{2} \text{ and } \cosh = \frac{e^x + e^{-x}}{2}.$$

large body of experimental literature, going back [Hull \(1932\)](#), who was inspired by an array of rat experiments, provides evidence for the hypothesis. A recent explosion of work followed [Kivetz et al. \(2006\)](#), who documented a collection of human experiments on the matter. The hypothesis has been used to explain effort patterns in a rich set of environments, from marketing, to the workforce, to athletics. Much of this literature views the effect as arising from deeply-ingrained psychological instincts.¹³ In contrast, the response to nearing goals emerges in our model as the result of an optimal search protocol.

Absent discounting, the agent optimally chooses the search scope to be fixed at a level independent of both global and local features of the process: the maximum value observed and the samples recently observed. Large companies face long activity horizons and conceivably less pronounced discounting on a quarter-to-quarter basis. When discounting is less prominent, our result indicates that innovations do not require ongoing changes in, say, development teams in terms of either size or work intensity. Indeed, many large companies do not fluctuate in their research and development expenditures over quarters.¹⁴ One setting for which detailed discovery data is more easily available is that of mineral prospecting. The short horizon over which such explorations take place suggests that discounting may have limited influence on decisions. 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.¹⁵

The intuition for Proposition 1 follows several steps. The first has to do with the impacts of small changes in search scope on the attained maximum. 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

¹³For example, in a 2020 interview for *Scientific American*, Oleg Urminky suggested a link between the goal-gradient hypothesis and present bias.

¹⁴This 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.

¹⁵The 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.

FIGURE 1: Independence of Local Choices on Established Maxima

time, with probability 1 the agent reaches a new value different than the current maximum.¹⁶ Formally, in Lemma A1 of the Appendix, we show that the controlled infinitesimal generator of the two-dimensional process operating on any C^2 function is almost surely equal to the controlled infinitesimal generator of the one-dimensional process X operating on that function.

The above arguments indicate that a marginal change in search scope affects the value of search only through its marginal impacts on local conditions. In particular, such a marginal change has no impact on the maximal value attained. Changing the search scope from 1 to σ at any small interval of time is tantamount to *speeding up* the (standard Brownian) process by a factor of σ^2 . When there is no discounting, on a given path, the agent effectively finds the efficient *speed* to hit the boundary $g(M_t)$ or a new maximum, surpassing 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^0 .

With discounting, the costs need to be adjusted. This adjustment is represented by the scaling factor $\mathbf{E} \left(e^{-r\mathbb{T}_{g(M),M}} | X \right)$. Intuitively, when current observations are distant from either the stopping boundary or the current maximum, foreseen costs close to a transition—either to search termination or to a new maximum—are heavily discounted. The agent then places a high weight on minimizing immediate costs, with less regard to speed. In contrast, when close to the stopping boundary or to a breakthrough, discounting is of less import and the agent places a high weight on the speed

¹⁶In 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 .

at which a transition will be reached, much as in the undiscounted case. As the discount rate increases, and the agent becomes more impatient, immediate costs play a more important role and the optimal search scope declines. Naturally, the optimal stopping boundary changes for different discount rates, which the following proposition characterizes.

Proposition 2 (Optimal Stopping Boundary). *The stopping boundary at any point t with a previously observed maximum M_t is given by: $g^r(M_t) = M_t - d^r$, where $d^0 = \frac{(\sigma^0)^2}{2c(\sigma^0)}$ and d^r solves:*

$$\frac{d^r}{d^0} \left(1 + 2 \frac{\sinh^2(\sqrt{r/2}d^r)}{\cosh(\sqrt{r/2}d^r)} \right) = 1.$$

Proposition 2 asserts that the stopping boundary corresponds to the currently-held maximum of the process minus a *fixed* amount, which naturally depends on search costs. Such stopping boundaries are often referred to as *drawdown stopping boundaries*, with the defining fixed amount termed the *drawdown size*. Going back to our example of product development, this result suggests that product release, occurring when search ceases, 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 for the drawdown nature of the optimal stopping policy is 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 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.

While the intuition, and proof, of the structure of the optimal search scope does not depend on the linearity of the agent's utility, the argument for the optimality of a drawdown stopping boundary most certainly does. In Section 6.1 and in the Supplementary Materials, we discuss an extension to general concave utilities. Our techniques can be directly extended, but the analysis becomes far more intricate. Nonetheless, for constant relative risk aversion (CRRA) utilities with parameter ρ , 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.

The precise derivation of the optimal drawdown size appears in the proposition's proof. Since the hyperbolic functions $\sinh$ and $\frac{\sinh}{\cosh}$ are both increasing, we have:

Corollary 2 (Optimal Retrospective Search without Discounting). *The optimal drawdown size d^r is decreasing in r .*

As discount rates increase, the agent becomes increasingly impatient and the value of future search is lowered. Consequently, the agent is more inclined to stop and demands more promising immediate discoveries to pursue search. The minimal search scope with any observed maximum M , given by $\sigma^r(M, \frac{M+g^r(M)}{2})$, is also decreasing in r . Thus, more impatient agents are less ambitious and quicker to quit.

4.2 Retrospective Search Outcomes

We now turn to the value of search, incorporating both the value of the implemented project and the costs accrued throughout the search process.

Proposition 3 (Expected Values of Retrospective Search). *For any $r \geq 0$, the expected payoff from optimal retrospective search is:*

$$\mathbf{E}(\Pi^r) = \frac{d^r}{d^0} \frac{\sinh(\sqrt{r/2}d^r)}{\sqrt{2r}}.$$

Since the process is symmetric, and can go up or down with equal probability, the agent might hope to gain up to d^r while searching. Formally, 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^* , $M_{\tau^*} - X_{\tau^*} = d^r$. It follows that, for any realized stopping time τ^* ,

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

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 ,

$$\mathbf{E}(M_\tau^*) = \mathbf{E}(M_\tau^* \mid \tau^* = t) = d'.$$

In particular, how long it takes for retrospective search to run its course is not indicative of the resulting expected value of the project.

The costs of search naturally attenuate the expected value of the implemented alternative. Accounting for these costs is generally more involved. To glean some intuition, consider the no-discounting case, where $r = 0$. As already mentioned, scaling of the 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, $\mathbf{E}(\tau^*) = \frac{(d^0)^2}{\sigma^2}$.

Corollary 3 (Expected Values without Discounting). *When $r = 0$, the expected project value upon optimal stopping is $\mathbf{E}(M_\tau^*) = \frac{(\sigma^0)^2}{2c(\sigma^0)}$ and the expected optimal search duration is $\mathbf{E}(\tau^*) = \frac{(\sigma^0)^2}{4c^2(\sigma^0)}$. Consequently, the expected payoff from optimal retrospective search is:*

$$\mathbf{E}(\Pi^*) = \mathbf{E}(M_\tau^*) - c(\sigma^*)\mathbf{E}(\tau^*) = \frac{(\sigma^0)^2}{4c(\sigma^0)}.$$

Corollary 3 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 the corollary 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.¹⁷

When $r = 0$, natural comparative statics emerge. Suppose search scope is exogenously fixed at a constant $\hat{\sigma}$ with associated search cost of $\hat{c}$. Keeping the cost fixed, as search scope increases,

¹⁷With 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, $\mathbf{E}(\Pi^*) = \frac{1}{2}\mathbf{E}(M_\tau^*)$.

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 Proposition 1 illustrates, the optimal policy depends both on the search cost's level and margins. We therefore need to impose further restrictions on cost functions in order to generate clear comparative statics.

As already noted, when cost functions are log-convex, an interior solution is unique.¹⁸ Consider two log-convex cost functions $c_1(\cdot)$ and $c_2(\cdot)$ such that $\frac{c_2(x)}{c_2'(x)} < \frac{c_1(x)}{c_1'(x)}$ for all x . Denote by $\sigma_i^*(\cdot, \cdot)$ and $\mathbf{E}(M_i^*)$ the optimal search scope and expected project value under cost function $c_i(\cdot)$. From Proposition 1, $\sigma_1^* > \sigma_2^*$ for any given stopping boundary. Costs naturally impact the stopping boundary as well. 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 $\gamma > 1$ and suppose there is no discounting, $r = 0$. Our discussion above suggests that the optimal search scope should decline with γ when γ is sufficiently low, namely $\gamma < 2e$. Indeed, the optimal search scope $\sigma^* = \left(\frac{2}{\gamma}\right)^{1/\gamma}$ declines for low values of γ and asymptotes at 1 as γ increases, see 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 value of the project. In particular, for small γ , the expected payoff from retrospective search declines. The reverse occurs for larger γ .

Formally, from Corollary 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

¹⁸The family of log-convex functions is rich and contains the family of exponential functions, Euler's Gamma function, etc.

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

$\mathbf{E}(M_\tau^*) = \frac{1}{2} \left(\frac{2}{e^\gamma} \right)^{2/\gamma}$, which is double the expected payoff from search, and is non-monotonic in γ . It is decreasing initially and then increasing indefinitely to an asymptote, as depicted in panel (c) of Figure 2.

In other words, increases in log-convexity as defined above make the agent less ambitious in terms of search scope, but also prolongs the search. Expected payoffs, however, 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 corresponds 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 $\bar{c}(\cdot)$ then satisfies:

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

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 return to when discussing 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 cost level parameter a and the marginal cost parameter b . Indeed, 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. For simplicity, we restrict our analysis to the case in which there is no discounting for neither the principal nor the agent.

We consider contracts that are comprised of a fixed wage $w \in \mathbb{R}$ 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.¹⁹

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:

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

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.

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 Propositions 1 and 2:

Corollary 4 (Optimal Commissioned Search). *With commission contract (w, α) , the agent’s optimal search scope is constant and solves:*

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

¹⁹See <https://www.sec.gov/edgar/search/> for a full set of reported contracts since 2001.

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 1. The stopping boundary, however, needs to be adjusted relative to that provided in Proposition 2, accounting for 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 Corollary 1 indicates, absent discounting, 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.²⁰ 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

²⁰Indeed, 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.

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.²¹ 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 $\mathbf{E}(\Pi_{w,\alpha}^P)$, and the agent, denoted by $\mathbf{E}(\Pi_{w,\alpha}^A)$. The precise formulation of these expected payoffs naturally generalizes our characterization in Proposition 2. 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 4 (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 $\mathbf{E}(\tau_{w,\alpha}) = \left(\frac{\alpha(\sigma^*)}{2(c(\sigma^*) - w)}\right)^2$. The expected returns for the agent and principal, respectively, are:*

$$\mathbf{E}(\Pi_{w,\alpha}^A) = G(\alpha, c(\sigma^*) - w) \quad \text{and} \quad \mathbf{E}(\Pi_{w,\alpha}^P) = G(1 - \alpha, w),$$

where

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

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:

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

Notice that σ^* is pinned down uniquely by a choice of w . In particular, if the principal wants

²¹Labs, 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.

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:²²

Proposition 5 (Optimal Commission Contract). *Whenever the principal's optimal commission contract (w^*, α^*) guarantees an interior search scope, it satisfies:*

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

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.

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.

²²The proof involves simple algebraic manipulations of the corresponding first-order conditions and is therefore omitted.

Using Proposition 5 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 4 suggests that the expected maximal value of the project is $\mathbf{E}[M_\tau^*] = \frac{8}{9} \frac{a}{b^2}$, while the expected search duration is $\mathbf{E}[\tau^*] = \frac{4}{9b^2}$. 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:

$$\mathbf{E}(\Pi_{w^*, \alpha^*}^A) = \mathbf{E}(\Pi_{w^*, \alpha^*}^P) = \frac{4}{9} \frac{a}{b^2}.$$

In other words, with linear costs, the agent and the principal split the 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 5, second-order terms of the cost function can generally play a role in the contractual solution.²³

²³Nonetheless, an analogous example with quadratic costs generates qualitatively similar messages.

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 U-shaped scope of search. The optimal search scope is flatter, and less responsive to recent discoveries, the more patient the agent is. Absent discounting, it is constant. There is a simple optimal stopping boundary: a retrospective searcher ceases search whenever the observed value is a certain fixed and constant distance below the maximal value observed. That fixed value is declining in the discount rate. 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, 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, for the features of the optimal search scope do 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 stopping boundary.

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. For simplicity, suppose $r = 0$. The smooth-pasting condition 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 Supplementary Material 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 Supplementary Materials, 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 Allowing for Drift

Throughout the paper, we assume the path of discovery has no drift. We do so for two reasons. First, it fits the search applications we have in mind: when looking for new ideas, products to acquire, etc., the mere passage of time does not generally 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, 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, the searcher chooses the maximal possible drift. For high enough costs, the searcher dispenses with drift.

6.3 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 one. 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 discoveries 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.

Similarly, when driving through a new city seeking a restaurant, one can certainly vary the driving speed to cover more ground, and it could be reasonable to drive back and enter a previously-seen restaurant. However, there is no point in going back to a restaurant in order to drive through the same streets again. The path of discoveries—here, the sequence of restaurants—does not change.

A Proofs of Main Results

A.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 an underlying Weiner process X_t that has 0 drift and standard deviation σ , which is controlled by the agent. For simplicity, we assume, as in the paper, that $X_0 = 0$.

We start with the undiscounted problem. 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 $c(\sigma)$. 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 yields:

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

Equivalently,

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

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 $\tilde{\sigma}_M(M_t, X_t, \sigma_t)$.²⁴ 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

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

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

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

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

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

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:

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

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

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, however, $V(Z_t) < u(M_t)$, 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 is useful to write the stopping

²⁴Standard 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.

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*.

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 \(1999\)](#), page 42, Theorem 6.3, and page 275, Theorem 6.2). This implies what is often termed *smooth pasting*, namely $V_x(g(M), M) = u_x(M) = 0$.²⁵ 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., [Yong and Zhou \(1999\)](#), pages 277-278, Theorem 6.6, Case 1).

The HJB above is derived without a discount factor, but discounting effectively translates into a termination rate of the process X , often referred to as the *killing rate*. The discounted problem is very closely related to the undiscounted problem as noted in Chapter 5 of [Itô et al. \(2012\)](#) and Chapter 2 of [Borodin and Salminen \(2012\)](#).

Consider the discounted search problem with discount rate $r > 0$. That is, consider an agent facing the following optimization problem:

$$\max_{\tau, \{\sigma_t\}_{t=0}^{\tau}} \mathbf{E} \left[e^{-r\tau} u(M_{\tau}) - \int_0^{\tau} [e^{-rt} c(\sigma_t)] dt \right]$$

It is well known ([Peskir and Shiryaev \(2006\)](#), chapters 5.4 and 6.3) that for any finite stopping time τ , and any continuous function $c(\cdot)$, the above problem can be equivalently defined for the process $\hat{Z} = (\hat{M}, \hat{X})$, the process $Z = (M, X)$ killed at rate r :

$$\max_{\tau, \{\sigma_t\}_{t=0}^{\tau}} \mathbf{E} \left(u(\hat{M}_{\tau}) - \int_0^{\tau} [c(\sigma_t)] dt \right).$$

We state the following facts about the relationship between the killed process and the unkilld one. A more detailed discussion of these facts can be found in [Borodin and Salminen \(2012\)](#) (pages 27-28). The proofs can be found in various sources including [Itô et al. \(2012\)](#) (pages 179-183).

1. The scale function and the speed measure of the killed process $\hat{X}$ equal the scale function and the speed measure of the unkilld process X .²⁶

²⁵The 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.

²⁶Recall that the scale function of a generic diffusion with drift μ variance σ is given by: $S(x) = \int_0^x e^{-\int_0^y \frac{2\mu(z)}{(\sigma(z))^2} dz} dy$.

2. The (controlled) infinitesimal generator of the process $\hat{Z} = (\hat{M}, \hat{X})$, denoted by $\mathcal{A}_{\hat{Z}}$, equals

$$\mathcal{A}_{\hat{Z}} = \mathcal{A}_Z - r.$$

Since the non-killed state space of $\hat{Z}$ and Z are the same, from here onwards, with a slight abuse of notation, we drop the hats and represent the HJB equation as follows:²⁷

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

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

A.2 Reducing Dimensionality

The following lemma, which we use throughout our analysis, suggests that a marginal change in search scope affects the value of search only through its marginal impacts on local conditions.

Lemma A1 (Reducing Dimensionality) *The infinitesimal generator satisfies the following:*

1. If $M_t > X_t$, then $\mathcal{A}_Z^{\sigma_t} = \mathcal{A}_X^{\sigma_t} = \frac{1}{2}(\sigma_t)^2 \frac{\partial^2}{\partial X^2}$.
2. If $M_t = X_t$, then $\frac{\partial V}{\partial M} = 0$.

That is, at any t , an infinitesimal change in the search scope, the control, σ_t has no effect via the current maximum M_t .

For completeness, we provide a proof below. Alternative proofs of Lemma A1, commonly known as “reflection on the diagonal,” can be found in various sources, including [Dubins et al. \(1994\)](#).

Proof of Lemma A1:

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.

and the speed measure of the same diffusion is given by $m(dx) = \frac{2dx}{S'(x)(\sigma(x))^2}$.

When the drift is equal to 0 the process is in the so called "natural scale" with $S(x) = x$ and $m(dx) = \frac{2dx}{(\sigma(x))^2}$

²⁷The HJB equation with discounting is usually directly derived by simply including the time derivative of the value function while applying Ito's lemma. It usually includes a normalization that leads directly to the second line up to normalization. Here, we use the killed process instead of normalizing to utilize the connection between the killed and unkilld diffusions.

For part 2, it is sufficient to show that for any C^2 function W of Z , for any t such that $X_t = M_t$, $\frac{\partial W(M_t, X_t)}{\partial M} = 0$.

Suppose the observed value at a date normalized to 0 coincides with maximal value: $X_0 = M_0 = M$. For any t , consider $W(M_t, X_t)$. In line with our description in Section A.1 above, applying Ito's formula and taking expectations,

$$\begin{aligned} \mathbf{E}_{M,M}(W(M_t, X_t)) &= W(M, M) + \mathbf{E}_{M,M} \left(\int_0^t \mathcal{A}_X^{\sigma_r} W(M_r, X_r) dr \right) + \mathbf{E}_{M,M} \left(\int_0^t \frac{\partial W(M_r, X_r)}{\partial M} dM_r \right) \\ &\quad + \frac{1}{2} \mathbf{E}_{M,M} \left(\int_0^t \frac{\partial W(M_r, X_r)}{\partial M \partial X} d\langle M_r, X_r \rangle + \int_0^t \frac{\partial^2 W(M_r, X_r)}{\partial M^2} d\langle M_r, M_r \rangle \right). \end{aligned}$$

Consider the terms involving the quadratic variation and quadratic covariance of M , the last two terms in the formula above. Since M_t is (weakly) increasing over any time interval, it has bounded variation, and thus has 0 quadratic variation. Therefore, those terms vanish. Dividing both sides by t , we have:

$$\frac{\mathbf{E}_{M,M}W(M_t, X_t) - W(M, M)}{t} = \frac{1}{t} \mathbf{E}_{M,M} \left(\int_0^t \mathcal{A}_X^{\sigma_r} W(M_r, X_r) dr + \int_0^t \frac{\partial W(M_r, X_r)}{\partial M} dM_r \right).$$

Taking the limit as $t \rightarrow 0$, by Dynkin's formula, the left-hand side converges to the infinitesimal generator of Z . The first term on the right-hand side reduces to the infinitesimal generator of X . Furthermore, $\frac{\partial W(M_r, X_r)}{\partial M} dM_r$ is the first-order term in the Taylor approximation of our function W and, hence, coincides with $\frac{\partial W(M, M)}{\partial M} \left(\lim_{t \rightarrow 0} \frac{\mathbf{E}_{M,M}(M_t - M)}{t} \right)$. Therefore,

$$\mathcal{A}_Z^{\sigma_r} W(M_r, X_r) dr = \mathcal{A}_X^{\sigma_r} W(M_r, X_r) dr + \frac{\partial W(M, M)}{\partial M} \left(\lim_{t \rightarrow 0} \frac{\mathbf{E}_{M,M}(M_t - M)}{t} \right).$$

From the Reflection Principle, $\mathbf{E}_{M,M}(M_t - M)$ is of the order $\sqrt{t}$. It follows that $\lim_{t \rightarrow 0} \frac{\mathbf{E}_{M,M}(M_t - M)}{t}$ goes to infinity. Therefore, it must be that $\frac{\partial W(M, M)}{\partial M} = 0$ for any C^2 function, including V . $\blacksquare$

Recall the HJB identifying the solutions to our problem, equation (1). The second term corresponds to the continuation choice of search. From Lemma A1 above, we can substitute $\mathcal{A}_X$ for $\mathcal{A}_Z$. Our HJB can then be written as follows:

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

A.3 Proofs

We proceed in two steps. First, we illustrate a recursive formulation of the value function. Using Lemma A1, we identify the optimal control. Then, in Lemma A2, we show that the optimal stopping boundary can be derived as the solution of an ordinary differential equation and provide its characterization. The optimal stopping boundary for linear utilities is described in our proof of Proposition 2, while the solutions for CRRA utilities and logarithmic utilities are relegated to the Supplementary Materials.

Let $\{\mathcal{F}_t^X\}_t$ denote the filtration generated by X . A control adapted to $\{\mathcal{F}_t^X\}_t$, also termed feedback control, is a control that is measurable with respect to the filtration $\{\mathcal{F}_t^X\}_t$. 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^r be an arbitrary adapted control when the discount rate is r . For notational simplicity we will suppress r until the final representation and use σ instead. Consider the following problem of choosing an optimal control and optimal stopping for the *killed process* (M_t, X_t) , with killing rate r , allowing for general utility functions satisfying the smoothness restrictions imposed in Section A.1:

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

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 $\tau^* = \inf\{t \geq 0 : X_t \leq g^r(M_t)\}$ for some differentiable function g and the optimal control takes the form of $\sigma(M, X)$. Again, we suppress r until the final representation and denote the stopping boundary by $g(M_t)$. We now use our simplified HJB equation, captured in (2) together with the smooth-pasting restrictions to establish the following three constraints:

$$\frac{(\sigma_t)^2}{2} \frac{\partial^2 V}{\partial X^2} = c(\sigma_t) - rV \text{ for } g(M) < x < M \quad (\text{Continuation Region})$$

$$V(M, X)|_{x=g(M)} = M \quad (\text{Value Matching})$$

$$\frac{\partial V(M, g(M))}{\partial X} = 0 \quad (\text{Smooth Pasting}).$$

Our next goal is to characterize $\sigma(M, X)$ and $g(\cdot)$. 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

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

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 the outcomes in $V(M_{\tau_{g(M),M}}, X_{\tau_{g(M),M}})$ by their respective probabilities, we have:

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

From [Revuz and Yor \(2013\)](#) (pages 304-305, Theorem 3.6 and Corollary 3.8), 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

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

where $S(x)$ denotes the scale function (which is in natural scale due to the lack of drift).

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

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

It is well known—again, see e.g. [Revuz and Yor \(2013\)](#)—that for any function f ,

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

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

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

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

Thus,

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

Reorganizing the above,

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

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

Letting x approach $g(M)$, we have

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

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

From [Borodin and Salminen \(2012\)](#) (Appendix 1, page 105), the Green's function of the killed process is given by:

$$G_{g(M),M}^r(x, y) = \begin{cases} \frac{\sinh(\sqrt{2r}(M-x)) \sinh(\sqrt{2r}(y-g(M)))}{\sqrt{2r} \sinh(\sqrt{2r}(M-g(M)))} & \text{if } M > x > y > g(M) \\ \frac{\sinh(\sqrt{2r}(M-y)) \sinh(\sqrt{2r}(x-g(M)))}{\sqrt{2r} \sinh(\sqrt{2r}(M-g(M)))} & \text{if } M > y > x > g(M) \end{cases}.$$

Letting X approach $g(M)$, we can use the smooth-pasting conditions to conclude

$$V(M, M) = u(M) + (M - g(M)) \int_{g(M)}^M \frac{\sinh(\sqrt{2r}(M-y))}{\sinh(\sqrt{2r}(M-g(M)))} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy. \quad (6)$$

We can then plug this back into (5) with the appropriate Green's function to get

$$\begin{aligned} V(M, X) &= u(M) + (X - g(M)) \int_{g(M)}^M \frac{\sinh(\sqrt{2r}(M-y))}{\sinh(\sqrt{2r}(M-g(M)))} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy \\ &\quad - \int_{g(M)}^M G_{g(M),M}^r(X, y) \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy. \end{aligned} \quad (7)$$

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

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

Using Lemma A1, this reduces to:

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

Replacing the supremum with the appropriate first-order condition,

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

Given its closed-form description, we can then take the second derivative of $V(M, X)$ with respect to $\hat{X}$:

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

We can plug this into the first-order condition above to get:

$$\sigma(M, X) \left(\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \int_{g(M)}^M 2rG_{g(M), M}^r(X, y) \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy \right) = c'(\sigma(M, X)).$$

In order to find the general solution, we re-arrange the continuation HJB as follows:

$$\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)} = \int_{g(M)}^M 2rG_{g(M), M}^r(X, y) \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy \quad (8)$$

Integrating both sides over X between $g(M)$ and M , we achieve:

$$\int_{g(\hat{M})}^{\hat{M}} \left(\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)} \right) dX = \int_{g(M)}^M \int_{g(M)}^M 2rG_{g(M), M}^r(X, y) \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy dX.$$

We can change the order of integration on the right-hand side. We can also factor out $\frac{2c(\sigma(M, y))}{\sigma(M, y)^2}$:

$$\int_{g(M)}^M \left(\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)} \right) dX = \int_{g(M)}^M \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} r \int_{g(M)}^M 2G_{g(M), M}^r(X, y) dX dy.$$

By definition, the Green's function is symmetric. That is, $G_{g(M), M}^r(X, y) = G_{g(M), M}^r(y, X)$. The inner integral is the integral of the Green's function, with speed measure 1. Letting $\mathbb{T}_{[g(M), M]}$ denote the time for a standard brownian motion (corresponding to $\sigma = 1$) to escape $[g(M), M]$, we have:

$$r \int_{g(M)}^M G_{g(M), M}^r(X, y) 2dx = r \mathbf{E} \left(\int_0^{\mathbb{T}_{[g(M), M]}} e^{-rt} dt | X \right).$$

Therefore,

$$\int_{g(M)}^M \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} r \mathbf{E} \left(\int_0^{\mathbb{T}_{[g(M), M]}} e^{-rt} dt | y \right) dy = \int_{g(M)}^M \left(\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)} \right) dX.$$

The integrals are equal for every continuous $g(M)$ and every $c(\sigma)$ if and only if they are pointwise

equal. Changing the integrand from y to X on the left-hand side yields:

$$\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} r \mathbf{E} \left(\int_0^{\mathbb{T}_{[g(M), M]}} e^{-rt} dt | X \right) = \left(\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)} \right).$$

Rearranging the terms produces:

$$\frac{2c(\sigma(M, X))}{c'(\sigma(M, X))} \left(1 - r \mathbf{E} \left(\int_0^{\mathbb{T}_{[g(M), M]}} e^{-rt} dt | X \right) \right) = \sigma(M, X).$$

Integration generates $r \mathbf{E} \left(\int_0^{\mathbb{T}_{[g(M), M]}} e^{-rt} dt | X \right) = 1 - e^{-r \mathbb{T}_{[g(M), M]}}$, which simplifies the formula:

$$\frac{2c(\sigma(M, X))}{c'(\sigma(M, X))} \mathbf{E}(e^{-r \mathbb{T}_{[g(M), M]}} | X) = \sigma(M, X).$$

Finally, the explicit expression for $\mathbf{E}(e^{-r \mathbb{T}_{[g(M), M]}} | X)$ can be found in [Borodin and Salminen \(2012\)](#), formula 3.0.1 (page 172):

$$\mathbf{E}(e^{-r \mathbb{T}_{[g(M), M]}} | X) = \frac{\sinh((M - X)\sqrt{2r}) + \sinh(X - g(M)\sqrt{2r})}{\sinh((M - g(M))\sqrt{2r})} = \frac{\cosh((M + g(M) - 2X)\sqrt{r/2})}{\cosh((M - g(M))\sqrt{r/2})}.$$

Putting this back into the previous formula yields:

$$\frac{2c(\sigma(M, X))}{c'(\sigma(M, X))} \frac{\sinh((M - X)\sqrt{2r}) + \sinh(X - g(M)\sqrt{2r})}{\sinh((M - g(M))\sqrt{2r})} = \sigma(M, X).$$

A direct implication of this derivation is that $\sigma(M, X)$ is symmetric around $(M + g(M))/2$. Furthermore, $\cosh$ is minimized at 0, so the multiplier is 1 at the boundaries, and equals $1/\cosh(d\sqrt{r/2})$ at the midpoint. It increases from the midpoint to the boundaries. Finally, $\lim_{r \rightarrow 0} \mathbf{E}(e^{-r \mathbb{T}_{[g(M), M]}} | X) = 1$, which coincides with the values at the boundaries. Therefore, $\sigma^r(M, M) = \sigma^r(M, g(M)) = \sigma^0(M, X) = \sigma^0$ is constant, and solves $\frac{2c(\sigma^0)}{c'(\sigma^0)} = \sigma^0$. ■

Proof of Proposition 2 and Corollary 2: The proof of Proposition 2 and Corollary 2 follows from the following lemma.

Lemma A2: *The optimal stopping boundary solves the ordinary differential equation (ODE):*

$$\begin{aligned}
& (g'(M) - 1) \left(\int_{g(M)}^M \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\sinh(\sqrt{2}\sqrt{r}(y - g(M)))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} dy \right) \\
& + (M - g(M)) \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\cosh(\sqrt{2}\sqrt{r}(M - g(M)))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} \frac{\sinh(\sqrt{2}\sqrt{r}(y - g(M)))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} dy \\
& + (M - g(M))g'(M) \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\sinh(\sqrt{2}\sqrt{r}(M - y))}{\sinh^2(\sqrt{2}\sqrt{r}(M - g(M)))} dy \\
& - (M - g(M)) \frac{2c(\sigma(M, M))}{\sigma(M, M)^2} + u'(M) = 0.
\end{aligned} \tag{9}$$

Proof of Lemma A2: In order to calculate the optimal stopping boundary, we differentiate equation (7) with respect to M , evaluate it at $X = M$, and set it equal to 0. Focus first on the term $\int_{g(M)}^M \frac{\sinh(\sqrt{2r}(M-y))}{\sinh(\sqrt{2r}(M-g(M)))} \frac{2c(\sigma(M,y))}{\sigma(M,y)^2} dy$. We can change variables, shifting to $\tilde{y} = M + g(M) - y$ and recalling that $\sigma(M, X)$ is symmetric around $(M + g(M))/2$. Thus, $\sigma(M, y) = \sigma(M, \tilde{y})$ within the integral range. We can then equivalently write equation (7), accounting for the sign change due to the variable replacement, as follows:

$$\begin{aligned}
V(M, X) = & u(M) - (X - g(M)) \int_{g(M)}^M \frac{\sinh(\sqrt{2r}(y - g(M)))}{\sinh(\sqrt{2r}(M - g(M)))} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy \\
& - \int_{g(M)}^M G_{g(M), M}^r(X, y) \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy.
\end{aligned} \tag{10}$$

Taking the derivative with respect to M and evaluating it at $X = M$ yields the following:

$$\begin{aligned}
& (g'(M) - 1) \left(\int_{g(M)}^M \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\sinh(\sqrt{2}\sqrt{r}(y - g(M)))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} \sinh(\sqrt{2}\sqrt{r}(y - g(M))) dy \right) \\
& + (M - g(M)) \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\cosh(\sqrt{2}\sqrt{r}(M - g(M)))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} \frac{\sinh(\sqrt{2}\sqrt{r}(y - g(M)))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} dy \\
& + (M - g(M))g'(M) \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\sinh(\sqrt{2}\sqrt{r}(M - y))}{\sinh^2(\sqrt{2}\sqrt{r}(M - g(M)))} dy \\
& - (M - g(M)) \frac{2c(\sigma(M, M))}{\sigma(M, M)^2} + u'(M) = 0.
\end{aligned} \tag{11}$$

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. ■

The ODE in Lemma A2 is non-linear. In general, such ODEs are not straightforward to solve in

analytical form, even when $r = 0$, due to the presence of $u'(M)$.²⁸ Nonetheless, in the Supplementary Materials, we illustrate this ODE's reduction to an alternative ODE that is more amenable to various classes of utilities studied in the economics literature when $r = 0$, as well as its application to the case of CRRA utilities, where stopping boundaries can be characterized analytically. In what follows, we solve this ODE directly for linear utilities. Set $u(M) = M$, the case analyzed in the text.

First, observe that we must have $M - g^r(M) = d^r$ for some d^r . Towards a contradiction, suppose this is not the case, so that there exist some $M, A \in \mathbb{R}$ such that $M - g^r(M) > M + A - g^r(M + A)$. From equation (6), we must have $V(M + A, M + A) - V(M, M) \neq A$. Now, if $V(M + A, M + A) - V(M, M) > A$, then $g^r(M)$ is suboptimal; if $V(M + A, M + A) - V(M, M) < A$, then $g^r(M + A)$ is suboptimal. We therefore get a contradiction.

Suppressing the r superscript for notational convenience, and using the fact that $g^r(M)$ is linear in M (with a derivative of 1), we can simplify equation (11) to:

$$\begin{aligned} & + (M - g(M)) \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\cosh(\sqrt{2}\sqrt{r}(M - g(M)))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} \frac{\sinh(\sqrt{2}\sqrt{r}(y - g(M)))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} dy \\ & + (M - g(M)) \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\sinh(\sqrt{2}\sqrt{r}(M - y))}{\sinh^2(\sqrt{2}\sqrt{r}(M - g(M)))} dy \\ & - (M - g(M)) \frac{2c(\sigma(\hat{M}, M))}{\sigma(\hat{M}, M)^2} + 1 = 0. \end{aligned}$$

Similar to the derivation of equation (10), we can introduce a change of variable for the first term:

$\tilde{y} = M + g(M) - y$ and use the identity $\frac{1}{\sinh(x)} - \frac{\cosh(x)}{\sinh(x)} = -\frac{\sinh(x/2)}{\cosh(x/2)}$, yielding:

$$\begin{aligned} & - (M - g(M)) \frac{\sinh(\sqrt{r/2}(M - g(M)))}{\cosh(\sqrt{r/2}(M - g(M)))} \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\sinh(\sqrt{2}\sqrt{r}(M - y))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} dy \\ & - (M - g(M)) \frac{2c(\sigma(M, M))}{\sigma(M, M)^2} + 1 = 0. \end{aligned} \tag{12}$$

Now, recall that the continuation HJB is:

$$\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)} = \int_{g(M)}^M 2r G_{g(M), M}^r(X, y) \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy.$$

We can divide both sides by $X - g(M)$ and take the limit as X goes to $g(M)$. The right-hand side

²⁸Peskir (1998) identified an equivalent ODE for the case of search without control and without discounting. 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.”

converges to:

$$\sqrt{2r} \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\sinh(\sqrt{2}\sqrt{r}(M - y))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} dy.$$

For the left-hand side, we use L'Hopital's rule to identify its limit as:

$$\left(\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)} \right) \frac{-\sigma_X(M, X)}{\sigma(M, X)} - \frac{\sigma_X(M, X)}{\sigma(M, X)} c''(\sigma(M, X)) \Big|_{X=g(M)}.$$

The first term in the parentheses equals 0 at $X = g(M)$ by Proposition 1. To calculate $\sigma_X(M, X)$, we use the identity

$$\frac{2c(\sigma(M, X))}{c'(\sigma(M, X))} \mathbf{E}(e^{-r\mathbb{T}_{[g(M), M]}} | X) = \sigma(M, X).$$

Let $\mathfrak{T}(X) = \mathbf{E}(e^{-r\mathbb{T}_{[g(M), M]}} | X)$. Taking the derivative of both sides with respect to X and simplifying yields:

$$\frac{2c(\sigma(M, X))c'(\sigma(M, X))\mathfrak{T}'(X)}{c'(\sigma(M, X))^2(1 - \mathfrak{T}(X)) + c(\sigma(M, X))c''(\sigma(M, X))\mathfrak{T}(X)} = \sigma_X(M, X).$$

By definition, $\mathfrak{T}(g(M)) = 1$ and $\mathfrak{T}'(g(M)) = -\sqrt{2r} \sinh(\sqrt{r/2}(M - g(M)))$. Plugging these back in and simplifying further yields:

$$\begin{aligned} & \frac{2c'(\sigma(M, g(M)))}{\sigma(M, g(M))} \sinh(\sqrt{r/2}(M - g(M))) \\ &= \int_{g(M)}^M \sqrt{2}\sqrt{r} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} \frac{\sinh(\sqrt{2}\sqrt{r}(M - y))}{\sinh(\sqrt{2}\sqrt{r}(M - g(M)))} dy \end{aligned}$$

Plugging this into equation (12),

$$\begin{aligned} & - (M - g(M)) \frac{\sinh(\sqrt{r/2}(M - g(M)))}{\cosh(\sqrt{r/2}(M - g(M)))} \frac{2c'(\sigma(M, g(M)))}{\sigma(M, g(M))} \sinh(\sqrt{r/2}(M - g(M))) \\ & - (M - g(M)) \frac{2c(\sigma(M, M))}{\sigma(M, M)^2} + 1 = 0. \end{aligned}$$

As r approaches 0, the hyperbolic terms vanish and we have, $M - g^0(M) = d^0 = \frac{(\sigma^0)^2}{2c(\sigma^0)}$. For positive r , we can use the fact that $\sigma^r(M, g(M)) = \sigma^0$ to identify the drawdown size $d^r = M - g^r(M)$:

$$\begin{aligned} & \frac{d^r}{d^0} \left(1 + 2 \frac{\sinh(\sqrt{r/2}d^r)}{\cosh(\sqrt{r/2}d^r)} d^r \sinh(\sqrt{r/2}d^r) \right) \\ &= \frac{d^r}{d^0} \left(1 + 2 \frac{\sinh^2(\sqrt{r/2}d^r)}{\cosh(\sqrt{r/2}d^r)} \right) = 1. \end{aligned}$$

Since both $\sinh$ and $\tanh$ are increasing in r , this implies that the optimal drawdown $d^r = M - g(M)$ is decreasing in r . ■

Proof of Proposition 3: Recall that when $u(M) = M$, equation (5) is given by:

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

From the proof of proposition 2, we know that $\int_{g(M)}^M \frac{\sinh(\sqrt{2r}(M - y))}{\sinh(\sqrt{2r}(M - g(M)))} \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy = \frac{2c'(\sigma(M, g(M)))}{\sigma(M, g(M))} \frac{\sinh(\sqrt{r/2}(M - g(M)))}{\sqrt{2r}}$. Plugging this in, the value function can be written as:

$$V(\hat{M}, \hat{X}) = M + (X - g(M)) \frac{2c'(\sigma(M, g(M)))}{\sigma(M, g(M))} \frac{\sinh(\sqrt{r/2}(M - g(M)))}{\sqrt{2r}} \\ - \int_{g(M)}^M G_{g(M), M}^r(X, y) \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy.$$

Similarly, from equation (8), we have

$$2r \int_{g(M)}^M G_{g(M), M}^r(X, y) \frac{2c(\sigma(M, y))}{\sigma(M, y)^2} dy = \frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)}.$$

Therefore, we can write the value function as:

$$V(M, X) = M + (X - g(M)) \frac{2c'(\sigma(M, g(M)))}{\sigma(M, g(M))} \frac{\sinh(\sqrt{r/2}(M - g(M)))}{\sqrt{2r}} - \frac{\frac{2c(\sigma(M, X))}{\sigma(M, X)^2} - \frac{c'(\sigma(M, X))}{\sigma(M, X)}}{2r}.$$

Recall that $\sigma(M, M) = \sigma(M, g(M)) = \sigma^0$, where σ^0 is the undiscounted optimal search scope. From the HJB, we know that $\frac{2c(\sigma(0, 0))}{\sigma(0, 0)^2} - \frac{c'(\sigma(0, 0))}{\sigma(0, 0)} = 0$. Thus,

$$V(0, 0) = (0 - g(0)) \frac{2c'(\sigma(0, g(0)))}{\sigma(0, g(0))} \frac{\sinh(\sqrt{r/2}(0 - g(0)))}{\sqrt{2r}}.$$

In general, we can write the above in terms of the drawdowns by noting that $d^r = 0 - g(0)$:

$$V(0, 0) = \frac{d^r}{d^0} \frac{\sinh(\sqrt{r/2}d^r)}{\sqrt{2r}}.$$
■

Proof of Corollary 3: From Proposition 1, σ^0 is constant and, from Proposition 2, the optimal stopping time is a drawdown stopping time with drawdown size $d^0 = \frac{(\sigma^0)^2}{2c(\sigma^0)}$. Recall that the optimal

stopping time is:

$$\tau^* = \inf \{t \geq 0 : M_t - X_t \geq d^0\}.$$

Since σ^0 is constant, from Taylor et al. (1975), the joint moment generating function and Laplace transform of X_{τ^*} and τ^* is given by:

$$\mathbf{E}[e^{X_{\tau^*} - c(\sigma^0)\tau^*}] = \frac{\beta e^{-d^0}}{\beta \cosh(\beta d^0) - \sinh(\beta)},$$

where $\beta = \sqrt{2c(\sigma^0)/(\sigma^0)^2}$.

The characterization of the distribution of M_{τ^*} then follows since $M_{\tau^*} = X_{\tau^*} + \frac{\sigma^0}{2c(\sigma^0)}$. 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, in addition to $V(M, X)$ that can be directly calculated via Proposition 3, we have

$$\mathbf{E}(\tau^*) = \frac{(d^0)^2}{(\sigma^0)^2} \text{ and } \mathbf{E}(M_{\tau^*}) = d^0 \text{ and } V(M, X) = d^0/2.$$

Furthermore, the distribution of M_{τ^*} is a standard exponential distribution with mean d . 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 ,

$$\mathbf{E}(M_{\tau^*}) = \mathbf{E}(M_{\tau^*} | \tau^* = t_1) = \mathbf{E}(M_{\tau^*} | \tau^* = t_2) = d^0.$$

■

Proof of Corollary 4: The proof is analogous to the proofs of Propositions 1 and 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

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

The stopping boundary is again linear and similar analysis to that in Claim 3 yields:

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

■

Proof of Proposition 4: The expressions describing $\mathbf{E}(M_{\tau})$ and $\mathbf{E}(\tau)$ follow directly from the proof of Corollary 4, 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 immediately. ■

Proof of Proposition 5: The principal's problem, taking the agent's solution from Corollary 4 as given, can be written as:

$$\begin{aligned} \max_{\alpha, w} & \frac{(1-\alpha)\alpha\sigma^{*2}}{2(c(\sigma^*)-w)} - w \left(\frac{\alpha\sigma^*}{2(c(\sigma^*)-w)} \right)^2 \\ \text{subject to} & \\ \sigma^* &= \frac{2(c(\sigma^*)-w)}{c'(\sigma^*)}. \end{aligned}$$

The optimal search scope 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 principal's problem 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:

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

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:

$$f_{11} \equiv -\frac{2(c(\sigma) - \frac{1}{2}\sigma c'(\sigma))}{c'^2} - \frac{2\sigma}{c(\sigma)},$$

$$f_{22} \equiv \frac{1}{2}\alpha \left(-\frac{4(\alpha-1)\sigma c'^2}{c(\sigma)^3} + \frac{2(\alpha-1)(\sigma c''(\sigma) + 2c'(\sigma))}{c(\sigma)^2} \right. \\ \left. + \frac{4\alpha c(\sigma)(c'''(\sigma)c'(\sigma) - 3c''^2)}{c'^4} + \frac{\alpha(2\sigma c''^2 + c'(\sigma)(4c'''(\sigma) - \sigma c''''(\sigma)))}{c'^3} \right),$$

$$f_{12} \equiv \frac{(2\alpha-1)\sigma c'(\sigma)}{c(\sigma)^2} + \frac{4\alpha c(\sigma)c''(\sigma)}{c'^3} - \frac{\alpha(\sigma c''(\sigma) + c'(\sigma))}{c'^2} + \frac{1-2\alpha}{c(\sigma)}.$$

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 invalid, 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(4), 1185–1209.
- Bergemann, D. and J. Valimaki (2008). *Bandit problems*. The New Palgrave Dictionary of Economics, 2nd ed. Macmillan Press.
- Borodin, A. N. and P. Salminen (2012). *Handbook of Brownian motion-facts and formulae*. Birkhäuser.
- Callander, S. (2011). Searching and learning by trial and error. *American Economic Review* 101(6), 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(1), 44–78.
- Dubins, L. E., L. A. Shepp, and A. N. Shiryaev (1994). Optimal stopping rules and maximal inequalities for bessel processes. *Theory of Probability & Its Applications* 38(2), 226–261.
- Garfagnini, U. and B. Strulovici (2016). Social experimentation with interdependent and expanding technologies. *Review of Economic Studies* 83(4), 1579–1613.
- Guo, Y. (2016). Dynamic delegation of experimentation. *American Economic Review* 106(8), 1969–2008.
- Halac, M., N. Kartik, and Q. Liu (2016). Optimal contracts for experimentation. *The Review of Economic Studies* 83(3), 1040–1091.
- Hull, C. L. (1932). The goal-gradient hypothesis and maze learning. *Psychological review* 39(1), 25.
- Itô, K., P. Henry Jr, et al. (2012). *Diffusion processes and their sample paths*. Springer Science & Business Media.
- Jensen, R. and M. Thursby (2001). Proofs and prototypes for sale: The licensing of university inventions. *American Economic Review* 91(1), 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.
- Kivetz, R., O. Urminsky, and Y. Zheng (2006). The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *Journal of marketing research* 43(1), 39–58.
- Kortum, S. S. (1997). Research, patenting, and technological change. *Econometrica*, 1389–1419.
- Manso, G. (2011). Motivating innovation. *The Journal of Finance* 66(5), 1823–1860.
- McCall, B. and J. McCall (2007). *The economics of search*. Routledge.
- Miller, R. A. (1984). Job matching and occupational choice. *Journal of Political economy* 92(6), 1086–1120.
- Moscarini, G. and L. Smith (2001). The optimal level of experimentation. *Econometrica* 69(6), 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(4), 634–650.
- Peskir, G. (1998). Optimal stopping of the maximum process: The maximality principle. *Annals of Probability*, 1614–1640.
- Peskir, G. (2005). Maximum process problems in optimal control theory. *International Journal of Stochastic Analysis* 2005(1), 77–88.
- Peskir, G. and A. Shiryaev (2006). *Optimal stopping and free-boundary problems*. Springer.
- Pissarides, C. A. (1984). Search intensity, job advertising, and efficiency. *Journal of labor Economics* 2(1), 128–143.
- Pissarides, C. A. (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(2), 127–146.
- Revuz, D. and M. Yor (2013). *Continuous martingales and Brownian motion*, Volume 293. Springer Science & Business Media.
- Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* 58(5), 527–535.
- Rogers, L. C. G. and D. Williams (2000). *Diffusions, Markov processes and martingales: Volume 2, Itô calculus*, Volume 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(2), 234–246.
- Urgun, C. and L. Yariv (2021). Constrained retrospective search. *AEA Papers and Proceedings* 111, 549–53.
- U.S. Department of Agriculture, F. S. (1995). Anatomy of a mine from prospect to production. Technical report.
- 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*, Volume 43. Springer Science & Business Media.
- Zhilkin, N. (1961). Feasible new methods of prospecting and exploration. *International Geology Review* 3(1), 71–77.